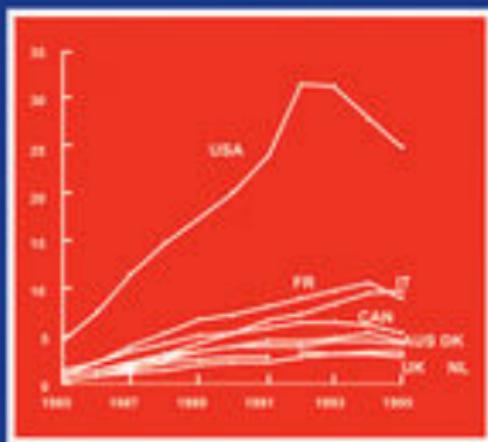


WILEY SERIES IN PROBABILITY AND STATISTICS

Empirical Model Building

Data, Models, and Reality

Second Edition



James R. Thompson



WILEY

Empirical Model Building

WILEY SERIES IN PROBABILITY AND STATISTICS

Established by **WALTER A. SHEWHART** and **SAMUEL S. WILKS**

Editors: *David J. Balding, Noel A. C. Cressie, Garrett M. Fitzmaurice, Harvey Goldstein, Iain M. Johnstone, Geert Molenberghs, David W. Scott, Adrian F. M. Smith, Ruey S. Tsay, Sanford Weisberg*
Editors Emeriti: *Vic Barnett, J. Stuart Hunter, Joseph B. Kadane, Jozef L. Teugels*

A complete list of the titles in this series appears at the end of this volume.

Empirical Model Building

Data, Models, and Reality

Second Edition

JAMES R. THOMPSON



A JOHN WILEY & SONS, INC., PUBLICATION

Copyright © 2011 by John Wiley & Sons, Inc. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representation or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print, however, may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Thompson, James R. (James Robert), 1938—

Empirical model building : data, models, and reality / James R. Thompson. — 2nd ed.

p. cm — (Wiley series in probability and statistics ; 794)

Includes bibliographical references.

ISBN 978-0-470-46703-9

1. Experimental design. 2. Mathematical models. 3. Mathematical statistics. I. Title.

QA279.T49 2011

519.5'7—dc23

2011013569

Printed in the United States of America

ePDF ISBN: 978-1-118-10962-5

eBook ISBN: 978-1-118-10965-6

10 9 8 7 6 5 4 3 2 1

*To: Marian Rejewski and Herman Kahn,
Master Empirical Model Builders*

Marian Rejewski cracked the Enigma Code in 1932,
sparing the United Kingdom strangulation by
German submarines in World War II.

The writings of Herman Kahn provided the basis for the
Reagan–Kohl Pershing II strategy which brought down the
“evil empire” without firing a shot.

Contents

Preface	xiii
1. Models of Growth and Decay	1
1.1. A Simple Pension and Annuity Plan	1
1.2. Income Tax Bracket Creep and the Quiet Revolution of 1980	8
1.3. Retirement of a Mortgage	16
1.4. Some Mathematical Descriptions of the Theory of Malthus	18
1.5. Metastasis and Resistance	29
Problems	37
References	38
2. Models of Competition, Survival, and Combat	39
2.1. An Analysis of the Demographics of Ancient Israel	39
2.2. The Plague and John Graunt's Life Table	44
2.3. Modular Data-Based Wargaming	48
2.3.1. Herman Kahn and the Winning of the Cold War	64
2.4. Predation and Immune Response Systems	65
2.5. Pyramid Clubs for Fun and Profit	69
Problems	72
References	74
3. Epidemics	75
3.1. Introduction	75
3.2. John Snow and the London Cholera Epidemic of 1854	77
3.3. Prelude: The Postwar Polio Epidemic	79
3.4. AIDS: A New Epidemic for America	80
3.5. Why an AIDS Epidemic in America?	82
3.5.1. Political Correctness Can Kill	84
3.6. The Effect of the Gay Bathhouses	85
3.7. A More Detailed Look at the Model	88
3.8. Forays into the Public Policy Arena	91
3.9. Modeling the Mature Epidemic	92

3.10. AIDS as a Facilitator of Other Epidemics	95
3.11. Comparisons with First World Countries	95
3.12. Conclusions: A Modeler's Portfolio	103
Problems	106
References	109
4. Bootstrapping	111
4.1. Introduction	111
4.2. Bootstrapping Analysis of Darwin's Data	112
4.3. A Bootstrapping Approximation to Fisher's Nonparametric Test	116
4.4. A Resampling Bassed Sign Test	119
4.5. A Bootstrapping Approach for Confidence Intervals	121
4.6. Solving Ill-Structured Problems	122
Problems	128
References	130
5. Monte Carlo Solutions of Differential Equations	131
5.1. Introduction	131
5.2. Gambler's Ruin	134
5.3. Solution of Simple Differential Equations	137
5.4. Solution of the Fokker-Planck Equation	138
5.5. The Dirichlet Problem	139
5.6. Solution of General Elliptic Differential Equations	142
5.7. Conclusions	144
Problems	145
References	146
6. SIMEST, SIMDAT, and Pseudoreality	147
6.1. Introduction	147
6.2. The Bootstrap: A Dirac-Comb Density Estimator	149
6.3. SIMDAT: A Smooth Resampling Algorithm	152
6.3.1. The SIMDAT Algorithm	153
6.3.2. An Empirical Justification of SIMDAT	153
6.4. SIMEST: An Oncological Example	156
6.4.1. An Exploratory Prelude	158
6.4.2. Models and Algorithms	158
Problems	170
References	172
7. Exploratory Data Analysis	175
7.1. Introduction	175
7.2. Smoothing	176
7.3. The Stem and Leaf Plot	187
7.4. The Five Figure Summary	188

7.5. Tukey's Box Plot	189
Problems	192
References	192
8. Noise Killing Chaos	193
8.1. Introduction	193
8.2. The Discrete Logistic Model	196
8.3. A Chaotic Convection Model	201
8.4. Conclusions	204
Problems	208
References	208
9. A Primer in Bayesian Data Analysis	209
9.1. Introduction	209
9.2. The EM Algorithm	210
9.3. The Data Augmentation Algorithm	213
9.4. The Gibbs Sampler	219
9.5. Conclusions	224
Problems	225
References	228
10. Multivariate and Robust Procedures in Statistical Process Control	229
10.1. Introduction	229
10.2. A Contamination Model for SPC	231
10.3. A Compound Test for Higher Dimensional SPC Data	236
10.4. Rank Testing with Higher Dimensional SPC Data	239
10.5. A Robust Estimation Procedure for Location in Higher Dimensions	243
Problems	246
References	248
11. Considerations for Optimization and Estimation in the Real (Noisy) World	249
11.1. Introduction	249
11.2. The Nelder–Mead Algorithm	253
11.3. The Box–Hunter Algorithm	256
11.4. Simulated Annealing	262
11.5. Exploration and Estimation in High Dimensions	264
Problems	270
References	271
12. Utility and Group Preference	273
12.1. Introduction	273
12.2. The St. Petersburg Paradox	274

12.3. von Neumann–Morgenstern Utility	282
12.4. Creating a St. Petersburg Trust	287
12.5. Some Problems with Aggregate Choice Behavior	290
12.6. Jeffersonian Realities	293
12.7. Conclusions	300
Problems	300
References	303
13. A Primer in Sampling	305
13.1. Introduction	305
13.1.1. Tracking Polls	308
13.2. Stratification	309
13.2.1. A Warehouse Inventory	311
13.3. The Saga of Happy Valley	312
13.3.1. The Problem with Standards	313
13.4. To Annex or Not	316
13.4.1. The Bootstrap: The World in a Drop of Water	316
13.4.2. Resampling and Histograms	317
13.5. Using Sampling to Estimate Total Population Size	321
13.5.1. Capture Recapture	322
Problems	324
References	325
14. Stock Market: Strategies Based on Data versus Strategies Based on Ideology	327
14.1. Introduction	327
14.2. Markowitz's Efficient Frontier: Portfolio Design as Constrained Optimization	329
14.3. Sharpe's Super Efficient Frontier: The Capital Market Line (CAPM)	333
14.4. The Security Market Line	335
14.5. The Sharpe Diagonal Model	337
14.6. Portfolio Evaluation and the Capital Asset Pricing Model	339
14.7. Views of Risk	342
14.7.1. Diversification	344
14.8. Stock Progression as Geometric Brownian Motion	344
14.8.1. Ito's Lemma	344
14.8.2. A Geometric Brownian Model for Stocks	346
14.9. Estimating μ and σ	346
14.10. The Time Indexed Distribution of Portfolio Value	348
14.11. Negatively Correlated Portfolios	351
14.12. Bear Jumps	354
14.13. "Everyman's" MaxMedian Rule for Portfolio Management	357
14.13.1. Investing in a 401k	360

14.14. Derivatives	363
14.14.1. Black-Scholes and the Search for Risk Neutrality	364
14.14.2. A Game Independent of the Odds	65
14.14.3. The Black-Scholes Hedge Using Binomial Trees	368
14.15. The Black-Scholes Derivation Using Differential Equations	374
14.16. Black-Scholes: Some Limiting Cases	376
14.17. Conclusions	381
Problems	382
References	384
Appendix A. A Brief Introduction to Probability and Statistics	387
A.1. Craps: An Intuitive Introduction to Probability and Statistics	387
A.1.1. Random Variables, Their Means and Variances	387
A.2. Combinatorics Basics	393
A.3. Bayesian Statistics	394
A.3.1. Bayes' Theorem	394
A.3.2. A Diagnostic Example	396
A.4. The Binomial Distribution	397
A.5. The Uniform Distribution	399
A.6. Moment Generating Functions	401
A.7. The Normal (Gaussian) Distribution	402
A.8. The Central Limit Theorem	405
A.9. The Gamma Distribution	406
A.10. Conditional Density Functions	407
A.11. The Weak Law of Large Numbers	409
A.12. The Multivariate Normal Distribution	411
A.13. The Wiener Process	411
A.14. The Poisson Process and the Poisson Distribution	412
A.14.1. Simulating Bear Jumps	414
A.15. Parametric Simulation	415
A.15.1. Simulating a Geometric Brownian Walk	416
A.15.2. The Multivariate Case	417
A.16. Resampling Simulation	420
A.16.1. The Multivariate Case	421
A.17. A Portfolio Case Study	422
Appendix B. Statistical Tables	425
Tables of the Normal Distribution	425
Tables of the Chi-Square Distribution	426
Index	427

Preface

That mathematics is not a science because it exists in a Platonic world of abstraction was well argued by the late John W. Tukey. Statistics, on the other hand, deals with data from the real world. Hence statistics can be claimed to be a science, to the extent that its practitioners focus on data analysis including model inferences based on data. Many scholars have believed that Tukey (himself a brilliant topologist) made a mistake in taking statistics out from under the umbrella of mathematics. Indeed, some members of the departments of statistics seldom, if ever, look at a set of data, contenting themselves with elegant data-free mathematical structures. Many a named chair has been awarded to a “statistician” who contented himself/herself with developing tests (seldom used in actual data analysis), and then proceeding to prove the asymptotic optimality of such tests under idealized conditions.

The father of Exploratory Data Analysis, John Tukey, took the position that individuals who avoid data from the real world, be they ever so elegant mathematically, are not practicing statistics but mathematics. Tukey went further and argued that those who did only apply to data standardized tests of models of other scientists were confining themselves unnecessarily to “confirmatory data analysis.” He wanted statisticians to be more original than that. It was better if data analysis were done in an exploratory fashion. In other words, he wanted statisticians to be key players in science. They should examine data without relying too much on existing theories, and try to make inferences about real world systems.

Statistics is older than mathematics. Moses did carry out a census of the number of warriors in the Jewish tribes in 1452 BC. He made inferences from his census about the logistics of caring for his people and the military conquest of Canaan. Furthermore, I would submit that Thucydides, who wrote a fairly objective account of the Peloponnesian War between 431 BC and 411 BC (the war continued until 404 BC, but Thucydides apparently did not live to see the end of it), should be considered the father of time series analysis, an important subcategory of statistics.

Centuries later, geostrategist Herman Kahn (who was a major player in Monte Carlo simulation) argued from historical patterns and extrapolations how best to overcome the Soviet threat to the Free World. He was going further than Thucydides in that he was not only talking about qualitative facts in the past, but was extrapolating into an unknown future what would be the likely results of various strategies first to contain and then to destroy politically the Soviet Union. In other words, he

engaged in Extrapolatory Data Analysis. The Reagan–Kohl Pershing II strategy was one almost taken right out of the pages of Kahn’s numerous books.

Yet many scholars would argue that neither Thucydides nor Kahn could be considered statisticians.

In my own professional experience, I have had the good fortune of working in the building of practical models in such fields as oncology, tall building construction, manufacturing, epidemiology, military strategy, remote sensing, public policy, market analysis, nonparametric data-based estimation of functions, and others. Seldom have I been able to sit inside the box of standardized theories, statistical or otherwise. I always had to climb outside the box and use data to build appropriate models for dealing with the problem at hand.

The purpose of this book is not only to add to the arsenal of tools for using data to build appropriate models, but to give the reader insights as to how he/she can take his/her own direction for empirical model building. From my own experience in obtaining a bachelor’s degree in engineering from Vanderbilt and a doctorate in mathematics from Princeton, I can say that I seldom use directly the tools learned in my university education. But it is also true that an understanding of those mechanisms has been of significant help in my career in data analysis and model building.

In Chapter 1, I consider topics in growth, including population growth, tax policy, and modeling tumor growth and its chemotherapeutic control. Parts of this chapter are simple from a mathematical standpoint, but the entire chapter has important implications for optimization and control. The so-called Malthusian theory is an example of a model that makes logical sense, absent data, but is seriously flawed when compared with reality. It is also noted that Malthus’s ideas led to Social Darwinism and the unnecessary horror of the Irish potato famine.

Chapter 2 starts with an attempt to use data from the Old Testament to make inferences about the growth of the Jewish population starting with 1750 BC and ending with the fall of the Kingdom of David in 975 BC. Then we examine John Graunt’s creation of the life table analysis of the demographic effects of the London plague of the sixteenth and seventeenth centuries. Then we delve into the data-based combat modeling started by Georg von Reisswitz and carried through with great success by Lanchester. Although easily computerized and data-modifiable, the U.S. Department of Defense, since the time of Robert McNamara has opted for non-data-based, noninteractive computer games such as Castforem. The work of the late Monte Carlo innovator, Herman Kahn, working without federal support in creating the strategy used by Reagan and Kohl to bring down the USSR, is also discussed.

Models for coping with contagious epidemics are presented in Chapter 3. We start with the laws of Moses concerning leprosy. We then proceed to John Snow’s halting of the East End cholera epidemic of the mid-1800s. Each of these epidemics was more or less successfully dealt with sociologically. We then suggest that the AIDS epidemic that has killed over 600,000 Americans could have been prevented in the United States if only the simple expedient of closing the gay bathhouses had been followed. We note that sociological control of epidemics should not be neglected. The United States has much the highest number of AIDS cases per hundred thousand in the First World (over ten times that in the United Kingdom). The failure

of the Centers for Disease Control to shut down the gay bathhouses is shown to be a plausible explanation of why AIDS continues to thrive in the United States. An argument is made that the United States may be “country zero” for the First World epidemic.

Chapter 4 deals with the bootstrap work of Julian Simon, Bradley Efron, and Peter Bruce. This computer-intensive resampling technique has revolutionized statistics. The classic zea mays data set from Fisher’s *The Design of Experiments* is reanalyzed using various bootstrapping techniques. Then we deal with some rather unusual problems successfully handled by bootstrap techniques.

In Chapter 5 we show the importance of simulation in solving differential equations which do not admit of closed-form solutions (i.e., most of them). Particularly for partial differential equations in dimensions higher than three, simulation becomes virtually our only recourse.

I have often been approached by clients who wanted me to increase the size of their data sets by statistical means. That is generally not doable. However, the SIMDAT algorithm can build a continuous nonparametric density estimation base of pseudo-data around and within the neighborhood of an existing set which avoids such anomalies as suggesting that ammunition be stored in a section of a tank where there was no data but is shown from the nonparametric density estimator approach to be very vulnerable indeed. For many problems it is easy to write down the plausible axioms which have generated a data set. However, it is rarely the case that these axioms lead to a ready representation of the likelihood function. The problem is that the axioms are written in the forward direction, but the likelihood requires a backward look. The SIMEST algorithm allows a temporally forward approach for dealing with the estimation of the underlying parameters. SIMDAT and SIMEST are developed in Chapter 6.

Chapter 7 is a brief survey of the exploratory data analysis paradigm of John Tukey. He viewed statistics not just as a device by which models developed by non-statisticians could be confirmed or rejected on the basis of data. He wanted statisticians to be on the cutting edge of discovery. He noted that exploration of data could be used to infer structures and effect inferences and extrapolations. EDA has greatly expanded the creativity horizons for statisticians as generalists across the scientific spectrum.

Chapter 8 is devoted to what I consider to be shortlived fads. Fuzzy logic and catastrophe theory have been shown to be inadequate tools. Chaos theory is not such a hot topic as formerly, but it still has its followers. The fads tend to build anti-Aristotelian structures not really sensitive to data. If we build a mathematical model that is not to be stressed by data, then we have entered the realm of postmodernism where everyone gets to create his/her own reality. Simply showing a mathematical structure does not indicate that structure conforms to anything in the real world. Nevertheless, I show that even if one does look at some of the artificial models of chaos theory, the addition of a slight amount of noise can frequently bring the chaotic model to something which does conform to real-world data bases.

Some professors of statistics believe that Bayesian data analysis is the only way to go. Bayesian theory has a lot to be said for it. For example, it gets around the

claim of Karl Popper that statistics can only demolish hypotheses, never confirm them. Over the years, the use of noninformative prior distributions has seemingly weakened a major *raison d'être* of Bayesian analysis. In Chapter 9, the author attempts to give a data-based exposition of Bayesian theory, including the EM algorithm, data augmentation, and the Gibbs sampler.

There used to be surveys performed to decide who the most important living statistician was. Sometimes John Tukey would come in first. At other times it would be Edward Deming, the developer of Statistical Process Control. In Chapter 10 we go beyond the normal low-dimensional analysis advocated by Deming to show how higher-dimensional control charts can be constructed and how nonparametric as well as parametric tests can be used.

In Chapter 11 we investigate procedures where optimization may be readily implemented in the real world, which is generally noisy. We particularly emphasize algorithms developed by the statisticians Nelder and Mead which are amazingly useful in this age of swift computing. Lawera and Thompson built upon the piecewise quadratic optimization technique used in the rotatable experimental designs of the statisticians Box and Hunter.

Chapter 12 shows how no lesser a person than the author of *The Declaration of Independence*, Thomas Jefferson, persuaded President Washington to veto the rather reasonable and transparent allocation rule of Alexander Hamilton for the allocation of congressmen across the various states in favor of Jefferson's rule favoring the large population states (at that time Jefferson's Virginia was the largest). This first use of the Presidential veto is simply an example of the fact (pointed out by Jefferson earlier) that one should take the statements of politicians with a very large grain of salt. We show the basis of the utility theory of Bernoulli as well as that of Morgenstern and von Neumann. Finally, we present the Nobel Prize winning Impossibility Theorem of Kenneth Arrow which demonstrates the fact that group decisions which make everybody happy can never be constructed.

Chapter 13 is a brief practicum in sampling theory. The author has experience consulting in this area and believes that a great deal can be achieved by the use of transparent and relatively simple strategies.

In Chapter 14 it is shown how efficient market theory including capital market theory and the derivative approaches which proceed from the Nobel Prize winning work of Black, Scholes, and Merton are inconsistent with market data. The efficient market hypothesis has dominated finance, as taught in schools of business, for decades, much to the disadvantage of investors and society as a whole. As an alternative, it is demonstrated how computer-intensive and momentum-based strategies may be created which significantly best the market cap Index Fund strategies that proceed from capital market theory. Serious doubt is cast on the practice of the vending of uncovered call options. Most importantly, this chapter attempts to show young investors how they can develop their own strategies for purchasing stocks in an age where wars of choice (based on information later declared to be false), bad Federal Reserve policy, and the financing of houses to persons unable to pay off the mortgages have produced conditions where inflation becomes almost a certainty. The author has no magic rule for making the reader rich, but he gives the kind of in-

formation which is assuredly useful in showing him/her how to plan for using the market as a vehicle for obtaining a measure of financial security.

This work was supported in part by the Army Research Office (Durham) (W911NF-04-1-0354) *Some Topics in Applied Stochastic Modeling, Risk Analysis and Computer Simulation*. I would like to thank my mentor John Tukey and my colleagues Katherine Ensor, Scott Baggett, John Dobelman, Ed Williams, David Scott, Chapman Findlay, Jacek Koronacki, Webster West, Martin Lawera, Marc Elliott, Otto Schwalb, William Wojciechowski, Steven Boswell, Rachel MacKenzie, Neely Atkinson, Barry Brown and Ricardo Affinito. At John Wiley & Sons, I would like to thank my editor, Stephen Quigley.

Finally and most importantly, I wish to thank my wife, Professor Ewa Thompson, for her continuing love and encouragement.

JAMES THOMPSON

*Houston, Texas
Easter 2011*

Chapter 1

Models of Growth and Decay

1.1 A Simple Pension and Annuity Plan

It is almost always the case that everything we plan for is, in actuality, an approximation to what might happen. The distinction between “science” and “engineering” frequently has to do with the difference between models of Newtonian precision and approximations to a reality only partially understood. The fact is, of course, that Newton’s laws are themselves approximations. It is true that we can much more accurately design an automobile than we can plan an economy. However, though the author absolutely believes in objective reality, he understands that he is unlikely to find it in any physical or social system. As St. Paul said, “we see through a glass darkly,” (1 Corinthians 13:12, King James Version). Of course some visions are cloudier than others. But, truth be told, if science is the study of precise reality and engineering is the study of approximations to that reality, then almost every scientist is actually an engineer. Frequently, physical scientists look down on the social sciences because the models in the physical sciences are more precise and more accurate. But, in reality, we are most of us empirical modelers doing the best we can, making logical inferences based on data, to understand what is going on and what will happen.

One easy introduction to the subject of growth models is obtained by considering an accounting situation, because such systems are relatively well defined. A pension and annuity plan is generally axiomatized clearly.

As a practical matter, pension funds in the United States are becoming much less generous than was the case in years past. A major reason for the distress of companies such as General Motors was the large liability built up

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

over years to provide pensions and health care plans for their workers. For many small companies, there now simply are no pension funds beyond the federally mandated Social Security. A company is obliged to pay slightly more than 7% of a worker's salary into Social Security. The worker must match this. Insofar as the company's contribution is concerned, it is tax deductible. The worker's portion is not tax deductible. So an average worker is faced with the necessity of a "payroll tax" for income that is put somewhere in the maze of federal funds and is frequently spent out as though it were ready cash money. This is a kind of double whammy. In a sense, the worker is taxed on money that she does not receive. Then, when she starts to collect Social Security, a middle class employee is taxed a third time on 85% of monies received.

Typically, many employers are very concerned about the welfare of their workers. We recall that in 1908 Henry Ford instituted the 40-hour week and the minimum daily wage of \$5. He also started at the same time a profit-sharing plan for his workers. He provided, at low mortgage rates, simple two-bedroom houses for his workers. We recall that this was 100 years ago when the majority of workers were not so blessed with so kindly an employer.

Unions tended to develop adversarial attitudes toward managers, whom they felt cared little about the welfare of the workers. Wage structures and benefit plans began to be increasingly divorced from economic reality. A plan that deferred some wage benefits into future retirement funds might be very attractive to a manager who was making promises that he would not have to meet in his own professional lifetime.

We will consider below a possible minimal pension fund. It is referred to as an Individual Retirement Account (IRA). Properly structured, the contribution of the worker and that of the employer are both tax deductible. At the time of retirement, the worker will receive monthly payments according to a mutually agreed upon plan. On these payments, he or she will pay taxes as though the annuity payments were ordinary income.

Suppose that we set up a pension plan for an individual who starts working for a firm at age N_1 , retiring at age N_2 . The starting salary is S_1 at age N_1 and will increase at rate α annually. The employer and employee both contribute a fraction $\beta/2$ of the salary each year to the employee's pension fund. The fund is invested at a fixed annual rate of return γ . We wish to find the value of the employee's pension fund at retirement. Furthermore, we wish to know the size of the employee's pension checks if he invests his pension capital at retirement in a life annuity (i.e., one that pays only during the employee's lifetime, with no benefit to his heirs at time of death). Let the expected death age given survival until N_2 be denoted by N_3 .

Many investigators find it convenient to consider first a pilot study with concrete values instead of algebraic symbols. For example, we might try $S_1 = \$40,000$; $N_1 = 21$; $N_2 = 65$; $\alpha = 0.02$; $\beta = 0.0705$; $\gamma = 0.05$; $N_3 = 75$. The values of α and γ are rather low. The value of β is roughly half the

value used at present in the U.S. Social Security System. The value of N_2 is the same as the present regular Social Security retirement age of 65. No allowance is made for annual losses from taxation, since pension plans in the United States generally leave the deposited capital untaxed until the employee begins his or her annuity payments (although the employee contributions to Social Security are taxed at full rate).

First, we note that at the end of the first year the employee will have approximately

$$P(1) = \beta S(1) = (0.0705) \$40,000 = \$2820 \quad (1.1)$$

invested in the pension plan. This will only be an approximation, because most salaried employees have their pension fund increments invested monthly rather than at the end of the year. We shall use the approximation for the pilot study.

At the end of the second year, the employee will have approximately

$$P(2) = \beta S(2) + (1 + \gamma) P(1) = (0.0705) S(2) + 1.05 (\$2820), \quad (1.2)$$

where

$$S(2) = (1 + \alpha) S(1) = 1.02 S(1) = 1.02 (\$20,000) = \$20,400.$$

Thus, we have that $P(2) = \$5837$. Similarly,

$$P(3) = \beta S(3) + (1 + \gamma) P(2) = (0.0705) S(3) + 1.05 (\$5837), \quad (1.3)$$

where

$$S(3) = (1 + \alpha) S(2) = 1.02 S(2) = 1.02 (\$20,400) = \$20,808.$$

So $P(3) = \$9063$ s. By this point, we see how things are going well enough to leave the pilot study and set up the recurrence relations that solve the problem with general parameter values. Clearly, the key equations are

$$S(j + 1) = (1 + \alpha) S(j) \quad (1.4)$$

and

$$P(j + 1) = \beta S(j) + (1 + \gamma) P(j), \quad \text{for } j = 1, 2, \dots, N_2 - N_1. \quad (1.5)$$

Moreover, at this point, it is easy to take account of the fact that pension increments are paid monthly via

$$S(j + 1) = (1 + \alpha) S(j), \quad \text{for } j = 1, 2, \dots, N_2 - N_1, \quad (1.6)$$

and

$$P_j(i + 1) = \frac{\beta}{12} S(j) + \left(1 + \frac{\gamma}{12}\right) P_j(i), \quad \text{for } i = 0, 1, 2, \dots, 11, \quad (1.7)$$

```

 $N_1$  = starting age of employment
 $N_2$  = retirement age
 $S$  = starting salary
 $P$  = starting principal
 $\alpha$  = annual rate of increase of salary
 $\beta$  = fraction of salary contributed by employee
 $\gamma$  = annual rate of increase of principal in fund
Year =  $N_1$ 
Month = 1
 $**P$  =  $\frac{\beta}{6}S + \left(1 + \frac{\gamma}{12}\right)P$ 
Month = Month + 1
Is Month = 13?
  If "no"  go to**
  If "yes" continue
    Year = Year + 1
     $S$  =  $S(1 + \alpha)$ 
    Month = 1
    Is Year =  $N_2 + 1$ ?
    If "no"  go to**
    If "yes" continue
      Return  $P$ 

```

Figure 1.1. Subroutine annuity ($N_1, N_2, S, \alpha, \beta, \gamma$).

where $P_{j+1}(0) = P_j(12)$. This system can readily be programmed on a handheld calculator or microprocessor using the simple flowchart in Figure 1.1. We find that the total stake of the employee at age 65 is a respectable \$450,298 (recall that we have used an interest rate and a salary incrementation consistent with a low inflation economy). We now know how much the employee will have in his pension account at the time of retirement. We wish to decide what the fair monthly payment will be if he invests the principal $P(N_2)$ in a life annuity. Let us suppose that actuarially he has a life expectancy of N_3 given that he retires at age N_2 . To do this, we first compute the value to which his principal would grow by age N_3 if he simply invested it in an account paying at the prevailing interest of γ . But this is easily done by using the preceding routine with $S = 0$ and setting N_1 equal to the retirement age N_2 and N_2 equal to the expected time of death N_3 . So we determine that using this strategy, the principal at the expected time of death is computed to be $P(N_3)$.

The monthly payments of the life annuity should be such that if they are immediately invested in an account paying at rate γ , then the total accrued principal at age N_3 will be $P(N_3)$. Let us suppose a guess as to this payment

X	=	guess as to fair monthly return
$P(N_2)$	=	principal at retirement
γ	=	annual rate of return on principal
N_2	=	retirement age
N_3	=	expected age at death given survival until age N_2
$P(N_3)$	=	$P(N_2) \left(1 + \frac{\gamma}{12}\right)^{N_3 - N_2}$

** Call Annuity $(N_2, N_3, X, 0, 6, \gamma)$
 Compare P with desired monthly payout
 Make new guess for X and return to**

Figure 1.2. Program trial and error.

is X . Then we may determine the total principal at age N_3 by using the flowchart in Figure 1.1 using $S = X, \alpha = 0, \beta = 6, \gamma = \gamma, N_1 =$ (retirement age) $N_2, N_2 =$ (expected age at death) N_3 . We can then find the fair value of monthly payment by trial and error using the program previously flowcharted in Figure 1.1.

We note that if the pensioner invests his principal at retirement into a fund paying 5% interest compounded monthly with all dividends reinvested, then at the expected death date he would have at that time a total principal of \$741,650.

$$P(N_3) = \$450,298 \left(1 + \frac{0.05}{12}\right)^{(75-65)12} = \$741,650. \quad (1.8)$$

Now on a monthly basis, the pensioner should receive an amount such that if he invested each month's payment at 5%, then at age 75 his stake would have increased from \$0 to \$741,650. As a first guess, let us try a monthly payout of \$5000. Using the program in Figure 1.2, we find a stake of \$776,414. Since this is a bit on the high side, we next try a payout rate of \$4,500—producing a stake at age 75 of \$737,593. Flailing around in an eyeballing mode gets us to within one dollar of $P(N_3)$ in a number of iterations highly dependent on the intuition of the user. Our number of iterations was nine. The equitable monthly payout rate is \$4,777.

It should be noted in passing that we have not taken into account the effect of inflation. What would the value of the first monthly payment be

in today's dollars if the inflation rate has progressed at the rate of 3% per year? The value of the first month's payout is then

$$\$4,777 \times (.97)^{41} = \$1,370$$

If the inflation rate should grow to 8%, then, the pensioner's first month check is in current dollars \$583. Social Security, on the other hand, is (supposedly) indexed for inflation. As John Bogle [2] has pointed out and as we shall demonstrate in Chapter 14, those who seek risk-free investment in fixed-rate bonds have failed to note that the actual value of a bond at maturity is an unknown because of inflation. The wise investor should put some of his or her annuity investment in stocks, since stocks, in a sense, self-adjust for inflation. Moreover, the current Social Security System has other benefits that the employee could elect to have incorporated into his payout plan, for example, survivorship benefits to a surviving spouse, and disability. Thus, our employer should really look further than bank returns. He needs to find something which is responsive to the cost of living. We will demonstrate later, in Chapter 14, how this might be achieved.

These additional add-ons would not cost sufficiently to lower the fair monthly payout below, say, \$3000 per month. And we recall that these are dollars in a low inflation economy. Considering that there is some doubt that an individual entering the current job market will ever receive anything from Social Security at retirement, a certain amount of indignation on the part of the young is perhaps in order. Furthermore, the proposed private alternative to Social Security would allow the investments by the employee and his employer in the plan to be utilized as capital for investment in American industry, increasing employment as well as productivity.

The method of trial and error is perhaps the oldest of the general algorithmic approaches for problem solving. It is highly interactive; that is, the user makes a guess (in the above case X) as to the appropriate "input" (control variable, answer, etc.) that he feeds into a "black box." An output result is spewed out of the black box and compared to the desideratum—in the earlier problem, P and $P(N_3)$.

We note that the above example is one in which we know the workings of the black box well. That is, we have a model for reality that seems to be precise. And little wonder, for annuity is a manmade entity and one should be able to grasp the reality of its workings far better than, say, daily maximum temperatures in Houston forecast 3 years into the future.

In actuality, however, even the pension fund example is highly dependent on a series of questionable assumptions. For example, the interest figures used assume negligible inflation—a fair assumption for the mid-1980s but a terrible one for the late 1970s. Objections as to the assumption that the pension fund will be soundly managed are not too relevant, since most such funds are broadly invested, essentially a random selection from the market. Objections as to the uncertainty of employment of the employee in his current company are also irrelevant, since it is assumed that the vesting

of such a fund is instantaneous, so that the employee loses no equity when he changes jobs. The plan suggested here is simply the kind of private IRA arrangement used so effectively by the Japanese as both a vehicle of retirement security and capital formation. Such a plan can reasonably be designed to track the performance of the overall economy. But the assumptions as to the specific yields of the plan will almost certainly be violated in practice. At a later time, we shall cover the subject of *scenario analysis*, in which the investigator frankly admits he does not fully understand the black box's workings very well and examines a number of reasonable sets of scenarios (hypotheses) and observes what happens in each. At this point, we need only mention that in reality we are always in a scenario analysis situation. We always see through a glass darkly.

Having admitted that even in this idealized situation our model is only an approximation to reality, we observe that a wide variety of algorithms exist for solving the problem posed. Usually, if we can come up with a realistic mathematical axiomatization of the problem, we have done the most significant part of the work. The trial-and-error approach has many advantages and is not to be despised. However, its relative slowness may be somewhat inefficient. In the present problem, we are attempting to pick X so that P is close to $P(N_3)$. It is not hard to design an automated algorithm that behaves very much like the human mind for achieving this goal. For example, in Figure 1.3, we consider a plot of $G = P - P(N_3)$ versus X . Suppose that we have computed G for two values X_{n-1} and X_n . We may then use as our next guess X_{n+1} , the intercept on the X axis of the line joining $[X_{n-1}, G(X_{n-1})]$ and $[X_n, G(X_n)]$.

Using the program in Figure 1.4, we use as our first guess a monthly output of 0, which naturally produces a stake at 75 years of age of 0. As our second guess, we use $X_2 = \$450,298/(7 \times 12)$. With these two starting values of X_n and G_n , the program converges to a value that gives $P = 741,650$ to within one dollar in three iterations. The equitable payout rate obtained by the secant method is, of course, the same as that obtained by trial and error—namely, \$4,777.

The Secant Method

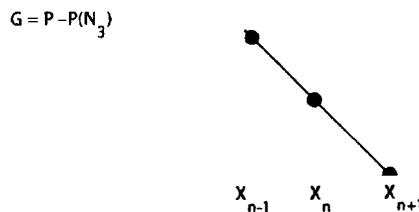


Figure 1.3. The secant method

$P(N_2)$ = principal at retirement
 γ = annual rate of return on principal
 N_2 = retirement age
 N_3 = expected age at death given survival until age N_2
 $P(N_3) = P(N_2) \left(1 + \frac{\gamma}{12}\right)^{12(N_3 - N_2)}$
 $X_{n-1} = 0$
 $G_{n-1} = -P(N_3)$
 $X_n = \frac{P(N_3)}{12(N_3 - N_2)}$
*Call Annuity $(N_2, N_3, X_n, 0, 6, \gamma)$
 $G_n = P - P(N_3)$
Slope = $= (G_n - G_{n-1}) / (X_n - X_{n-1})$
 $X_{n+1} = X_n - \frac{G_n}{\text{Slope}}$
Call Annuity $(N_2, N_3, X_{n+1}, 0, 6, \gamma)$
 $G_{n+1} = P - P(N_3)$
Is $G_{n+1} < 1$?
If yes, print X_{n+1} and stop
If no continue
 $X_{n-1} = X_n$
 $G_{n+1} = G_n$
 $X_n = X_{n+1}$
 $G_n = G_{n+1}$
Go to *

Figure 1.4. Program secant method

1.2 Income Tax Bracket Creep and the Quiet Revolution of 1980

Many, with some justification, believe that our leaders in Washington have something approaching a proprietary interest in the wealth of America's citizens. That works up to a point, but beyond some hard-to-establish threshold taxes can cause revolutions. The American Revolution is a case in point. The high tariffs on exported cotton and imported textiles led, in large measure, to the American Civil War. The discussion here concerns a quiet revolution that removed President Carter from office by the election of 1980.

Facing inflation of more than 10% and interest rates that reached 20%, President Carter gave a speech on July 15, 1979, in which he blamed the problem on a "general malaise" of the American people. This was heady stuff, for most Americans did not feel they were part of such a general shiftlessness. The General Malaise Speech, as it came to be called, flew in the face of the famous maxim of the long dead economist and sociologist Vilfredo Pareto. One form of Pareto's Maxim is that the catastrophically many failures are not due to a general malaise but to a small number of

assignable causes [[11, p.10].

It seemed to many that perhaps President Carter himself, rather than the general population, had allowed things to go very wrong with the nation's economy. The election of Ronald Reagan in November of 1980 had as its most important accomplishment the abolition of the Soviet Union. But the electorate was not dreaming of such a result. They were concerned about the deterioration in their standards of living. We will demonstrate below how the tax system contributed to the economic problem, which was the real reason for Carter's defeat.

Beginning in 1981, there were several changes in the U.S. income tax laws. The major reason advanced for these modifications in the tax regulations was something called "bracket creep." This is a phenomenon of the progressive income tax which causes an individual whose income increases at the same rate as inflation to fall further and further behind as time progresses. There was much resistance on the part of many politicians to this indexing of taxes to the inflation rate, since the existing tax laws guaranteed a 1.6% increase in federal revenues for every 1% increase in the inflation rate. Another problem that was addressed by the tax changes in the early 1980s was the fact that a professional couple living together in the unmarried state typically paid a few thousand dollars less in taxes than if they were married. Those who felt the need for indexing and some relief from the "marriage tax" had carried out several "if this goes on" type scenario analyses. We consider one such below. All the figures below use typical salary rates for 1980 and an inflation rate a bit below that experienced at that time. The tax brackets are those of the 1980 IRS tables.

Let us consider the case of John Rickenik who accepts a position with a company that translates into a taxable income of \$20,000. Let us project John's earning profile in the case where both inflation and his salary increase at an annual rate of 7%. First of all, we see that John's income will grow annually according to the formula

$$\text{income} = 20,000(1.07)^{\text{year} - 1980}. \quad (1.9)$$

The tax required to be paid in any given year is easily determined from Table 1.1. Inflation, on an annual basis, can be taken care of by expressing all after tax amounts in 1980 dollars according to the formula

$$\text{value in 1980 dollars} = \frac{\text{nominal amount}}{1.07^{\text{year} - 1980}}. \quad (1.10)$$

To determine John's after tax profile, we need to examine the 1980 tax tables for single taxpayers.

We will see that that John is not holding his own against inflation despite of the fact that his salary is increasing at the same rate as inflation. This is due to the fact that his marginal increases in salary are being taxed at rates higher than the average rate for the total tax on his earnings. The purpose

of indexing is to see that the boundaries for the rate changes increase at the annual rate of inflation.

We can readily compute the 6-year horizon table for John Rickenik's after tax income in 1980 dollars (Table 1.1).

Table 1.1. Rates for Single Taxpayers.

If taxable income

is not over \$2,300...	Over	But not over	0	of amount over
	\$2,300	\$3,400	14 %	\$2,300
	\$3,400	\$4,400	\$154 + 16 %	\$3,400
	\$4,400	\$6,500	\$314 + 18 %	\$4,400
	\$6,500	\$8,500	\$692 + 19 %	\$6,500
	\$8,500	\$10,800	\$1,072 + 21 %	\$8,500
	\$10,800	\$12,900	\$1,555 + 25 %	\$10,800
	\$12,900	\$15,000	\$2,059 + 26 %	\$12,900
	\$15,000	\$18,200	\$2,605 + 30 %	\$15,000
	\$18,200	\$23,500	\$3,565 + 34 %	\$18,200
	\$23,500	\$28,800	\$5,367 + 39 %	\$23,500
	\$28,800	\$34,100	\$7,434 + 44 %	\$28,800
	\$34,100	\$41,500	\$9,766 + 49 %	\$34,100
	\$41,500	\$55,300	\$13,392 + 55 %	\$41,500
	\$55,300	\$81,800	\$20,982 + 63 %	\$55,300
	\$81,800	\$108,300	\$37,677 + 68 %	\$81,800
	\$108,300	...	\$55,697 + 70 %	\$108,300

Table 1.2. After-Tax Income.

Year	Nominal	Tax	Tax Income	After-Tax Income (1980 dollars)
1980	\$20,000	\$4,177	\$15,823	\$15,823
1981	\$21,400	\$4,653	\$16,747	\$15,651
1982	\$22,898	\$5,162	\$17,736	\$15,491
1983	\$24,501	\$5,757	\$18,744	\$15,300
1984	\$26,216	\$6,426	\$19,790	\$15,098
1985	\$28,051	\$7,142	\$20,909	\$14,908

Let us now investigate the "marriage tax." Suppose that John Rickenik marries his classmate, Mary Weenie, who has the same earnings projections as does John—that is, 7% growth in both salary increments and inflation. You might suppose that computing the after tax income of the Rickenik family is trivial. All one has to do is to double the figures in Table 1.1. This is, in fact, the case if John and Mary live together without being legally married.

There is another table that applies to John and Mary if they are legally husband and wife. See Table 1.3.

In Table 1.3, we compare the tax John and Mary must pay if they are living in common-law marriage as compared to that if they are legally married. All figures are given in nominal dollar amounts (i.e., in dollars uncorrected for inflation).

Table 1.3. 1980 Tax Schedule for Married Couples.

Over	But not over		of the amount over
\$3,400	\$5,500	14 %	\$3,500
\$5,500	\$7,600	\$294 + 16 %	\$5,500
\$7,600	\$11,900	\$630 + 18 %	\$7,600
\$11,900	\$16,000	\$1,404 + 21 %	\$11,900
\$16,000	\$20,200	\$2,265 + 24 %	\$16,000
\$20,200	\$24,600	\$3,273 + 28 %	\$20,200
\$24,600	\$29,900	\$4,505 + 32 %	\$24,600
\$29,900	\$35,200	\$6,201 + 37 %	\$29,900
\$35,200	\$45,800	\$8,162 + 43 %	\$35,200
\$45,800	\$60,000	\$12,720 + 49 %	\$45,800
\$60,000	\$85,600	\$19,678 + 54 %	\$60,000
\$85,600	\$109,400	\$33,502 + 59 %	\$85,600
\$109,400	\$162,400	\$47,544 + 64 %	\$109,400
\$162,400	\$215,400	\$81,464 + 68 %	\$162,400
\$215,400	...	\$117,504 + 70 %	\$215,400

Table 1.4. The Marriage Tax.

Year	Comd. Income	Tax Married	Tax Unmarried	Marrg. Tax
1980	\$40,000	\$10,226	\$8,354	\$1,872
1981	\$42,800	\$11,430	\$9,306	\$2,124
1982	\$45,796	\$12,717	\$10,324	\$2,394
1983	\$49,002	\$14,289	\$11,514	\$2,775
1984	\$52,432	\$15,970	\$12,852	\$3,118
1985	\$56,102	\$17,768	\$14,284	\$3,484

We note that, even corrected for inflation, the marriage tax is increasing year by year. For example, when measured in constant dollars, the 1980 marriage tax of \$1872 grows to \$2484 in 1985.

Despite the tax disadvantages, John and Mary decide to get married. Their living expenses, until they can buy a house, are \$27,000 per year (in 1980 dollars). To buy a house, the Riceniks need a down payment of \$10,000 (in 1980 dollars). Assume that at the end of each year until they can make a down payment they invest their savings at an annual rate of 8% in short-term tax-free bonds. How many years must the Riceniks save in order to acquire their home if their wages and inflation increase at an annual rate of 7% and if the tax tables for 1980 had been kept in place? We answer this question by creating Table 1.5. As we note from the table, the Riceniks actually would never be able to afford their house under the

conditions given. From 1985 on, they would actually see the diminution of their savings.

Next, let us give a savings profile of the Ricensiks with conditions as given earlier, except with the change that income tax levels are indexed by inflation, and the marriage penalty has been eliminated. Here, the relevant taxes in 1980 dollars can be obtained by doubling \$4177, the tax for a single person earning \$20,000 per annum. This then gives us Table 1.6 (in 1980 dollars).

Table 1.5. The Savings Profile of John and Mary Ricensik.

Year	Income	Tax	After Tax Inc.	Savings	Accum.	Savings 1980 Dollars
1980	\$40,000	\$10,226	\$29,774	\$2,274	\$2,274	\$2,274
1981	\$42,800	\$11,430	\$31,370	\$1,945	\$4,401	\$4,113
1982	\$45,796	\$12,718	\$33,078	\$1,593	\$6,346	\$5,543
1983	\$49,001	\$14,289	\$34,712	\$1,023	\$7,877	\$6,430
1984	\$52,432	\$15,970	\$36,462	\$415	\$8,922	\$6,807
1985	\$56,102	\$17,768	\$38,334	-\$236	\$9,400	\$6,702

Table 1.6. Savings Profile with Indexing and No Marrg. Tax.

Year	Income	Tax	Income	Savings	Savings
1980	\$40,000	\$8,354	\$31,646	\$4,146	\$4,146
1981	\$40,000	\$8,354	\$31,646	\$4,185	\$8,331
1982	\$40,000	\$8,354	\$31,646	\$4,224	\$12,555
1983	\$40,000	\$8,354	\$31,646	\$4,263	\$16,818
1984	\$40,000	\$8,354	\$31,646	\$4,304	\$21,122
1985	\$40,000	\$8,354	\$31,646	\$4,343	\$25,465

Thus, if the income tax were indexed to inflation, the Ricensiks could afford the down payment on their home in less than 3 years even if their salaries only kept pace with inflation. In the above case, the only participation the Ricensiks see in the assumed increase of productivity made possible by technology is through the interest on their capital. Even without this interest, they would be able to move into their home in 1982. It is perhaps interesting to note that most people who go through the computations in Table 1.6 do not believe the results when they create the table for the first time.

Next, let us briefly consider one of the consequences of the pre-1987 U.S. income tax—the tax shelter. The top tax level on salaried income was reduced in the early 1980s to 50%—a reform instituted during the Carter administration. However, it is not surprising that many people would search for some legal means of avoiding paying this “modest” rate of taxation. This can be achieved by noting that income which was gained by making an investment and selling it at least 6 months later was discountable in reporting total income by 60%. Thus, a property that is purchased for

\$100 and sold for \$200 results in a profit of \$100, but a taxable profit of only \$40. Thus, the after-tax profit is not \$50 but \$80. Not surprisingly, this "loophole" made possible a large tax avoidance industry. We shall see below an example of what should by all reason be deemed a bad real estate investment which turns out to result in approximately the same after tax profit as that obtained by a "good" conventional investment. The example is one simple indication of how taxation polity can affect, very dramatically, the increase of investments which do little good to society over those which provide the engine of capitalism.

Ms. Brown is an engineer whose income level puts her in the 50 % bracket on the upper level of her income tax. She decides to purchase a real-estate acreage for a nominal price of \$100,000. The terms are \$10,000 down, and \$10,000 payable at the beginning of each year with 10 % interest on the balance payable 1 year in advance. Let us suppose Ms. Brown sells the property at the end of the fifth year. If the selling price is \$135,000, how well has Ms. Brown done? Let us examine Ms. Brown's cash outflow during the 5 years (see Table 1.7).

Table 1.7. A Real Estate Investment Outflow Profile.

Year	Principal Investment	Interest Investment
1	\$10,000	\$9,000
2	\$10,000	\$8,000
3	\$10,000	\$7,000
4	\$10,000	\$6,000
5	\$10,000	\$5,000

Ms. Brown has invested \$85,000 and still owes \$50,000 on the property at the time of sale. Thus, it would seem that the result of five years' investment is that Ms. Brown gets back exactly what she has put into the investment. Furthermore, Ms. Brown must pay capital gains tax of 20% on the profit of \$35,000. So, Ms. Brown's after-tax stake at the end of the five years is \$78,000. Clearly, Ms. Brown has not done well in her investing.

Let us suppose that Ms. Brown had followed a more traditional investment strategy. Suppose that she had taken the available salary income and invested in a money market fund at an annual rate of 10%. Now, we must note that the \$9,000 interest payment which she made in the first year of the real estate investment is deductible from her gross income. If she had not paid it in interest, she would have had to pay half of it to the federal government in taxes. So she would not have had \$19,000 to invest in a market fund in year 1—only \$14,500. Moreover, the interest which she would earn in year 1, \$1,450, would have been taxed at the 50% rate. Consequently, the after tax capital of Ms. Brown using the money market strategy would be given as shown in Table 1.8.

Note that Ms. Brown has approximately the same after tax capital using either the "imprudent" tax shelter or the "prudent" money market. It

is not that the real estate investment has become good; it is rather that the tax system has made the money market investment bad. This is an example of the means whereby individuals have been forced by the tax system into making unproductive investments rather than depositing the money with lending agencies who would, in turn, lend out the money for capital development.

Table 1.8. After-Tax Principal With Money Market Strategy.

Year	After tax capital at year's end
1	\$15,225
2	\$30,686
3	\$46,396
4	\$62,365
5	\$78,609

In September of 1986, the Rostenkowski-Packwood Tax Reform eliminated the capital gains preference, thus removing the incentive for the kind of bad investment mentioned here. Since much of the capital gains investment had favored the rich, Packwood and Rostenkowski set the marginal for families in the \$150,000 per year and up range at 28%. Marginal tax rates for the upper middle class, on the other hand, were set at 33%. So tax treatment which favors the wealthy and shafts the middle class was frankly and straightforwardly institutionalized instead of concealed with a capital gains exclusion. Other Rostenkowski-Packwood reforms included a suspension of the marriage deduction and that for state sales taxes. Moreover, the inflation indexing provisions of the reforms of the early 1980s appeared to be put at the hazard, "since the marginal rates are now so low." The "transitional" year marginal rates for the upper middle class of 38% were proposed to be made permanent by the Speaker of the House, Jim Wright. Some have cynically argued that, having wiped out a host of deductions in exchange for lower rates, the Congress will now gradually raise the new rates.¹

One thing that was not changed by any of the tax reforms of the last decade is the Social Security tax on income below roughly \$40,000 per worker (a ceiling raised steadily by legislation passed years ago). This tax

¹For whatever reasons one might conjecture, both Rostenkowski and Packwood were pilloried and forced out of office over matters that seem less serious than the punishments they received. Rostenkowski actually went to prison (his crime seems to have been renting a property owned by his sister for his Chicago office). Packwood appears to have been something of a Don Juan, but not even in the same league as Bill Clinton or John Kennedy. Nor is it recorded that he was responsible for the death of any of his conquests, as was the case with Ted Kennedy. It would appear that Rostenkowski and Packwood may have angered some influential people with their reforms. At any rate, President Clinton raised the marginal tax rates and restored the capital gains reduction. The point of this section has been to show how tax rates can remove from office a President and cause a vendetta against tax reformers.

is approximately 7% from the employee and 7% from the employer, for a pooled total of slightly more than 14%. The self-employed pay roughly 10%. Many citizens pay more in Social Security tax than in income tax. And, of course, the monies paid by the worker in Social Security tax are also subject themselves to federal income tax. In addition to the two federal income taxes mentioned above, most states also levy income taxes—amounting to sums as high as 25% of the federal income tax. Furthermore, almost all states and municipalities levy sales taxes, real estate taxes, and so on. The proportion of an American citizen's income which goes to pay taxes is now several times that two centuries ago when taxes were a major cause of the American Revolution against Great Britain. Still, things could have been worse; the British taxation rates today are much more severe even than those in the United States.

We have gone through this section to demonstrate how a relatively simple and well-defined system such as the U.S. income tax is by no means easy to grasp until one goes through some computations with actual figures. When we leave the comfortable realm of well-axiomitized man-made systems and go into the generally very imperfectly understood systems of the social and natural sciences, we may well expect our difficulties of comprehension to increase dramatically.

1.3 Retirement of a Mortgage

In 2009, in response to plummeting house prices, the federal government made it possible for mortgages to drop below 5%. However, so much dubious activity had been undertaken by American banks that several of the prudently paranoid tried to pay off the remaining portion of their house mortgages. The reason is that in the event of crises, the homeowner might find it hard to get the cash to make a mortgage payment. "Bank holidays" can occur arbitrarily at the behest of the government. If a banker is considering which of two houses each valued at \$175,000, on which to foreclose, he will no doubt prefer to foreclose the one on which only \$15,000 is owed rather than the one on which the full \$175,000 is owed.

Let us now consider the case of Mr. and Mrs. Jones who have 10 years to run on their mortgage. They now owe a principal of \$15,000 on which they pay 5.5% interest on the outstanding balance. Suppose recent mortgage rates are in the 4% range. How much should the Joneses have to pay to retire the mortgage? The naive answer is "obviously \$15,000." But is the mortgage really worth \$15,000 to the bank who holds it? Obviously not. If the bank had the \$15,000 in hand, they could lend it out at only 4%. We wish to find the fair market value of the Jones mortgage.

One answer that might be proposed is that the fair market value is that quantity—say x —which, when compounded at 4%, will be worth the same as \$15,000 compounded at 5.5%. This amount may easily be computed via

$$X(1.04)^{10} = \$15,000(1.055)^{10}. \quad (1.11)$$

or

$$X = \$15,000 \left(\frac{1.055}{1.04} \right)^{10} = \$17,309.42 \quad (1.12)$$

But this answer is also wrong, for we recall that under the terms of a mortgage, the principal is retired over the entire term of the mortgage—not at the end. As the mortgage company receives the principal, it can loan it out at the prevailing mortgage rate—say 4%. The usual rule of principal retirement for older mortgages is that monthly payout is determined so that each installment of principal retired plus interest on unpaid balance is equal. Thus, the Joneses are to pay out their mortgage over 120 months according to the rule

$$\begin{aligned} y &= p_1 + \$15,000 \frac{0.055}{12} \\ &= p_2 + (\$15,000 - p_1) \frac{0.055}{12} \\ &= p_3 + (\$15,000 - p_1 - p_2) \frac{0.055}{12} \\ &= \dots = p_{120} + (\$15,000 - p_1 - p_2 - \dots - p_{119}) \frac{0.055}{12}, \end{aligned} \quad (1.13)$$

where p_i is the principal retired on month i . By the 120th payment, all the \$15,000 will have been retired. Hence,

$$\sum_{i=1}^{120} p_i = \$15,000. \quad (1.14)$$

Using the fact that monthly payments are equal, we have

$$p_1 + \$15,000 \left(\frac{0.055}{12} \right) = p_2 + (\$15,000 - p_1) \left(\frac{0.055}{12} \right). \quad (1.15)$$

This gives

$$p_2 = \left(1 + \frac{0.055}{12} \right) p_1. \quad (1.16)$$

Similarly, from the second and third equalities above, we have

$$p_3 = \left(1 + \frac{0.055}{12} \right) p_2. \quad (1.17)$$

or

$$p_3 = rp_2, \quad (1.18)$$

where

$$r = 1 + \frac{0.055}{12}.$$

Thus, we have

$$p_i = r^{i-1} p_1. \quad (1.19)$$

Using the fact that $\sum p_i = \$15,000$, we have

$$p_1(1 + r + r^2 + \dots + r^{119}) = \$15,000. \quad (1.20)$$

Recalling the formula for the sum of a geometric progression, we have

$$p_1 \left(\frac{1 - r^{120}}{1 - r} \right) = \$15,000. \quad (1.21)$$

This yields a p_1 value of \$94.04. Any other p_i value is immediately available via $p_i = r^{i-1} \$94.04$. The monthly payment is given simply by

$$\$94.04 + \$15,000 \left(\frac{0.055}{12} \right) = \$162.79.$$

Moreover, the interest payment for the i th month is simply obtained via

$$\left(\$15,000 - \sum_{j=1}^{j=i-1} p_j \right) \frac{0.055}{12}. \quad (1.22)$$

But this is just

$$\left(\$15,000 - p_1 \sum_{j=0}^{j=i-1} r^{j-1} \right) \frac{0.055}{12}. \quad (1.23)$$

or

$$\left(\$15,000 - p_1 \frac{1 - r^i}{1 - r} \right) \frac{0.055}{12}. \quad (1.24)$$

We are now in a position to answer the question as to the fair value for an instantaneous payment of the Jones mortgage. If the company that holds the Jones's mortgage lends out the monthly current mortgage payment at an annual rate of I , then the total value of monies received over the ten years will be

$$\$162.79(1 + s + s^2 + \dots + s^{119}) = \$162.79 \frac{1 - s^{120}}{1 - s}. \quad (1.25)$$

where

$$s = 1 + I/12.$$

Suppose the interest were the 5.5% the Jones currently pay. Then the total value of monies received over the ten years would be \$25,966.24. To compute present value of the mortgage, X , we solve

$$Xs^{120} = \$25,966.24. \quad (1.26)$$

Of course, we get the figure of \$15,000. But if the lender is only able to make a return of 4% on a new mortgage, then from (1.24), the total cash he will receive from the alternative mortgage is \$23,941.87. And, from (1.25), the present value of the alternative mortgage to the lender is \$17,412.34. Unless the borrower had no cost for early repayment written into his mortgage, owning his home without mortgage liability will cost at least \$17,412.34. Given the kindness of bankers, \$20,000 might be closer to the mark.

1.4 Some Mathematical Descriptions of the Theory of Malthus

In this chapter on models of growth, we began with compound interest because, historically, this seems to be the oldest of the quantified growth models. People have been lending out money at interest rates since the dawn of civilization. Because such financial transactions are essentially human constructs, it is not surprising that the modeling of such a process has been well understood for so long. In a sense, it is as though the process proceeds directly from an idealized model. Of course, even here, in what should be the most straightforward kind of modeling situation, there are many problems. Inflation rates are not generally predictable; wars and revolutions disrupt the orderly process of commerce; bankruptcy laws are created to protect debtors from their creditors, and so on. The point is that even in the most simple kind of modeling—that is, the modeling of processes based on idealized models—we are unable to describe precisely what is going on.

When it comes to modeling processes that do not proceed directly from a man-made model, we shall expect to be peering through even muddier water. Still, as we shall see, the attempt to try and understand at least in part what is happening is well worth the effort. In fact, these attempts to conceptualize portions of the world around us have existed as long as humankind. It is our ability of the last 300 years to mathematize our conceptions and the commercial impetus to do so which have been largely responsible for the rapid scientific and technological progress characterizing this period. And it is this progress that has been responsible for the dramatic improvement in the material standard of living in those countries where this progress has been permitted.

Let us now consider the revolutionary work of the Reverend Thomas Malthus [9], *An Essay on the Principle of Population*. This book (published in 1798 and available as a free download on the Internet), generally speaking, was not as explicitly mathematical as one might have wished. Malthus used words rather than equations. His basic thesis is succinctly, if ambiguously, given by, “Population, when unchecked, increases in a geometrical ratio. Subsistence increases only in an arithmetical ratio.” If the qualifying phrase “when unchecked” had been omitted, we could have summarized these two sentences very simply, using the symbols P for population and F for food by the two differential equations

$$\frac{dP}{dt} = \alpha P \quad \text{and} \quad \frac{dF}{dt} = \beta. \quad (1.27)$$

The solutions to these equations are given simply by

$$P(t) = P(0)e^{\alpha t} \quad \text{and} \quad F(t) = F(0) + \beta t. \quad (1.28)$$

The consequence would then be that the population was increasing at an exponential rate without any constraint from the more slowly growing food supply. But Malthus coupled the two processes, population and food (“subsistence”), with the use of the phrase “when unchecked.” This simple qualifying phrase implies that a shortage of food checks the population growth. How should this checking effect be incorporated into our mathematical model? We might first decide that our constant α is not a constant at all but rather a function of F and P . To do so, however, we have to express population and food in similar units. This might be accomplished by using as one unit of food that amount required to sustain one person for a given unit of time. We then arrive at the coupled differential equation model

$$\frac{dP}{dt} = \alpha(F, P)P \quad \text{and} \quad \frac{dF}{dt} = \beta, \quad (1.29)$$

where $\alpha(F, P)$ is read “ α of F and P .” Note that whereas F affects P , there is no effect of P on F . Accordingly, we may write

$$F(t) = F(0) + \beta t. \quad (1.30)$$

Suppose we argue that the growth rate per person is proportional to the difference between the food available and the food consumed. Then we have

$$\frac{1}{P} \frac{dP}{dt} = a [F(0) + \beta t - P]. \quad (1.31)$$

In the case where the food supply is unchanging (i.e., $\beta = 0$), this is the *logistic* model of Verhulst (1844). A simple integration by parts for this special case gives

$$P = \frac{FP(0)e^{atF}}{P(0)e^{atF} + F - P(0)}. \quad (1.32)$$

Consequently, we shall refer to (1.31) as the *generalized logistic* model. The solution is given by

$$P(t) = \frac{P(0) \exp [a(F(0)t + \frac{1}{2}\beta t^2)]}{1 + aP(0) \int_0^t \exp [aF(0)\tau + \frac{1}{2}a\beta\tau^2] d\tau}. \quad (1.33)$$

The present application provides some promise from a curve fitting standpoint. For example, for t large (1.33) is approximately given by

$$P(t) \approx F(0) + \beta t. \quad (1.34)$$

For a long established society, we might expect the Malthusian argument would yield a population that tracks the supposedly linear growth of the food supply. This is consistent with the population growth in England and Wales from 1800 to 1950 as shown in Figure 1.5. Using 1801 as the time origin, and using population in units of 1 million, we can fit the line by eye to obtain

$$P(t) \approx 6.4 + 0.2(t - 1801). \quad (1.35)$$

Naturally, there is little chance of determining a , since population records in Britain in the remote past, when population growth may have been less limited by the growth in the food supply, are not available. However, in the remote past, when t is "small," becomes effectively

$$P(t) \approx P(0)e^{aF(0)t}. \quad (1.36)$$

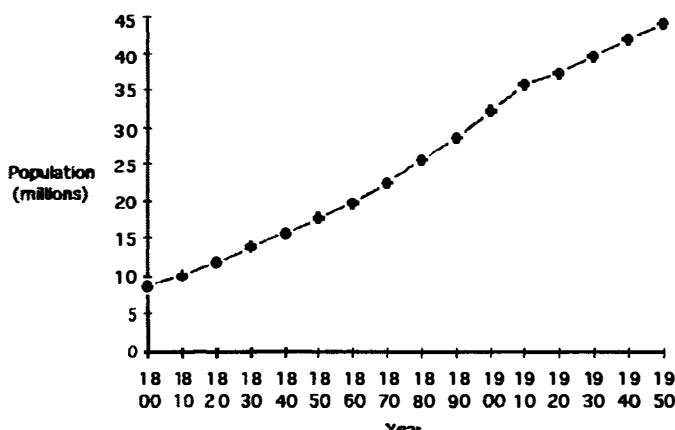


Figure 1.5. Population growth in England and Wales.

Looking at the population growth of the United States in Figure 1.6, we note that whereas 20th-century population growth is approximately linear, growth prior to, say, 1880 is much faster than linear. The only functional curve that can readily be identified by humans is the straight line. Accordingly, we take natural logarithms of both sides of (1.35). This gives us the equation of a straight line in t

$$\log [P(t)] \approx \log(P(0)) + aF(0)t. \quad (1.37)$$

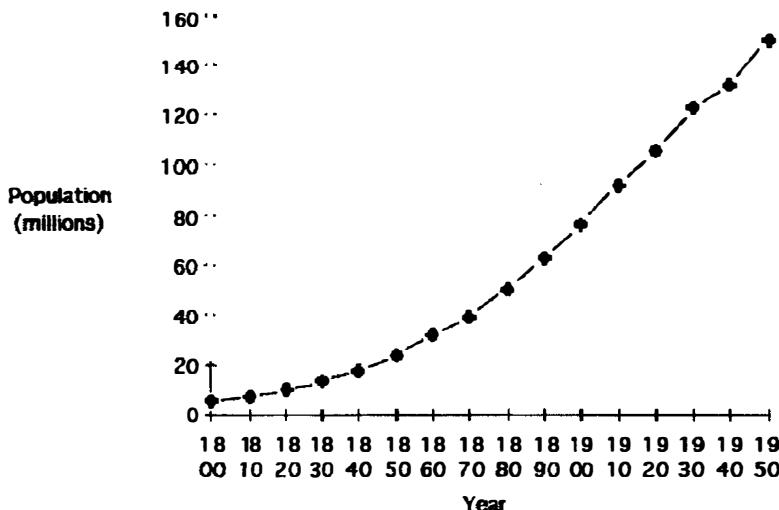


Figure 1.6. U.S. population growth.

If (1.36) holds approximately over some time interval, then we should expect that a plot of $\log(P)$ versus time would yield very nearly a straight line. As we note from Figure 1.6, it seems that for the first 70 years of the 19th century, population growth is consistent with Figure 1.7. We shall not at this point go through the argument for estimating all the parameters in (1.37). But let us suppose that (1.37) could be made to fit closely both the American and British data. Would we have “validated (1.32) that is, would we have satisfied ourselves that (1.32) is consistent with the data and our reason? Now a monotone data set can be fit with many curve families particularly if they have three free parameters—as does (1.37). A good curve fit is a necessary condition but not a sufficient condition for establishing the “validity” (“plausibility” might be a better word) of a model.

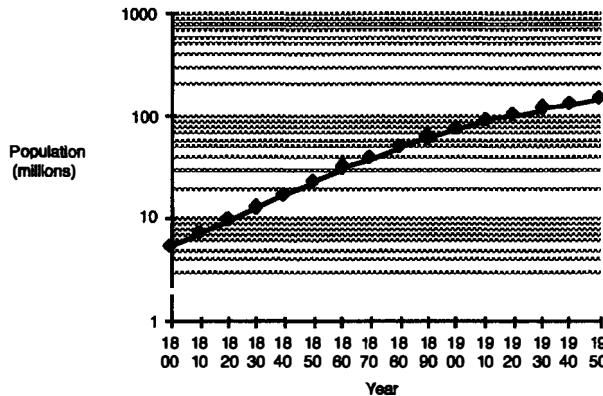


Figure 1.7. Logarithm of U.S. population growth.

Let us reexamine (1.37) in light of the hypotheses of Malthus and our own reasoning. Is it reasonable to suppose that the per-person population rate of increase will depend on the total availability of food? It would appear that if the food supply will support a population of 1,100,000, when the actual population is 1,000,000, then the per-person rate of increase should be the same for a population of 100,000 which has an available food supply for 110,000. Accordingly, (1.37) is not appropriate. Rather, we should prefer

$$\frac{1}{P} \frac{dP}{dt} = \frac{a[F(0) + \beta t - P]}{P}. \quad (1.38)$$

The solution to this equation is

$$P(t) = \left(P(0) + \frac{\beta}{a} - F(0) \right) e^{-at} + F(0) - \frac{\beta}{a} + \beta t. \quad (1.39)$$

It is interesting to note that in (1.39), for large values of time, the population curve is given essentially by

$$P(t) = F(0) - \frac{\beta}{a} + \beta t. \quad (1.40)$$

Thus, the model is consistent with the British population growth figures and with those from the United States after, say, 1910.

To examine the behavior of $P(t)$ for small values of t , we expand the exponential term neglecting all terms of order t^2 or higher (i.e., $O(t^2)$) to give

$$P(t) \approx \left[P(0) + \frac{\beta}{a} - F(0) \right] (1 - at) + F(0) - \frac{\beta}{a} + \beta t; \quad (1.41)$$

that is,

$$P(t) \approx P(0) + [F(0) - P(0)]at. \quad (1.42)$$

Now we note that (1.38) departs from the spirit of Malthus, since growth consists essentially of two pieces linear in time with a transition in-between. The differential equation (1.37), moreover, has the problem that there is no "threshold" phenomenon insofar as per-capita food supply is concerned. If there is an excess capacity of food per person of 20, then people will multiply 20 times as fast as if the per capita excess capacity were 1. This is clearly unreasonable. However, if we fit the model and find out that the per capita excess food capacity is small throughout the extent of the data, then we might find the model satisfactory. As it is not too much trouble to try a rough-and-ready fit, we shall attempt one, using U.S. data in Table 1.9.

Table 1.9. U.S. Population 1800–1950.

Date	t	$P(t)$ in millions
1800	0	5.308483
1810	10	7.239881
1820	20	9.638453
1830	30	12.866020
1840	40	17.069453
1850	50	23.191876
1860	60	31.443210
1870	70	38.558371
1880	80	50.155783
1890	90	62.947714
1900	100	75.994575
1910	110	91.972266
1920	120	105.710620
1930	130	122.785046
1940	140	132.169275
1950	150	150.697361

First of all, using 1800 as the time origin, we have $P(0) = 5.308483$. Then, using the population figures for 1800 and 1810 as base points, we have from (1.42)

$$F(0) - \frac{0.19314}{a} = 5.308483. \quad (1.43)$$

Then, using (1.41) for $P(120)$ and $P(150)$, we have

$$\beta = 1.49956. \quad (1.44)$$

Finally, extending the line from $P(150)$ with slope 1.49956 back to the intercept on the population axis, we have from (1.44)

$$F(0) - \frac{1.49956}{a} = -74.9340. \quad (1.45)$$

Solving (1.43) and (1.45) together, we find

$$F(0) = 17.1715 \quad (1.46)$$

and

$$a = 0.016281. \quad (1.47)$$

Already, we should have some concern about the size of $F(0)$. This leads to a per capita excess food supply in 1800 of around 2, probably an unreasonably high figure, and one that is high enough that the absence of a threshold effect in excess food supply might very well render our model less useful than one might hope.

Using these “rough and ready” parameter values in (1.38), we see in Figure 1.8 a comparison between the model values and the actual population figures. The quality of the fit is not spectacular. However, with three parameters to juggle— $F(0)$, a , β —we could no doubt arrive at a very good fit with a little work. We might, for example, find the least squares solution by minimizing

$$S(a, \beta, F(0)) =$$

$$\sum_{j=0}^{j=15} \left[P(10j) - \left(P(0) + \frac{\beta}{a} - F(0) \right) e^{-10aj} - F(0) - \frac{\beta}{a} - 10j\beta \right]^2. \quad (1.48)$$

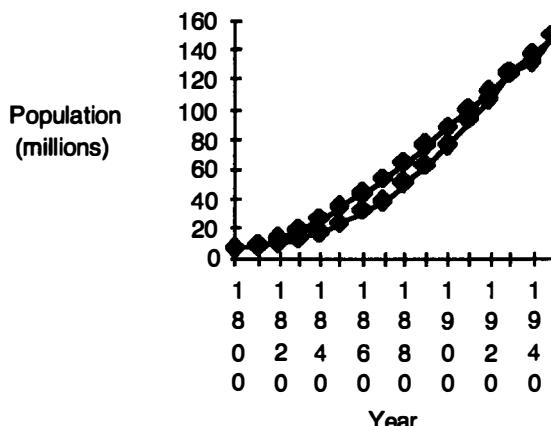


Figure 1.8. U.S. population growth and model values.

Now, let us go to a model more consistent with the apparent intent of Malthus. We need a model that will not penalize exponential growth when the excess per capita capacity of food is sufficiently large but will cause growth to be proportional to excess food capacity when the “affluence threshold” has been reached. There are an infinite number of models that will do this. Consequently, we here borrow a maxim from the Radical Pragmatist William of Ockham and try to pick a simple model that is consistent with Malthus’ conjecture and with the facts before us.

$$\frac{dP}{dt} = aP, \quad \text{when } \frac{F(0) + \beta t - P}{P} > \frac{1}{k}. \quad (1.49)$$

and

$$\frac{dP}{dt} = ak[F(0) + \beta t - P], \quad \text{otherwise.} \quad (1.50)$$

Examining the logarithmic curve in Figure 1.8, it seems that population growth is close to log-linear until around 1860. So we shall guess that the time where we switch from (1.30) to (1.31) is $t^* = 1860$. From the data at 1800 and 1860, we can estimate a via

$$a = \frac{1}{60} \ln \left(\frac{31.44321}{5.30843} \right) = 0.029648. \quad (1.51)$$

We know what the solution to (1.49) is from (1.38), namely,

$$P(t) = \left[P(t^*) + \frac{\beta}{ak} - F(t^*) \right] e^{-ak(t-t^*)} + F(0) + \beta t - \frac{\beta}{ak}. \quad (1.52)$$

Going to the essentially linear part of the model, we have, as when fitting (1.52), that $\beta = 1.49954$. By examining the intercept when $t = t^*$, we find that

$$F(t^*) = F(0) + \beta t^* = 15.73696 + \frac{50.57879}{k}. \quad (1.53)$$

We still must estimate $F(0)$. In this case, we shall guess that a value around 10 will be close to the truth. From (1.53), this gives us a k value of 0.600437. Finally, then, we have the estimated model

$$P(t) = P(0)e^{0.029648t}, \quad \text{before 1860,} \quad (1.54)$$

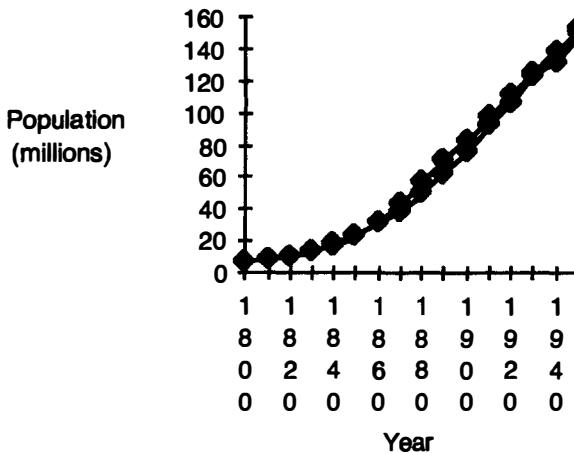


Figure 1.9. U.S. population growth vs. model values.

and

$$P(t) = 15.70611e^{-0.017801(t-60)} - 74.2530 + 1.40053t, \text{ thereafter. (1.55)}$$

We see the results of our fit in Figure 1.9. With a bit more careful juggling of the four parameters a , k , $F(0)$, and β , we could essentially make the fitted curve nearly repeat precisely the actual data set.

Have we then established the “validity” of Malthus’s model? By no means. But even with the crude fitting strategy employed, there is some reason for acknowledgment if not of the validity of the Malthusian model, then at least of its value as a framework within which to consider population data. Recall that population increases of nations have been for many years measured by per capita rate of increase, or its equivalent—just as Malthus said we should. Reasoning from the standpoint of Malthus, we might argue that in the recent stages of the United States, its population growth was linear. This linear growth was also exhibited in the British data. These facts are directly consistent with what one would predict from Malthus’ hypotheses. Moreover, it is of interest to note that all these data were collected after Malthus’ book was published. A model is generally more believable if it can predict what will happen, rather than “explain” something that has already occurred. (Note, however, that by fitting our model using data from the entire time range of the data base, we have engaged in a bit of “sharp practice.” Really to “validate” a model in time, we should obtain the parameters using values from early time and see if the model can then predict the later behavior.)

However, if we really wished to give a thorough investigation of the validity of Malthus’s hypotheses, we should incorporate more factors into the

model. For example, in the case of the American population, we should take some cognizance of the effects of immigration. (It would be possible to eliminate immigration as a significant demographic factor from some rapidly growing populations, such as that of 19th-century India.) Moreover, we have not directly checked Malthus' assumption as to the linear growth of the food supply. Malthus's implicit assumption that food strongly affects population growth, but that the change in the population has little feedback effect on the change in the food supply, would also need to be checked. Moreover, we must face the fact that the use of food as the sole limiting factor is too restrictive. Indeed, even a generalization of the concept of "food" to something like per capita productivity has problems. We know, for example, that if we substratify the American population into economic groups, we would find a much higher rate of population growth for the lower income levels than for the middle class. This is not consistent with Malthus's hypotheses. And, indeed, we cannot simply assume that the hypotheses have been invalidated by the welfare state. This kind of high reproductive rate of the poor relative to that of the well to do also exists in Mexico, for example, which has very little in the way of welfare. It would appear that in many contemporary societies the Malthusian hypotheses have been, in a sense, stood on their head with material well-being rather than with poverty acting as a brake on population growth. Almost without exception, the rapid rate of population growth in the Third World (before the AIDS epidemic) is not due to increases in the birth rate, but rather to a lowering of infant mortality using the cheap end of First World medical technology (vaccination, etc.) without the compensating decreases in the birth rate which have occurred in the First World. On the other hand, we do see in large stretches of Africa, an apparent return to the Malthusian model, as various governments there have managed to wreck their economies to the extent that even the most rudimentary trappings of modern medicine (e.g., sterile instruments) are disappearing and the food production even in fertile areas has been seriously damaged. It would appear that for subsistence economies, Malthus's model retains a great deal of validity; but that for First World societies (which, of course, postdate Malthus), radical modifications to his model would be required.

In the above, we note one of the primary values of a mathematical model—namely, as a framework in which to view data. The model gives us a benchmark by which to gauge our observations. It forces us to ask questions about the implications of the data in the light of our "best guess" made prior to the data analysis itself. This will generally lead to modifications of the original model, which will give us a new framework requiring that other data be collected, and so on.

If we take the position that the scientist should clear his or her mind of all prejudices about the generating mechanism of the data and "let the data speak for itself," we lose much. The human mind simply does not ever "start from zero" when analyzing phenomena. Rather, it draws on

instinct plus a lifetime of experience and learning to which is added the current information. Empirical model building is simply a formulation of this natural learning process.

On the other hand, models, such as that of Malthus, which were conjectured based on *logical* feelings as the way things ought to be instead of basing things on data, are almost always flawed, frequently disastrous. During the famous Irish potato famine, the Malthusian model was used as an excuse by the English establishment to let the Irish starve during the potato blight rather than simply giving them a portion of the Irish grain crop which was perfectly healthy and plentiful but raised on huge estates owned by the English and worked by poor Irish peasants. The argument was that if one fed the starving Irish, it would only encourage them to have more children and make things worse. This variant of Malthusianism is sometimes referred to as *Social Darwinism* and it is absolutely contrary to the Christian ethos of Western civilization even if the model were correct (it was not over the long term). In our own times we see much in human demographics to contradict Malthus's model. Some of the world's most rapid human population growth is among people who are desperately poor. Indeed, it may well be the case that the high population growth amongst, say, the Palestinians is due to a notion that it is their only viable weapon against the Israelis, whose economy and weaponry are far better than those of the Palestinians. And in most of the First World, with its relatively abundant resources, most countries show a negative population growth. It would seem that Malthus did not come up with anything like a universal law. During the 210 years since he presented his model, it really has not performed very well. And it has been used by some governments with very cruel and disastrous effects.

In Chapter 14, we shall discuss the efficient market hypothesis (EMH), which was developed in the 1960s with reliance on *bright ideas* instead of on data, which stood in stark contradiction to the EMH. As Sherlock Holmes used to remark frenziedly to Dr. Watson, "Data, Watson. I must have data."

1.5 Metastasis and Resistance

The killing effect of cancer is a result of the fact that a cancerous tumor grows, essentially, without limit and that it can spread to other parts of the body in a cascading sequence of *metastatic* events. On the other hand, if we can remove the tumor before it spreads via *metastatic progression*, that is, the breaking off of a cell from the primary that lodges in sites remote from it, then the chances of cure are excellent. On the other hand, once the primary has produced metastases to sites remote from the primary, surgical removal of the primary alone will generally not cure the patient's cancerous condition. If these metastases are widely spread, then, since each

surgical intervention involves a weakening of the patient and the removal of possibly essential tissue, alternatives to surgery are required. Radiotherapy is also contraindicated in such cases, since radiation, like surgery, generally should be focused at a small region of tissue. Bartoszyński et alia [1] have postulated that the tendency of a tumor to throw off a metastasis in a time interval $[t, t + \Delta t]$ is proportional to its size (in tumor cells) $n(t)$:

$$P(\text{metastasis in } [t, t + \Delta t]) = \mu n(t) \Delta t, \quad (1.56)$$

where μ represents the tendency to metastasize.

For all intents and purposes, only chemotherapy is left to assist a patient with diffuse spread of a solid tumor. Happily, we have a vast pharmacopoeia of chemotherapeutic agents that are very effective at killing cancerous tissue. Unfortunately, the general experience seems to be that within most detectable tumors, cells exist that resist a particular regimen of chemotherapy. Originally, it was thought that such resistance developed as a feedback response of the malignant cells to develop strains that resisted the chemotherapy. Currently, it is believed that resistance develops randomly over time by mutation and that the eventual dominance of resistant cells in a tumor is simply the result of destruction of the nonresistant cells by the chemotherapeutic agent(s), a kind of survival-of-the-fittest cells. It is postulated that, during any given time interval $[t, t + \Delta t]$,

$$P(\text{resistant cell produced during } [t, t + \Delta t]) = \beta n(t) \Delta t, \quad (1.57)$$

where β represents the tendency to develop irreversible drug resistance and $n(t)$ is the size of the tumor at time t .

Now if the resistant cells are confined to the primary tumor, a cure would result if the primary were excised and a chemotherapeutic agent were infused to kill the *nonresistant* cells that might have spread metastatically to other sites in the body. Unfortunately, it might very well happen that some resistant cells would have spread away from the primary or that originally nonresistant metastases would have developed by the time of the beginning of therapy.

Accordingly, it might be appropriate in some cases to follow surgery immediately by a regime of chemotherapy, even though no metastases have been discovered. The reason is that this kind of "preemptive strike" might kill unseen micrometastases which had not yet developed resistance to the chemotherapeutic agent(s). Such a strategy is termed *adjuvant chemotherapy*.

At the time of presentation, a patient with a solid tumor is in one of three fundamental states:

no metastases; (1.58)

metastases, none of which contain resistant cells; (1.59)

metastases at least one of which contains resistant cells. (1.60)

Both simple excision of the primary and the adjuvant regime of excision plus chemotherapy will cure a patient in state (1.58). A patient in state (1.60) will not be cured by either regime. Of special interest to us is the probability that, at the time of discovery (and removal) of the primary tumor, the patient is in state (1.59). The probability that a patient has metastases, all nonresistant, at the time of presentation, gives us an indication as to the probability that a patient will be cured by adjuvant therapy but not by simple excision of the primary tumor.

For most types of solid tumors it is a close approximation to assume that the rate of growth is exponential; that is,

$$n(t) = e^{\alpha t}. \quad (1.61)$$

Because we may describe the time axis in arbitrary units, we lose no generality by making our primary tumor growth our "clock," and hence, we may take α as equal to 1. Thus, for our purposes, we can use $n(t)$ as given simply by e^t , since for the parameter ranges considered, the amount of tumor mass removed from the primary to form the metastatic mass and/or the resistant mass is negligible (of relative mass, when compared to the primary of 1 part per 10,000). Furthermore, we assume that backward mutation from resistance to nonresistance is negligible.

Our task is to find a means of measuring the efficacy of adjuvant chemotherapy. We shall try to find estimates of the marginal improvement in the probability of cure of a solid tumor (with no apparent metastases) as a function of the size of the tumor and the parameters μ and β . We shall examine two approaches to this problem and suggest a third. We note that the problem suffers from the fact that the relatively simply Poisson axioms that describe the process go forward in time. However, to try to obtain the expression for marginal improvement in cure requires essentially a "backward" look. The first process will be an approximate one based on an argument used for a related problem of Goldie et al. [5,6]. First, let us look at the number of cells in resistant clones which develop in the primary tumor. This number, of course, is stochastic, but we shall approximate the number of resistant cells (R) by its expected value $E(R)$.

$$\frac{d(E(R))}{dt} \approx E(R) + \beta[e^t - E(R)] \quad (1.62)$$

has the solution

$$E(R) \approx e^t - e^{t(1-\beta)}. \quad (1.63)$$

Then

$$P(\text{no metastasis thrown off by resistant clone by } N) \quad (1.64) \\ \approx \exp \left[-\mu(N-1) + \frac{\mu}{1-\beta}(N^{1-\beta} - 1) \right].$$

Similarly,

$$P(\text{no metastasis which develops resistant clone by } N) \quad (1.65) \\ \approx \exp \left[-\beta(N-1) + \frac{\beta}{1-\mu}(N^{1-\mu} - 1) \right].$$

Summarizing,

$$P(\text{no resistant metastasis by total mass } N) \approx \quad (1.66) \\ \exp \left[-(\beta + \mu)(N-1) + \frac{\mu}{1-\beta}(N^{1-\beta} - 1) + \frac{\beta}{1-\mu}(N^{1-\mu} - 1) \right].$$

By differentiating (1.65) and (1.66) with respect to N , we note that for β very large relative to μ , the chances are that any formation of a resistant metastasis after tumor discovery would most likely be the result of spread from a resistant clone in the primary. Here, the standard protocol of removing the primary before beginning chemotherapy would be indicated. However, if the force of metastasis is much stronger than that of mutation to resistance, then chemotherapy might appropriately precede surgical intervention.

Now the approximation in (1.63) may be a serious underestimate for values of the stated probability less than 0.25. The reason is that by using $E(R)$ instead of R , we are allowing clones of size less than one cell to form metastases, an obvious biological absurdity. Similarly, by replacing M by $E(M)$, we are allowing metastatic clones of less than one cell to develop and be at risk to resistance.

It is possible to obtain an exact expression for the marginal improvement possible from adjuvant chemotherapy, but only as the result of a fair amount of reflection. Let us consider the two events

$A(N)$ = the event that by the time the total tumor mass equals (1.67)

N cells a nonresistant metastasis develops in which a resistant population subsequently develops before a total tumor mass of N

$B(N)$ = the event that by the time the total tumor mass equals (1.68)

N cells a resistant clone develops from which a metastasis develops before a total tumor mass of N).

We shall seek to compute $P[A^c(N)]$ and $P[B^c(N)]$. This can easily be computed by using (1.55) and (1.56).

$P(\text{metastasis occurs in } [t, t + \Delta t] \text{ followed by a resistant subclone before } T)$

$$= \mu e^t \left[1 - \exp \left(- \int_0^{T-t} \beta^{\tau} d\tau \right) \right] \Delta t = \mu e^t \left(1 - e^{\beta} e^{-\beta N/n} \right) \Delta t. \quad (1.69)$$

where $T = \ln(N)$. But then,

$$\begin{aligned} P[A^c(N)] &= \exp \left(- \int_0^T \mu e^t (1 - e^{\beta} e^{-\beta N/n}) dt \right) \\ &= \exp \left[\mu - \mu e^{\beta(1-N)} + \mu \beta e^{\beta} N \{Ei(-\beta) - Ei(-\beta N)\} \right]. \end{aligned} \quad (1.70)$$

Here, the *exponential integral* $Ei(\cdot)$ is defined for negative x by

$$Ei(x) = - \int_{-x}^{\infty} \frac{e^{-t}}{t} dt. \quad (1.71)$$

Similarly, we obtain

$$P[B^c(N)] = \exp \left[\beta - \beta e^{\mu} e^{-\mu N} + \mu \beta e^{\mu} N \{Ei(-\mu) - Ei(-\mu N)\} \right]. \quad (1.72)$$

Thus, we can compute the probability that no resistant metastases have been formed by a total tumor size of N via

$$P(\text{no resistant metastases}) = P[A^c(N)]P[B^c(N)]. \quad (1.73)$$

We have been able to obtain (1.69) and (1.71), exact solutions to the special two-stage branching processes defined by hypotheses (1.55), (1.56, and (1.57) (generally deemed intractable and requiring approximation), by exploiting the special structure of the case of interest to us here: namely, the nonappearance of any second-stage events. The exact expression in (1.72) is easily computed using, for example, the IMSL routine MMDEI [7] for computing exponential integrals. However, for the magnitudes of N (very large) and μ and β (very small) in the present application, we obtain essentially the same result [using approximations for small arguments of $Ei(x)$] with

$$P(\text{no resistant metastases by } N) = (\mu \beta \gamma^2)^{\mu \beta N}. \quad (1.74)$$

where $\gamma = 1.781072$ is e raised to Euler's constant. The probability of no metastases at all by tumor size N is given simply by

$$P(\text{no metastases by } N) = \exp \left(- \int_0^T \mu e^t dt \right) = e^{-\mu(N-1)}. \quad (1.75)$$

Thus, the probability that a patient with a tumor of size N is curable by adjuvant therapy but not by simple excision of the primary tumor alone is given by

$$P(\text{metastases, none of them resistant by } N) = (\mu\beta\gamma^2)^{\mu\beta N} - \exp[-\mu(N-1)]. \quad (1.76)$$

In Figure 1.10, we show the probability of nonoccurrence of metastases.

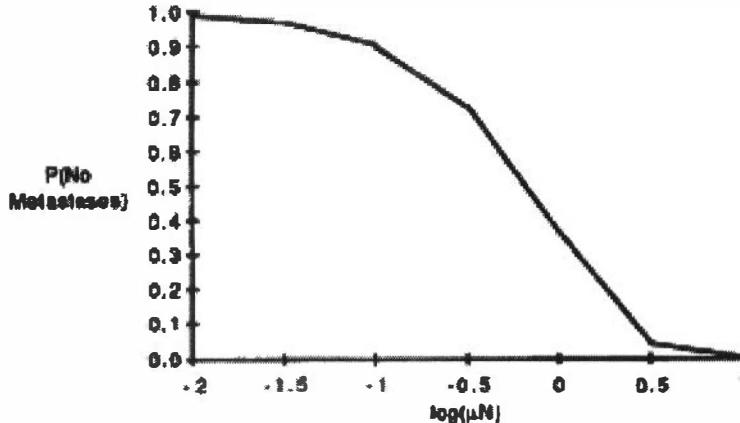


Figure 1.10. Nonoccurrence of metastases.

Typical values for β for mammalian cells are 10^{-6} to 10^{-4} . A typical tumor size at time of detection is 10^{10} cells (roughly 10 cubic centimeters).

In Figure 1.11, we show the probability of nonoccurrence of resistant metastases versus $\log(N)$ for various values of $\log(\mu\beta)$.

In Figure 1.12, we show the probability there exist metastases but no resistant metastases at a tumor mass of 10^{10} for various β and μ values. We note, for example, that for $\beta = 10^{-6.5}$ to 10^{-10} , the probability a patient presents with a condition curable by adjuvant therapy but not by simple excision of the primary is at least 40 %. Similar results hold for early detection (10^9 cells) as shown in Figure 1.13. This wide range of μ values includes that reported for breast cancer by Bartoszyński et al. [1].

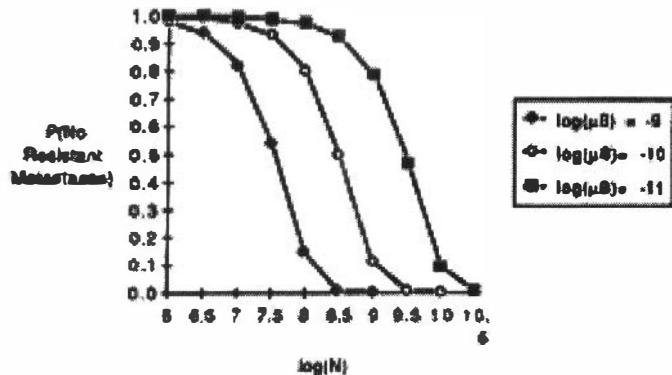


Figure 1.11. Nonoccurrence of resistant metastases.

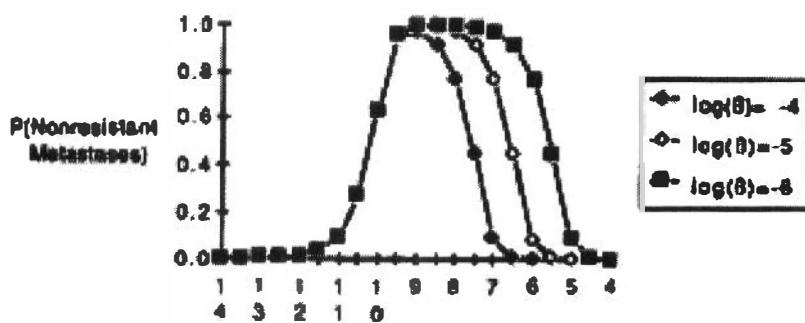


Figure 1.12. Probability of metastases but no resistant metastases when $\log(\text{tumor cell count}) = 10$.

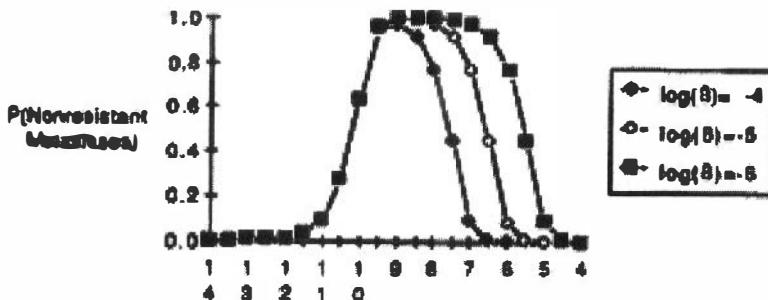


Figure 1.13. Probability of metastases but no resistant metastases when $\log(\text{tumor cell count}) = 9$.

As an aside, (1.71), and (1.72) can be used for the situation where one is interested in obtaining the probability that resistance to both of two independent chemotherapeutic agents will not have developed by tumor mass of size N . If the two parameters of resistance are β_1 and β_2 , then the probability of no doubly resistant cells is given by

$$\begin{aligned}
 & P(\text{no doubly resistant cells}) \quad (1.77) \\
 & = \exp[\beta_1 + \beta_2 - \beta_1 \exp(\beta_2 - \beta_2 N) \\
 & \quad - \beta_2 \exp(\beta_1 - \beta_1 N) + \beta_1 \beta_2 \exp(\beta_1) N \{Ei(-\beta_1) - Ei(-\beta_1 N)\} \\
 & \quad + \beta_1 \beta_2 \exp(\beta_2) N \{Ei(-\beta_2) - Ei(-\beta_2 N)\}].
 \end{aligned}$$

For typical values of N , β_1 , and β_2 , this is essentially given by

$$P(\text{no doubly resistant cells}) = (\beta_1 \beta_2 \gamma^2)^{\beta_1 \beta_2 N}. \quad (1.78)$$

In the case of breast cancer that has metastasized at least to local nodes, it has been reported by Buzdar et al. [3] that the use of adjuvant chemotherapy decreases disease mortality by as much as 54% when compared to surgery alone. In order to estimate μ and β clinically, we need randomized trials on tumors (which have exhibited no metastases at presentation) using surgery followed by adjuvant chemotherapy compared with surgery alone. Such a clinical data base is not currently available. However, since we are, at this stage, really seeking rough estimates of μ and β , animal experiments may be appropriate.

We examine below how such experiments might be used to estimate μ and β . Let us suppose we have stratified our data by tumor size at presentation. Consider the 10^{10} primary cell stratum. Suppose that the control group (surgical excision only) exhibits a cure rate of 5% and that the adjuvant therapy group exhibits a 95% cure rate. Then we can estimate μ :

$$\exp[-\hat{\mu}(N-1)] = \exp[-\hat{\mu}(10^{10} - 1)] = .05. \quad (1.79)$$

This yields a $\hat{\mu}$ value of 3×10^{-10} .

Next, we can estimate β from (178):

$$(\hat{\mu}\hat{\beta}\gamma^2)^{\hat{\mu}\hat{\beta}N} = (3 \times 10^{-10}(\hat{\beta}\gamma^2)^{3\hat{\beta}} = .95. \quad (1.80)$$

This gives (using Newton's method) $\hat{\beta} = 0.0006$. [The same estimate is obtained by the use of (1.66).]

Although we are here essentially concerned only with obtaining rough estimates of μ and β , it is clear that several resampling techniques (e.g., the jackknife or the bootstrap [4]) can be used to determine the variability of the estimates. Let us suppose, for example, that we have N_1 individuals in the control group of whom n_1 are cured and N_2 individuals in the adjuvant group of whom n_2 are cured. Using the coding of 1 for a cure and 0 for a noncure, we repeatedly sample (say, M times, where M is very large) N_1 individuals with probability of success n_1/N_1 and N_2 individuals with probability of success n_2/N_2 . For each such j th sampling we obtain $\hat{\mu}_j$ and $\hat{\beta}_j$ as above. Then we have as ready bootstrap estimates for $\text{Var}(\hat{\mu})$, $\text{Var}(\hat{\beta})$, and $\text{Cov}(\hat{\mu}, \hat{\beta})$, $\Sigma(\hat{\mu}_j - \bar{\mu})^2/M$, $\Sigma(\hat{\beta}_j - \bar{\beta})^2/M$, and $\Sigma(\hat{\mu}_j - \bar{\mu})(\hat{\beta}_j - \bar{\beta})/M$, respectively. An unacceptably large 95 % confidence ellipsoid would cause us to question the validity of our model, and the comparative suitability of competing models might be judged by examining the "tightness" of the 95 % confidence ellipsoids of each using the same data set.

In our discussion, note that both (1.65) and (1.73) are easy to compute. Equation (1.65) is much easier to think out than (1.73) but if we use (1.65), we know we are employing an approximation whose imprecision is hard to assess unless we have the more precise formula (1.73), which is time consuming to derive. It would clearly be good to have the advantage of the accuracy of (1.73) without the necessity for much effort spent in cogitation. We shall show later how we can employ simulation to go in the forward direction, pointed to by the axioms in SIMEST in Chapter 6.

Problems

1.1. Assume a society in which there is no inflation. Savings accounts pay an interest rate of 3% per year. A young professional begins investing into an IRA (tax exempt) at the end of each year the amount of \$50,000. How long must she invest before she arrives at the capital necessary to pay \$20,000 per year forever?

1.2. In 1790, Benjamin Franklin set up a two hundred year trust for the University of Pennsylvania paying 3% interest annually. The average rate of inflation for the United States since its founding has been a remarkably low (by modern standards) 1%. If the original stake was \$100,000, find (a) the value of the trust in 1990 in 1990 dollars; (b) the value of the trust in 1990 in 1790 dollars.

1.3. An individual has a stake which, in January of 2011, is worth \$10,000. Assume the interest is continuous and at the annual rate of 6.25%. Assume that all income will be taxed continuously at the 33% level, and that all after tax income will be continuously reinvested. Assume that inflation is 5%. Find, in 2011 dollars, the value of the individual's account on January 1 of 2027.

1.4. Pick two or more human populations (selected by geographical region, ethnicity, religion, or some other criterion) and examine their growth and/or decline over a period of at least one hundred years. Plot the populations versus time and/or other interesting independent variables. By argument, supported by models, discuss the pluses, incompletenesses,etc., of the Malthusian approach to the populations you are considering.

References

- [1] Bartoszyński, R., Brown, B. W., and Thompson, J. R. (1982). "Metastatic and systemic factors in neoplastic progression," in *Probability Models and Cancer*, Lucien LeCam, and Jerzy Neyman, eds.. Amsterdam: North Holland, pp. 253–264, 283–285.
- [2] Bogle, J.C. (1999). *Common Sense and Mutual Funds: New Imperatives for the Intelligent Investor*. New York: John Wiley & Sons.
- [3] Buzdar, A. U., Hortobagyi, G. N., Marcus, C. E., Smith, T. L., Martin, R., and Gehan, E. A. (1986). "Results of adjuvant chemotherapy trials in breast cancer at M.D. Anderson Hospital and Tumor Institute," *NCI Monographs*, 1, 81–85.

- [4] Efron, B. (1979). "Bootstrap methods—another look at the jackknife," *Annals of Statistics*, **7**, 1–26.
- [5] Goldie, J. H. and Coldman, A. J. (1979). "A mathematical model for relating the drug sensitivity of tumors to their spontaneous mutation rate," *Cancer Treatment Reports*, **63**, 1727–1733.
- [6] Goldie, J. H., Coldman, A. J., and Gudauskas, G. A. (1982). "Rationale for the use of alternating non-cross-resistant chemotherapy," *Cancer Treatment Reports*, **66**, 439–449.
- [7] IMSL, Inc. (1985). Subroutine MMDEI.
- [8] Jahnke, E. and Emde, F. (1945). *Tables of Functions*, New York: Dover, pp. 1–3.
- [9] Malthus, T. E. (1798). *An Essay on the Principle of Population as It Affects the Future Improvement of Society, with Remarks on the Speculations of Mr. Godwin, M. Condorcet, and Other Writers*, London.
- [10] Thompson, J. R. and Brown, B. W. (1986). A stochastic model providing a rationale for adjuvant chemotherapy, Rice Technical Report 86-19, Department of Statistics.
- [11] Thompson, J.R. and Koronacki, J. (1992). *Statistical Process Control for Quality Improvement*. New York: Chapman & Hall.

Chapter 2

Models of Competition, Survival, and Combat

2.1 Analysis of the Demographics of Ancient Israel based on Data in the Biblical Books of Exodus, Numbers, Judges, and II Samuel

Let us consider a simple model of population growth:

$$\frac{dY}{dt} = \alpha Y, \quad (2.1)$$

where Y is the size of the population at time t and α is the (constant) rate of growth. The solution is simply

$$Y(t) = Y_0 e^{\alpha t}, \quad (2.2)$$

where Y_0 is the size at $t = 0$. We note first that this model has several natural reasons for its appeal.

1. The model follows directly from the microaxiom:

$$P(\text{individual gives birth in } [t, t + \Delta t]) = \alpha \Delta t. \quad (2.3)$$

The expected increase in the interval is given simply by multiplying by Y . As Y becomes large, we can replace Y by its expectation to obtain

$$\Delta Y = Y \alpha \Delta t. \quad (2.4)$$

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

2. The model has the happy property that if we choose to divide the Y population into subgroups, which we then watch grow, when we recombine them, we get the same result as when we use pooled Y all along.

Clearly, there are bad properties of the model, too. For example:

1. The model is more appropriate for single-cell animals than for people; there is no child-bearing age range considered in the model, no allowance for the fact that only women have children, and so on. However, for large samples, this objection is academic. We could use the number of men (as we shall of necessity do thanks to the presentation of our data) assuming rough proportionality between men and women. Age stratification will also not make a noticeable difference for large populations.
2. There is no factor for diminishing growth as the people fill up the living space. However, for the population to be considered here, it would seem that death by armed violence was a much more important factor than a limit to growth factor (although it can surely be argued that death in war was a function of “crowding”).
3. There is no factor for death by war. We shall add such a factor to our model shortly.

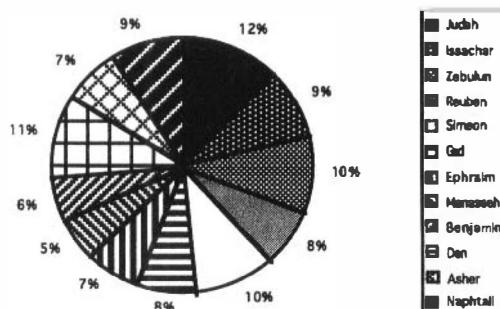
Around 1700 B.C., there were 12 male Jews fathering children. These were, of course, the sons of Jacob, his two wives (Leah and Rachel), and two concubines (Bilhah and Zilpah). The entire family disappeared into Egypt shortly thereafter. Before the conquest of Canaan, in 1452 B.C., Moses conducted a census of the total military force of the Jews. The results are shown in Table 2.1. As the Levites are always excluded from censuses, we must inflate the figure by 12/11 if we are to come up with a figure that can be used to estimate α . This gives us a comparable total of 656,433 and an estimate of $\alpha = \ln(656,433/12)/248 = 0.044$. To put this in perspective, this growth rate is roughly that of Kenya today, and Kenya has one of the fastest growing populations in the world. Certainly, it was an impressive figure by Egyptian standards and is given as the reason (Exodus 1) for the subsequent attempts (using such devices as infanticide by midwives) of a few later Pharaohs to suppress the growth in Jewish population. We note that this growth rate would seem to indicate that the bondage of Israel was of relatively short duration and that, on the whole, the Jews had lived rather well during most of their time in Egypt. This is also consistent with later writings.

Table 2.1. Moses's Census of 1452 B.C. |

Tribe	Number of Warriors
Judah	74,600
Issachar	54,400
Zebulun	57,400
Reuben	46,500
Simeon	59,300
Gad	45,650
Ephraim	40,500
Manasseh	32,200
Benjamin	35,400
Dan	62,700
Asher	41,500
Naphtali	53,400
Total	603,550

In Deuteronomy 23, there is an injunction to treat the Egyptians (along with the Edomites, descendants of Jacob's brother Esau) relatively well. They were allowed to become full participants in all forms of Jewish life after only three generations. (Contrast this with the Moabites who were not allowed membership in the congregation under any circumstances "even to their tenth generation." Obviously, the harsh Mosaic law was not always strictly observed, because King David and King Solomon had the Moabitess Ruth as their great grandmother and great great grandmother, respectively.) Also, the Jews, having gained control of Canaan, were very much inclined to enter into alliances with the Egyptians against their enemies to the north and east. The current guarded friendship between Egypt and Israel has ancient precedent.

Israeli Defense Forces (1452 B.C.)

**Figure 2.1. Total Israeli militia (1452 B.C.).**

Now had Moses had a Macintosh, he could have represented this data graphically, using a pie chart as shown in Figure 2.1 . We note that a

pie chart is somewhat appropriate, since the circular representation does not give any notion of ordering of the attributes of tribalship (i.e., it does not demand that we say that Issachar > Dan). It does give false notions of adjacency; for example, Asher is “between” Dan and Naphtali for no good reason. Less appropriate would have been a column chart (Figure 2.2) which does “order” one tribe versus another for no good reason.

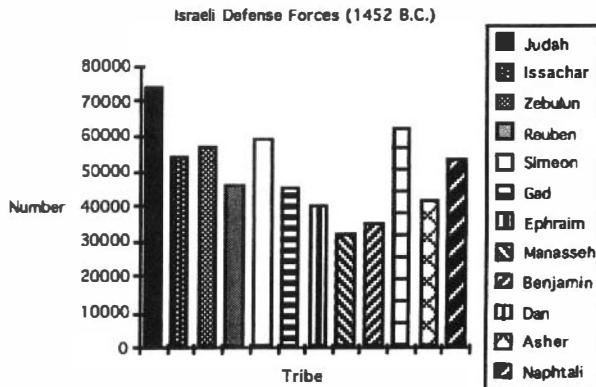


Figure 2.2. Bar chart of Israeli defense forces (1452 B.C.).

Still less appropriate would have been a line chart (Figure 2.3), which presents information more misleadingly and less informatively than the other two. (Note, however, that any one of the charts could have been constructed from either of the others.)

In truth, for the present data set, we may be better off to stick with a table, since the graph may encourage us to make subconscious inferences that are accidents of presentation rather than the result of intrinsic attributes of the data.

Returning to our attempt to make some modeling sense of the data at hand, we note (Judges 20) that after the conquest of Canaan in 1406 B.C. we have more population information at our disposal. The total warrior population had declined to 427,000. Putting in the correction to account for the Levites, we come up with a total of 465,000. Recall that before the conquest, the comparable figure was 656,000. It is obvious now that we need to include a term to account for the losses in warfare during this period of 46 years:

$$\frac{dY}{dt} = \alpha Y - \lambda, \quad (2.5)$$

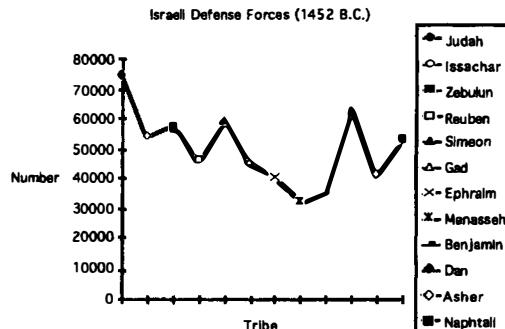


Figure 2.3. Line chart of Israeli defense forces (1452 B.C.).

where λ is the annual loss to the male population due to warfare. The solution is simply

$$Y(t) = \frac{\lambda}{\alpha} + \left(Y_0 - \frac{\lambda}{\alpha} \right) e^{\alpha t}. \quad (2.6)$$

Now the figures for the two base years of 1452 B.C. and 1406 B.C. give us a means of estimating λ via

$$\lambda = \frac{\alpha}{e^{\alpha t} - 1} (Y_0 e^{\alpha t} - Y), \quad (2.7)$$

where the time origin is taken to be 1452 B.C.

If we use our previous estimate of $\alpha = .044$, we find $\lambda = 30,144$, giving total losses of males during the 46-year period of 1,386,606. Clearly, the conquest of Canaan had caused heavy casualties among the Jews. Moreover, total war was pursued, with the Jews killing everyone: men, women, and children. Assuming their opponents retaliated in kind, the above figure might be close to 3,000,000.

Now unsettled times can be expected to lower fecundity, so perhaps the use of $\alpha = 0.044$ is unrealistic. We do have a means of estimating the growth rate between 1406 and 1017, since we have another census at that later date (the warrior age male population exclusive of Levites was 1,300,000). Unfortunately, there was intermittent, frequently very intense, warfare with the Philistines, Midianites, Ammonites, and so on, during that period, so a new value of λ is really needed. But if we swallow the later warfare casualties in α , we come up with an estimate of 0.0029. If we use this value in our equation for estimating the total casualties during the conquest of Canaan, we come up with a λ of 5777 and a total warrior age male loss figure of 265,760. This figure is no doubt too small. But translated to U.S. population values, if we go between the two extreme figures, we should be talking about losses at least as great as 30,000,000.

One observation that ought to be made here is that in 1017 B.C. David had an army of 1,300,000 after 450 years of war, during which great sacrifices had been made and borne by the Jews. The period of peace that

started about that year continued essentially until 975 B.C. After that date, the Kingdom was generally in subjugation and vassalage to some other state. What happened? Were the Jews overwhelmed by armies so massive that they could not match them?

To answer this question, we might consider some sizes of the largest field armies of the ancient world: see Table 2.2.

The notion of Israel as a numerically tiny nation, unable to cope with attacks by massive invading armies is seen to be false. At the time of the Kingdom of David (1017 till 975), Israel had a larger army than Rome had 700 years later. We know, not only from Biblical accounts but also from other records, that the Jews were one of the largest “national” groups in the ancient world. Their militia was vast [1,300,000—larger than the maximum size (750,000) ever attained by the total military forces of the Romans]. What then caused the decline (which set in immediately after Solomon)? The short answer is *high taxes*. Whether these taxes were mainly the result of costs of the Temple of Solomon or the costs for the luxuriant life-style of Solomon, we cannot answer definitively. However, we know that in 975 B.C. a delegation pleaded with Solomon’s son Rehoboam to lower the rate of taxation. While acknowledging that taxes had been high under his father (“My father has chastised you with whips . . .”) Rehoboam promised little in the way of tax relief (“But I will chastise you with scorpions.”). Upon reception of this news, the larger part of the Jews seceded from the Kingdom of Judah, setting up the Kingdom of Israel. And within these kingdoms, there then began further quarrels and divisions. Frequently, the two kingdoms would go to war with each other, inviting help from outside powers. From the time of the death of Solomon, the Kingdom of David was on a slippery slope to destruction. The total time span of the Kingdom was a scant 42 years.

Table 2.2. Field Armies of the Ancient World.

Tribe	Number of Warriors
Hannibal’s army at Cannae	32,000
Largest Assyrian army	200,000
Xerxes’ army in Greece	200,000
Alexander’s army at Issus	30,000
Scipio’s army at Zama	43,000
Rome’s total forces under Augustus	300,000

2.2 The Plague and John Graunt’s Life Table

In A.D. 1662, some 3113 years after the census of Moses, an obscure haberdasher, late a captain in the loyalist army of Charles II, published an analysis on data originally collected by Thomas Cromwell, 127 years earlier, dealing with age at the time of death in London. The data had been

collected at the request of the merchants of London who were carrying out the most basic kind of marketing research; that is whether potential customers (i.e., live people) were on the increase or decrease. Interestingly enough, the question originally arose because of the fact that the bubonic plague had been endemic in England for many years. At times, there would be an increase of the incidence of the disease, at other times a decrease. It was a matter of sufficient importance to be attended to by Chancellor Cromwell (also Master of the Rolls). Without any central data bank, a merchant might put a shop in an area where the decline in population had eliminated any potential opportunity, due to market saturation.

Cromwell's data base consisted in records of births and deaths from the Church of England to be carried out and centrally located by the clergy. Before John Graunt, all analyses of the data had suffered the usual "can't see the forest for the trees" difficulty.

Graunt solved this problem and started modern statistics by creating Table 2.3.

Graunt also tabulated $1 - F$, where F is the cumulative distribution function. We can easily use this information to graph F (see Figure 2.4).

From Graunt's table, it is an easy matter to compute the life expectancy—18 years. It would seem the plague was in full force. Note that Graunt's brilliant insight to order his data made possible a piece of information simply not available to the ancient Israeli statisticians, the sample average. How obvious Graunt's step seems to us today. Yet, it would appear he was the first to take it. It is also interesting to note how application frequently precedes theory. Graunt's sample cumulative distribution function predates any notion of a theoretical cdf.

Table 2.3. Graunt's Life Table
Age Interval P (death in interval)

0–6	.36
6–16	.24
16–26	.15
26–36	.09
36–46	.06
46–56	.04
56–66	.03
66–76	.02
76–86	.01

Following our earlier graphical analyses, we note that the pie chart is rather less informative than the bar chart, near histogram), which is slightly less useful than the near histopolygon line chart). The pie chart draws the viewer's attention to a periodicity which simply does not exist. If we divide the probabilities by the width of the age interval, we could get a true histogram as shown in (with unequal intervals).

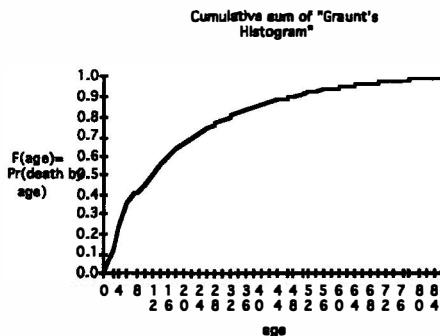


Figure 2.4. Cumulative sum of Graunt's histogram.

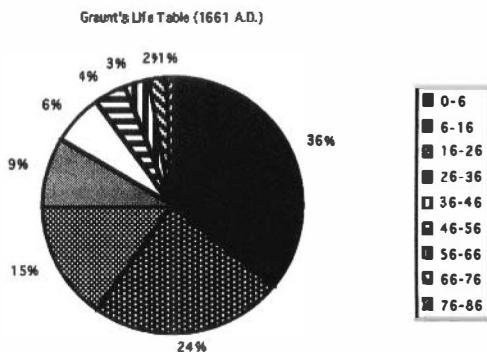


Figure 2.5. Graunt's life table (A.D. 1661).

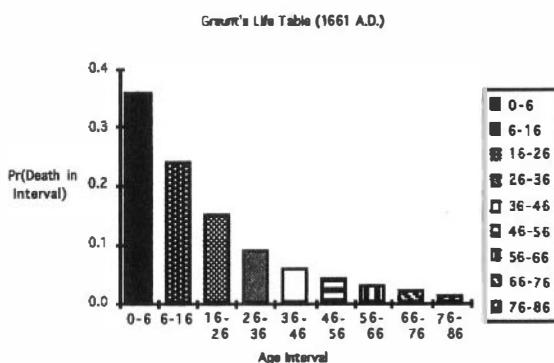


Figure 2.6. Graunt's life table (A.D. 1661).

We observe that Graunt's table cries out to be graphed, as the demographic data from ancient Israel did not. What is the difference? In the

Israeli data, there is no natural measure of proximity of tribal attributes. The covariate information is completely qualitative. Dan cannot be said, in general, to be "closer" to Naphtali than to Benjamin. A 5-year-old Englishman is very like a 7-year-old and very different from a 70-year-old. Graunt had empirically discovered a practical realization of the real number system before the real numbers were well understood. In so doing, he also presented the world with its first cumulative distribution function.

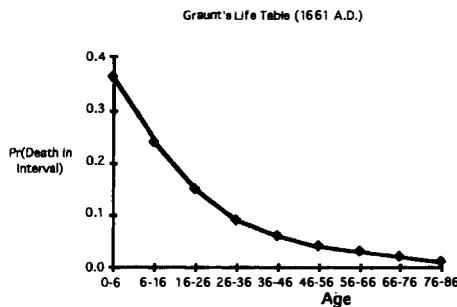


Figure 2.7. Graunt's life table (A.D. 1661).

Had he graphed his table, he might have been tempted to draw a tangent and then graph that. A pity he did not, so that a statistician could have been credited with discovering derivatives. And yet, for all the things he did not do, we must give Graunt enormous credit for what he did do.

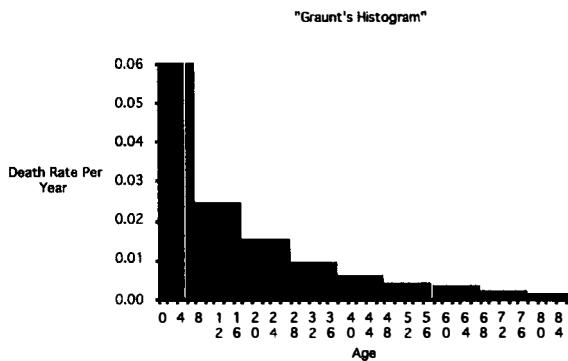


Figure 2.8. Graunt's histogram (A.D. 1661).

He brought, empirically, the notion of continuity into data analysis. By his tabulation of the cumulative distribution function, Graunt brought forth the modern science of statistics. No longer would stochastics be simply a plaything for the gentleman hobbyist. It would be the fundamental grammar of empirical science. Graunt gave us the first rationally constructed spreadsheet. As we know from Pepys' journal and other sources, Graunt

died destitute and apparently dropped from his membership (the first modern statistician had been inducted by the command of Charles II over the grumblings of other members) in the Royal Society of London.

2.3 Modular Data-Based Wargaming

Checkerboard-based games are of ancient origin, being claimed by several ancient cultures. One characteristic of these games is the restricted motion of the pieces, due to the shape and tiling of the playing field. This is overcome, in measure, in chess by giving pieces varying capabilities for motion both in direction and distance. Another characteristic of these games is their essential equality of firepower. A pawn has the same power to capture a queen as the queen to capture a pawn. Effectiveness of the various pieces is completely a function of their mobility.

The directional restrictions of square tiles are a serious detriment to checkerboard games if they are to be reasonable simulations of warfare. The most satisfactory solution, at first glance, would be to use building blocks based on circles, since such tiles would appear to allow full 360° mobility. Unfortunately, as we observe, circles cannot be satisfactory tiles, since they leave empty spaces between the tiles.

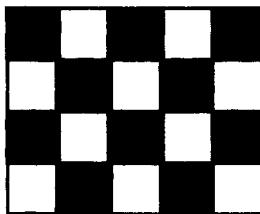


Figure 2.9. Tiling with squares.

A natural first attempt to overcome the difficulty of circles as tiles would be to use equilateral octagons, since these allow motion to the eight points of the compass, N, NE, E, SE, S, SW, W, NW. Unfortunately, as we see in Figure 2.11, this still leaves us with the empty space phenomenon.

None of the ancient games is particularly apt as an analogue of combat after the development of the longbow, let alone after the invention of gunpowder. Accordingly, the Prussian von Reiswitz began to make suitable modifications leading in 1820 to Kriegspiel. The variants of the Prussian game took to superimposing a hexagonal grid over a map of actual terrain (see Figure 2.12).

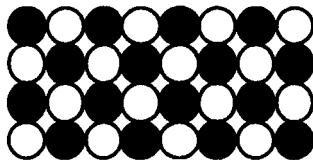


Figure 2.10. Tiling with circles.

Motion of various units was regulated by their capabilities in their particular terrain situation. The old notion of “turns” was retained, but at each turn, a player could move a number of units subject to a restriction on total move credits. Combat could be instituted by rules based on adjacency of opposing forces. The result of the combat was regulated by the total firepower of the units involved on both sides in the particular terrain situation. A roll of the dice followed by lookup in a combat table gave the casualty figures together with advance and retreat information.

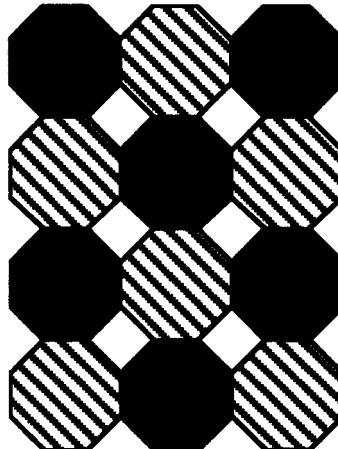


Figure 2.11. Tiling with octagons.

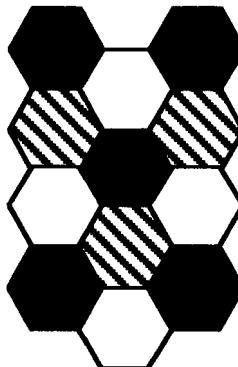


Figure 2.12. Tiling with hexagons.

The Prussian game, together with later American variants, such as Stratego, was validated against actual historical combat situations. In general, these games were excellent in their ability to simulate the real-world situation. Their major difficulty was one of bookkeeping. Frequently, a simulated combat could take longer to play than the actual historical battle. If the masking of movements and questions of intelligence gathering were included in the game, a large number of referees was required.

In attempting to take advantage of the computer, the creators of many modern military war-games have attempted to go far beyond resolution of the bookkeeping problems associated with Kriegspiel. Very frequently, these games do not allow for any interaction of human participants at all.

Initial conditions are loaded into a powerful mainframe computer, and the machine plays out the game to conclusion based on a complex program that may actually look at the pooled result of simulations of individual soldiers firing at each other, even though the combat is for very large units. Any real-time corrections for imperfections in the game are accordingly impossible. Any training potential of such games is obviously slight.

Furthermore, the creators of many of these games may disdain to engage in any validation based on historical combat results. Such validation as exists may be limited to checking with previous generations of the same game to see whether both gave the same answer.

If we know anything about artificial intelligence (and admittedly we know very little), it would seem to be that those simulations work best that seem to mimic the noncomputerized human system. Attempts to make great leaps forward without evolution from noncomputerized systems are almost always unsuccessful. And it is another characteristic of such a nonevolutionary approach that it becomes quickly difficult to check the results against realistic benchmarks. Before anyone realizes it, a new, expensive, and, very likely, sterile science will have been created soaking up time and treasure and diverting us from the real-world situation.

My own view is that it is better to use the computer as a means of alleviating the bookkeeping difficulties associated with Kriegspiel-like board

games. In the late 1970s and early 1980s, I assigned this task to various groups of students at Rice University. Experience showed that 200 person hours of work generally led to games that could emulate historical results very well.

At least another 500 person hours would have been required to make these games user friendly, but the rough versions of the games were instructive enough. One criticism made against historical validation is that technology is advancing so rapidly that any such validations are meaningless. It is claimed that the principal function of wargaming should be predictions of what will happen given the new technologies. While not agreeing that parallels between historical situations and future conflicts are irrelevant (and I note here that the strategy and tactics hobbyists generally make games ranging from Bronze Age warfare to starship troopers), I agree that the predictive aspect, in the form of scenario analyses, is very important.

Accordingly, one student created a game for conflict between an American carrier task force and a Soviet missile cruiser task force. Given the relatively close-in combat that would be likely, it seemed that if the Soviet commander is willing to sacrifice his force for the much more costly American force, he can effect an exchange of units by a massive launch of missiles at the outset of the conflict. Clearly, such a playout could have serious technological implications, for example, the desirability of constructing a system of jamming and antimissile defenses that is highly resistant to being overwhelmed by a massive strike. Or, if it is deemed that such a system could always be penetrated by further technological advances on the Soviet side, it might be appropriate to reconsider task forces based around the aircraft carrier. In any event, I personally would much prefer an interactive game in which I could see the step-by-step results of the simulation.

Also, a validation using, say, data from the Falkland conflict could be used to check modular portions of the game. World War II data could be used to check other parts. The validation would not be as thorough as one might wish, but it would be a good improvement on no validation at all. Some "supersophisticated" unvalidated computer simulation in which the computer simply played with itself and, at the end of the day, told me that existing antimissile defenses were sufficient would leave me neither comforted nor confident.

An integral part of any Kriegspiel computerization should deal with the resolution of the likely results of a conflict. A ready means of carrying this out was made via the famous World War I opus of Lanchester. Let us suppose that there are two forces, the Blue and the Red, each homogeneous, and with sizes u and v , respectively.

Then, if the fire of the Red force is directed, the probability a particular Red combatant will eliminate some Blue combatant in time interval $[t, t + \Delta t]$ is given simply by

$$P(\text{Blue combatant eliminated in } [t, t + \Delta t]) = c_1 \Delta t, \quad (2.8)$$

where c_1 is the Red coefficient of directed fire. If we wish then to obtain the total number of Blue combatants eliminated by the entire Red side in $[t, t + \Delta t]$, we simply multiply by the number of Red combatants to obtain

$$E(\text{change in Blue in } [t, t + \Delta t]) = -vc_1\Delta t. \quad (2.9)$$

Replacing u by its expectation (as we have the right to do in many cases where the coefficient is truly a constant and v and u are large), we have

$$\frac{\Delta u}{\Delta t} = -c_1 v. \quad (2.10)$$

This gives us immediately the differential equation

$$\frac{du}{dt} = -c_1 v. \quad (2.11)$$

Similarly, we have for the Red side

$$\frac{dv}{dt} = -c_2 u. \quad (2.12)$$

This system has the time solution

$$u(t) = u_0 \cosh \sqrt{c_1 c_2} t - v_0 \sqrt{\frac{c_1}{c_2}} \sinh \sqrt{c_1 c_2} t. \quad (2.13)$$

$$v(t) = v_0 \cosh \sqrt{c_1 c_2} t - u_0 \sqrt{\frac{c_2}{c_1}} \sinh \sqrt{c_1 c_2} t. \quad (2.14)$$

A more common representation of the solution is obtained by dividing (2.11) by (2.12) to obtain

$$\frac{du}{dv} = \frac{c_1 v}{c_2 u}, \quad (2.15)$$

with solution

$$u^2 - u_0^2 = \frac{c_1}{c_2} (v^2 - v_0^2). \quad (2.16)$$

Now u and v are at “combat parity” with each other when

$$u^2 = \frac{c_1}{c_2} v^2. \quad (2.17)$$

A special point needs to be made here. Such parity models assume that both sides are willing to bear the same proportion of losses. If such is not the case, then an otherwise less numerous and less effective force can still emerge victorious. For example, suppose that the Blue force versus the Red force coefficient is 0.5 and the Blue force has only 0.9 the numerosity of the Red force. Then if Blue is willing to fight until reduced to 0.2 of his original strength, but Red will fight only to 0.8 of his original strength,

then using (2.17), we find that by the time Red has reached maximal acceptable losses, Blue still has 30% of his forces and thus wins the conflict. This advantage to one force to accept higher attrition than his opponent is frequently overlooked in wargame analysis. The empirical realization of this fact has not escaped the attention of guerrilla leaders from the Maccabees to the Mujaheddin.

Accordingly, it is interesting to note that if there is a doubling of numbers on the Red side, Blue can only maintain parity by seeing to it that c_2/c_1 is quadrupled, a seemingly impossible task.

Lanchester's formula for undirected fire follows from similar Poissonian arguments. The probability that a Red combatant will eliminate some Blue combatant in $[t, t + \Delta t]$ is given by

$$\begin{aligned} P(\text{a Blue eliminated by a Red in } [t, t + \Delta t]) \\ = P(\text{shot fired in } [t, t + \Delta t]) P(\text{shot hits a Blue}) \Delta t. \end{aligned} \quad (2.18)$$

Now, the probability that a shot aimed at an area rather than at an individual hits someone is proportional to the density of Blue combatants in the area, and hence proportional to u . Thus, we have

$$P(\text{Blue eliminated in } [t, t + \Delta t]) = d_1 u \Delta t. \quad (2.19)$$

The expected number of Blues eliminated in the interval is given by multiplying the above by the size of the Red force, namely, v . So the differential equations are

$$\frac{du}{dt} = -d_1 u v \quad \text{and} \quad \frac{dv}{dt} = -d_2 u v. \quad (2.20)$$

This system has the time solution

$$u(t) = \frac{(d_2/d_1)u_0 - v_0}{(d_2/d_1) - (v_0/u_0)e^{-(d_2u_0-d_1v_0)t}}. \quad (2.21)$$

and

$$v(t) = \frac{(d_1/d_2)v_0 - u_0}{(d_1/d_2) - (u_0/v_0)e^{-(d_1v_0-d_2u_0)t}}. \quad (2.22)$$

Here, when dividing the equations and solving, we obtain the parity equation:

$$u - u_0 = \frac{d_1}{d_2}(v - v_0). \quad (2.23)$$

In such a case, a doubling of Red's parity force can be matched by Blue's doubling of d_2/d_1 .

In attempting to match either law (or some other) against historical data, one needs to be a bit careful. In 1954, Engel claimed to have validated the

applicability of Lanchester's directed fire law for the Battle of Iwo Jima. He used no records for Japanese casualties and simply adjusted the two parameters to fit the record of American casualty data. According to Engel's model, Japanese resistance was completely eliminated in 36 days. But American data reveal that resistance continued for some time after the battle was over, with 20 Japanese killed in a single banzai charge on day 37 and up to 70 prisoners taken on some days up to 1 month after day 36. More significantly, there is available partial Japanese force data delivered by the Japanese commander, General Kuribayashi, by radio to Tokyo at times well before day 36. For example, on day 21 of the conflict, when Engel's model gives a figure of 8550 for the Japanese forces on the island, Kuribayashi gives the actual figure of 1500. Using the partial Japanese casualty records shows that the directed fire model gave answers much at variance with the data (sometimes off the Japanese total effectives by a factor of 4) and that the undirected fire model appeared to work much more satisfactorily. It is possible to track very closely the American force levels using either the directed or undirected fire models. But the undirected fire model has the additional attribute of closely tracking the partial force information for the Japanese. We have exhibited both the directed and undirected fire models above in Figure 2.13.

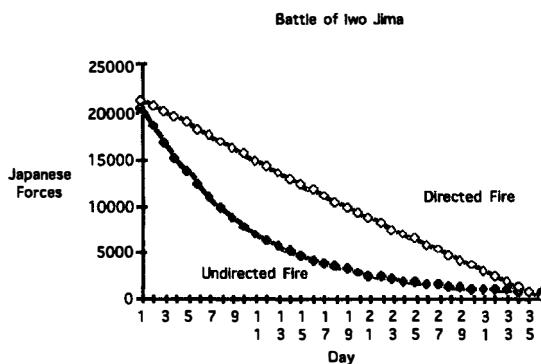


Figure 2.13a. Battle of Iwo Jima with 21,500 defenders.

Suppose, however that the Japanese had had 50,000 defenders. We note, in this case, a catastrophe for the Americans. One can only wonder at the excellence of U.S. Naval Intelligence which showed the Japanese had only around half that number. One gets the impression that, time and again, U.S. Naval Intelligence gets things right more often than other United States intelligence services.

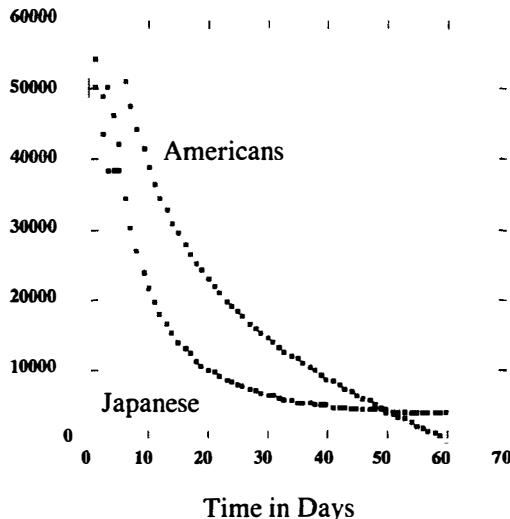


Figure 2.13b. Battle of Iwo Jima with 50,000 defenders.

In the Iwo Jima scenario, considering the losses of the Japanese forces, it is rather clear that the undirected fire model is to be preferred over the directed one. In the case of the American attrition, the directed fire model is the more appropriate. However, any homogeneous force model would probably not be as satisfactory as a heterogeneous force model in an engagement in which naval gunfire together with marine assault both played important roles. We shall address the heterogeneous force model problem shortly. In a much broader context of combat simulation, we note that a model which appears at first glance to do an excellent job of “prediction” may become seriously deficient as more data are made available.

The year 2011 marked the 175th anniversary of the Battle of the Alamo. This battle gives an example of a situation in which a mixture of the two models is appropriate. Since the Texans were aiming at a multiplicity of Mexican targets and using rifles capable of accuracy at long range (300 m), it might be appropriate to use the directed fire model for Mexican casualties. Since the Mexicans were using less accurate muskets (100 m) and firing against a fortified enemy, it might be appropriate to use the undirected fire model for Texan casualties. This would give

$$\frac{du}{dt} = -d_1uv \quad \text{and} \quad \frac{dv}{dt} = -c_2u. \quad (2.24)$$

The parity equation is given by

$$v^2 - v_0^2 = \frac{2c_2}{d_1}(u - u_0). \quad (2.25)$$

The Texans fought 188 men, all of whom perished in the defense. The Mexicans fought 3000 men of whom 1500 perished in the attack. By plugging in initial and final strength conditions, it is an easy matter to compute $c_2/d_1 = 17,952$. However, such an index is essentially meaningless, since the equations of combat are dramatically different for the two sides. A fair measure of man for man Texan versus Mexican effectiveness is given by

$$\frac{\frac{1}{u} \frac{du}{dt}}{\frac{1}{v} \frac{dv}{dt}} = \frac{c_2}{d_1 u}. \quad (2.26)$$

This index computes the rate of destruction of Mexicans per Texan divided by the rate of destruction of Texans per Mexican. We note that the mixed law model gives a varying rate of effectiveness, depending on the number of Mexicans present. At the beginning of the conflict, the effectiveness ratio is a possible 96; at the end, a romantic but unrealistic 17,952.

The examination of this model in the light of historical data should cause us to question it. What is wrong? Most of the Mexican casualties occurred before the walls were breached. Most of the Texan casualties occurred after the walls were breached. But after the walls were breached, the Mexicans would be using directed fire against the Texans.

We have no precise data to verify such an assumption, but for the sake of argument, let us assume that the Texans had 100 men when the walls were breached, the Mexicans 1800. Then (2.26) gives $c_2/d_1 = 32,727$. The combat effectiveness ratio $c_2/d_1 u$ goes then from 174 at the beginning of the siege to 327 at the time the walls were breached. For the balance of the conflict we must use (2.21) and (2.22) with the combat effectiveness ratio $c_2/c_1 = 99$. Personally, I am not uncomfortable with these figures. The defenses seem to have given the Texans a marginal advantage of around 3. Those who consider the figures too "John Wayneish" should remember that the Mexicans had great difficulty in focusing their forces against the Alamo, whereas the Texans were essentially all gainfully employed in the business of fighting. This advantage to a group of determined Palikari to defend a fortified position against overwhelming numbers of a besieging enemy is something we shall return to shortly.

Another example of the effect of fortifications in combat is obtained from the British-Zulu War of 1879. On January 22, at Isandhlwana, an encamped British column of 1800 British soldiers and 1000 native levies was attacked by 10,000 warriors of the Zulu King Cetawayo. The suggestion of the Afrikaaner scouts to laager (roughly drawing the wagons into a circle) was rejected by the British commander. Consequently, even though the British troops had the benefit of modern breech-loading rifles, they were quickly engaged in hand to hand combat by the Zulus. The result of the conflict was that only around 55 British and 300 of the native levies survived. We do not have precise knowledge of the Zulu losses; but we do know that, on the evening of January 22, 5000 of the same Zulu force attacked a British force of 85 at Rorke's Drift. The British commander (a lieutenant of

engineers, John Chard, later a colonel) had used the few hours warning to laager his camp with overturned wagons and sacks of meal. On January 23, the Zulus withdrew, leaving 400 dead on the field. British losses were 17 killed and 10 seriously wounded. Here we have an example of nearly identical types of forces on the attack and on the defense in both engagements. Since the Zulus always fought hand-to-hand, we shall use (2.24) in both battles for both sides. If we assume (a popular notion of the day) that the native levies made no contribution, and that 5000 Zulus were incapacitated by the Isandhlwana engagement, the combat effectiveness of British soldier versus that of Zulu soldier computes to be 23.17. (The assumptions here obviously tend hugely to inflate the actual British versus Zulu combat effectiveness.) At Rorke's Drift, the combat effectiveness ratio goes to 994.56. Thus, the advantage given to the British defenders of Rorke's Drift by the hastily constructed defenses was at least $994.56/23.17 = 42.92$. The advantage was not primarily an increased combat effectiveness of the British soldiers, but rather a diminution of the combat effectiveness of the Zulus. Having transmitted some feeling as to the advantages of commonsense utilization of the method of Lanchester (borrowed in spirit from Poisson), we shall now take the next step in its explication: namely, the utilization of heterogeneous force equations.

Let us suppose that the Blue side has m subforces $\{u_j\}; j = 1, 2, \dots, m$. These might represent artillery, infantry, armor. Then, we have

$$\frac{du_j}{dt} = - \sum_{i=1}^n k_{ij} c_{1ij} v_i \quad (2.27)$$

and

$$\frac{dv_i}{dt} = - \sum_{j=1}^m l_{ji} c_{2ji} u_j. \quad (2.28)$$

Here, k_{ij} represents the allocation (a number between 0 and 1 such that $\sum_{j=1}^m k_{ij} \leq 1$) of the i th

Red subforce's firepower against the j th Blue subforce. c_{1ij} represents the Lanchester attrition coefficient of the i th Red subforce against the j th Blue subforce. Similar obvious definitions hold for $\{l_{ji}\}$ and $\{c_{2ji}\}$.

Beyond the very real utility of the Lanchester combat laws to describe the combat mode for war-games, they can be used as a model framework to gain insights as to the wisdom or lack thereof of proposed changes in defense policy. For example, a dismantling of intermediate range missiles in eastern and central Europe throws additional responsibility on the effectiveness of NATO conventional forces, since a conventional Soviet attack is no longer confronted with a high risk of a Pershing missile attack from West Germany against Russia. The rather larger disparity in conventional forces between the Soviet block and NATO forces can roughly be addressed by a consideration of Lanchester's directed fire model. As we have observed in

(2.26), in the face of a twofold personnel increase of Red beyond the parity level, Blue can, assuming Lanchester's directed fire model, maintain parity only by quadrupling c_2/c_1 . This has usually been perceived to imply that NATO must rely on its superior technology to match the Soviet threat by keeping c_2 always much bigger than c_1 .

Since there exists evidence to suggest that such technological superiority does not exist at the conventional level, it appears that the Soviets kept out of western Europe because of a fear that a conventional juggernaut across western Europe would be met by a tactical nuclear response from West Germany, possibly followed, *in extremis*, by a strategic attack against population centers in Russia: thus, the big push by the Soviets and their surrogates for "non first use of nuclear weapons" treaties and their enthusiastic acceptance of American initiatives to remove intermediate range missiles from eastern and central Europe. It is not at all unlikely that the Soviets could have taken western Europe in a conventional war in the absence of intermediate range missiles in West Germany if the current disparate numerical advantage of conventional Soviet forces in Europe were maintained. Thus, one practical consequence for NATO of the dismantling of intermediate range missiles in Europe might be attempts by the Western powers to bring their conventional forces to numerical parity with those of the Soviets. This might require politically difficult policy changes in some NATO countries, such as reinstitution of the draft in the United States.

In my paper, "An Argument for Fortified Defense in Western Europe" [9], I attempted to show how the c_2/c_1 ratio could be increased by using fortifications to decrease c_1 . Whether or not the reader judges such a strategy to be patently absurd, it is instructive to go through the argument as a means of explicating the power of Lanchester's laws in scenario analysis.

My investigation was motivated, in part, by the defense of the Westerplatte peninsula in Gdańsk by 188 Polish soldiers from September 1 through September 7 in 1939, and some interesting parallels with the much lower tech siege of the Alamo 100 years earlier. (Coincidentally, the number of Polish defenders was the same as the number of Texans at the Alamo.) The attacking German forces included a battalion of SS, a battalion of engineers, a company of marines, a construction battalion, a company of coastal troops, assorted police units, 25 *Stukas*, the artillery of the Battleship Schleswig-Holstein, eight 150 mm howitzers, four 210 mm heavy mortars, 100 machine guns, and two trainloads of gasoline (the Germans tried to flood the bunkers with burning gasoline).

The total number of German troops engaged in combat during the 7-day seige was well over 3000. Anyone who has visited Westerplatte (as I have) is amazed with the lack of natural defenses. It looks like a nice place for a walkover. It was not.

The garrison was defended on the first day by a steel fence (which the Germans and the League of Nations had allowed, accepting the excuse of the Polish commander, Major Sucharski, that the fence was necessary to

keep the livestock of the garrison from wandering into Gdańsk), which was quickly obliterated. Mainly, however, the structural defenses consisted in concrete fortifications constructed at the ground level and below. Theoretically, the structural fortifications did not exist, because they were prohibited by the League of Nations and the peninsula was regularly inspected by the Germans to ensure compliance. However, extensive "coal and storage cellars" were permitted, and it was such that comprised the fortifications. The most essential part of the defenses was the contingent of men there. Unlike the Texans at the Alamo who realized they were going to die only after reinforcements from Goliad failed to arrive and the decision was made not to break through Santa Anna's encirclement, the Polish defenders of Westerplatte realized, long before the conflict, that when the German invasion began, they would be doomed. It is interesting to note the keen competition that existed to gain the supreme honor of a posting to Westerplatte. Perhaps "no bastard ever won a war by dying for his country" but the defenders of the Alamo and those of Westerplatte consciously chose their deaths as an acceptable price for wreaking a bloody vengeance on the enemies of their people.

Ever since the abysmal failure of the Maginot Line in 1940, it has been taken for granted that any strategy based on even the partial use of fixed defenses is absurd. I question this view. Historically, fixed defenses have proved more effective as islands rather than as flankable dikes. The Maginot Line was clearly designed as a dike, as was the Great Wall of China, and both proved failures. It is unfortunate that the dike-like tactics of trench warfare had proved so effective in World War I. Otherwise, the French would undoubtedly have noted that they were basing their 1940 defense on a historically fragile strategy. Dikes generally can withstand force only from the front, as the Persians (finally) discovered at Thermopolae. If the dikes are sufficiently narrow and thick, however, they may be effective islands and very difficult to outflank. It was conceded by the panzer innovator, von Manstein, that Germany absolutely could not have taken the Sudetenland defenses in 1938 had they been used. This brings up another interesting point. An effective system of fixed defenses is very much dependent on the will of the people using them.

Historical examples, modern as well as ancient, of successful use of constructed defensive positions can be given ad infinitum. Among the crusading orders, the Templars and Hospitalers early discovered that they could maintain an effective Christian presence in the Holy Land only by concentrating a large percentage of their forces in a number of strongly fortified castles. This gave them sufficient nuisance value to cause concessions by the Muslim leaders. Most of the military disasters to the orders were the result of their willingness to strip their castle defenses and join the crusader barons in massive land battles against numerically overwhelming odds—as at Hattin. For more than 1000 years, some of the Christian peoples in the Near East, for example, the Armenians and the Maronites, maintained

their very identity by mountain fortifications.

It is interesting to note that one crusader fortress—Malta—never fell to the Muslims and was only taken (by treachery) by Napoleon in 1798. In the second World War, the connection between the resistance of Malta and the ultimate destruction of the Afrika Korps is well remembered. Even light, hastily constructed defenses, manned by people who do not know they are supposed to surrender when surrounded, can be extremely effective in slowing down the enemy advance, as proved by the 101st Airborne during the Battle of the Bulge. In the examples above, there seem to be some common points. First, fortified defense gives a ready means of increasing the ratio of the Lanchester coefficients in favor of the Blue side. One natural advantage to this type of defense is the fact that the defender can increase his Lanchester attrition ratio by a policy of construction over a period of time. This may be a more fruitful policy than placing all one's hopes on increasing ones Lanchester ratio by the design of new weapons systems.

Second, fortified defense should rely on adequate stores of supplies located within the "fortress perimeter." It should be assumed by the defenders that they will be completely surrounded by the enemy for long periods of time. (In their fortress at Magdeburg, the Teutonic Knights always kept 10 years' provisions for men and horses.)

Third, fortified defense is a task best undertaken by well-trained professionals with strong group loyalty.

Fourth, fortified defense is most effective when there are allied armies poised to strike the enemy at some future time and place. The fortress and the mobile striking force complement each other in their functions. The function of the fortress is to punish, harass, and divide the enemy and to maintain a presence in a particular area. In general, however, offensive activities must be left to the mobile forces. The deployment of enemy forces to take fortified positions will weaken their ability to withstand mobile offensive operations. Let us now examine modified versions of (2.21) and (2.22):

$$\frac{du}{dt} = -c_1^* v \quad (2.29)$$

and

$$\frac{dv}{dt} = -c_2^* u. \quad (2.30)$$

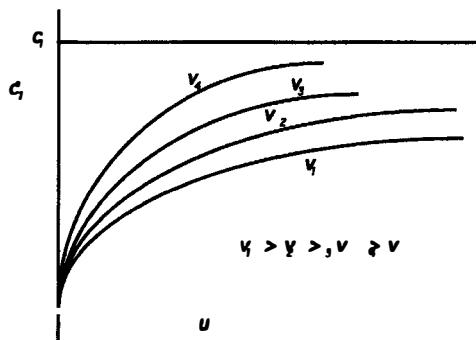


Figure 2.14. Lanchester combat.

The attrition to Blue coefficient is represented by the variable $c_1^* = c_1^*(u, v)$ and is demonstrated graphically in Figure 2.14. In the above, we assume that c_1^* never exceeds c_1 , the attrition constant corresponding to nonfortified combat. Clearly, the functions c_1^* and c_2^* are functions of the manner in which the fortress has been constructed. It may be desirable to design the fortifications so that c_1^* is small, even at the expense of decreasing c_2^* . Generally, one might assume that c_2^* is close to the nonfortified attrition rate of u against v , since the defenders will have removed potential cover for the Red side. In fortress defense, the solution in time is likely to be important, since a primary objective is to maintain a Blue presence for as long as possible. Next, consider a linear approximation to the v -level curves of $c_1^*(u, v)$ Figure 2.15. Then we would have

$$\frac{du}{dt} = -g(v)uv - c_1^{**}u, \quad (2.31)$$

where $c_1^*(u, v) = g(v)u$ and c_1^{**} is the Blue coefficient of internal attrition. (We note that this analysis has moved us, quite naturally, to an undirected fire model for the defenders' losses. The model thus derived is essentially that used earlier for the Alamo.) We might reasonably expect that the besieging forces would maintain more or less a constant number of troops in the vicinity of the redoubt. Hence, we would expect

$$\frac{dv}{dt} = -c_2^*u - c_2^{**}v + P(u, v) = 0, \quad (2.32)$$

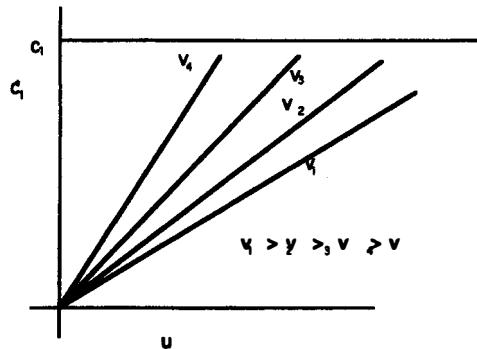


Figure 2.15. Linear approximations to Lanchester combat.

where $P(u, v)$ is the rate of replacement necessary to maintain constant v strength and c_1^{**} is the Red coefficient of internal attrition. We might expect that $c_2^{**} \gg c_1^{**}$, since inadvertent self-inflicted casualties are a well-known problem for the besieging force. Then

$$u(t) = u_0 \exp[-(g(v)v + c_1^{**})t]. \quad (2.33)$$

The enemy attrition by time t is given by

$$\int_0^t P(u, v) dt = c_2^{**}tv + c_2^{*}u_0 \frac{1 - \exp[-(g(v)v + c_1^{**})t]}{g(v)v + c_1^{**}}. \quad (2.34)$$

If the Blue defense can hold out until $u = \gamma u_0$ (where $0 < \gamma < 1$), then the time until the end of resistance is given by

$$t^* = \frac{\ln(\gamma)}{g(v)v + c_1^{**}}. \quad (2.35)$$

We have then that the total losses to the Red side by the time the defense falls are given by

$$\frac{c_2^{*}u_0(1 - \gamma) - c_2^{**}v \ln(\gamma)}{g(v)v + c_1^{**}}. \quad (2.36)$$

It is interesting to note that if $c_2^{**} = 0$, then the minimization of Red casualties seems to be consistent with the minimization of t^* . This might indicate that an optimum strategy for Red is to overwhelm the Blue fortifications by sheer weight of numbers. This would not be true if beyond some value of v , $d(g(v)v)/dv \leq 0$, implying that beyond a certain strength, additional Red forces would actually impair Red's ability to inflict casualties on the Blue side. As a matter of fact, the history of fortified defense seems to indicate that such a "beginning of negative returns" point in the v space does exist. It is generally the case for the besieging force that $c_2^{**} \gg 0$ and that it is increasing in v . This is particularly true if the besieged forces are able from time to time to conduct carefully planned "surprises"

to encourage increased confusion and trigger happiness on the part of the besiegers.

In the heterogeneous force model for fortified defense, we have

$$\frac{du_j}{dt} = - \sum_{i=1}^n k_{ij} g_{ij}(v_i) v_i u_j - c_{1j}^{**} u_j \quad (2.37)$$

and

$$\frac{dv_i}{dt} = - \sum_{j=1}^m l_{ji} c_{2ji}^* u_j - c_{2i}^{**} v_i. \quad (2.38)$$

The size of the j th Blue subforce at time t is given by

$$u_j(t) = u_j(0) \exp \left(-t \sum_{i=1}^n k_{ij} g_{ij}(v_i) v_i + c_{ij}^{**} \right). \quad (2.39)$$

The total attrition to the i th enemy subforce at time t is given by

$$\begin{aligned} & \int_0^t P_i(u, v) d\tau \\ &= \sum_{j=1}^m l_{ji} c_{2ji}^* u_j(0) \int_0^t \exp \left(-\tau \sum_{i=1}^n k_{ij} g_{ij}(v_i) v_i + c_{ij}^{**} \right) d\tau + c_{2i}^{**} t v_i \\ &= \sum_{j=1}^m l_{ji} c_{2ji}^* u_j(0) \frac{1 - \exp \left(-t \sum_{i=1}^n k_{ij} g_{ij}(v_i) v_i \right)}{\sum_{i=1}^n k_{ij} g_{ij}(v_i) v_i} + c_{2i}^{**} t v_i. \end{aligned} \quad (2.40)$$

Suppose that the effectiveness (at time t) of the Blue defender is measured by

$$T(t) = \sum_{j=1}^m a_j u_j(t), \quad (2.41)$$

where the a_j are predetermined relative effectiveness constants. If we assume that the fortress is lost when the effectiveness is reduced to some fraction γ of its initial value, that is, when

$$T(t) < \gamma T(0), \quad (2.42)$$

then we can use (2.42), in straightforward fashion, to solve for the time of capture.

This model gives some indication of the power of the simple Lanchester "laws" in analyzing a "what if" scenario. It is, in large measure, the lack of "gee-whizziness" of Lanchester's models that renders them such a useful

device to the applied worker. Generally, after a few hours of self-instruction, potential users can reach the level of sophistication where they can flowchart their own war-games or other forms of scenario analysis.

2.3.1 Herman Kahn and the Winning of the Cold War

John Tukey once told the author that he considered Herman Kahn to be the world leader and pioneer in Monte Carlo simulation. Kahn turned his attention in the 1950s to strategies for dealing with Soviet Russia. Kahn was a model for effective big picture statisticians. His book *On Thermonuclear War* was required reading for members of President Kennedy's cabinet. Kahn invented the *escalation ladder*, the use of which was generally agreed upon both by the United States and Russia. It was a discrete multi-step upward and downward progression of hostilities. Its use gave both the great powers time to think before rushing precipitously to full scale nuclear war. The escalation ladder alone was worth a great deal.

But Kahn went much further. He noted that in the Soviet Union nothing worked very well except the military. Even before the Communists seized Russia, it would be a fair analysis to observe that Russia was the successor to the Mongol Golden Horde, and that it lived by conquest. Its expansion rate was equivalent to a Belgium-sized country per year. There really was very little indication that Russia was a normal country. It would either continue to be a gigantic kleptocracy or its historic model would simply fail.

The Leonid Brezhnev strategy was straightforward. At a time of its own choosing, led by 50,000 tanks, the Soviets start a march to the Rhine. Because of the logistic advantages of shorter supply lines than those of the Americans, by bloody attrition, the Russians drive the Americans to the point of being conventionally overwhelmed. At this point, the Americans are driven to the use of battlefield nuclear weapons. At this point, Brezhnev calls the West German government and notes that their country is about to become a nuclear landscape. He suggests the Germans leave NATO and order the Americans to leave West Germany. An arrangement is achieved whereby the countries east of the Rhine become quasi-independent satellites, which pay Russia annual tribute.

Herman Kahn understood well that statistics was not only about crunching numbers. It was about logical conclusions based on facts. And these facts might be historical precedent. Kahn died in the earliest days of the Reagan Administration. But his books lived past his lifespan. Reagan's personal library reveals that Reagan read and reread them all. In coordination with German Chancellor Helmut Kohl, Reagan placed Pershing II missile sites in West Germany. These missile sites were an innovation. Earlier, none of the sites in West Germany contained nuclear weapons which could reach Moscow or Krasnoyarsk. Thus, an American nuclear attack would have to be launched by SAC, American nuclear submarines, Minute-

man sites in America and so on. In other words, if NATO escalated, full scale nuclear war would result.

However, Kahn proposed a distributed system. If the Soviets launched a panzer blitz, then when they would get within 10 minutes of a U.S. Pershing silo near Kleindorf, the American commander Colonel Miller would know he had 20 minutes to live. He also would know that his family in Kleindorf would have 35 minutes to live. His German counterpart, Oberst Miller, would be in the same situation. It is a capital offense for a NATO commander to launch a missile strike without NATO authorization. But what will Colonel Miller do? Nobody knows. Facing such a situation, perhaps Brezhnev would simply continue his attack. But not Gorbachev.

Thus, after the Pershing II's are installed, Helmut Kohl offers Gorbachev \$60 billion to remove the Red Army from East Germany. Gorbachev, desperate for cash, agrees, assuming this is the beginning of perpetual tribute from the West.

On June 4, 1989, a Solidarnosc government is formed in Poland. On November 9, 1989, the Berlin Wall falls. On December 12, 1991, the Soviet Union dissolves.

2.4 Predation and Immune Response Systems

Let us consider Volterra's predator-prey model and some consequences for modeling the human body's anticancer immune response system. For the classical shark-fish model, we follow essentially Haberman [5]. Suppose we have predators, say sharks, whose numbers are indicated by S , who prey on, say fish, whose numbers are indicated by F . In the 1920s, it was brought to the attention of Volterra that there seemed to be a periodic pattern in the abundance of certain food fish in the Adriatic, and that this pattern did not seem to be simply seasonal. Volterra attempted to come up with the simplest logical explanation of this periodicity.

We might suppose that the probability a typical shark gives birth to another shark (for reasons of simplicity, we treat the sharks as though they were single-cell creatures) is given by

$$P(\text{birth in } [t, t + \Delta t]) = (\lambda F)\Delta t. \quad (2.43)$$

Here the assumption is that the probability of reproduction is proportional to the food supply, that is, to the size of the fish population. The probability a shark dies in the time interval is considered to be a constant $k \Delta t$. Thus, the expected change in the predator population during $[t, t + \Delta t]$ is given by

$$E[\Delta S] = S(\lambda F - k)\Delta t. \quad (2.44)$$

As we have in the past, we shall assume that for a sufficiently large predator population, we may treat the expectation as essentially deterministic.

This gives us the differential equation

$$\frac{dS}{dt} = S(\lambda F - k). \quad (2.45)$$

Similarly, the probability that a given fish will reproduce in $[t, t + \Delta t]$ minus the probability it will die from natural causes may be treated like

$$P(\text{birth in } [t, t + \Delta t]) = a\Delta t. \quad (2.46)$$

We have assumed that the fish have essentially an unlimited food supply. The death by predation, on a per fish basis, is obviously the number of sharks multiplied by their fish eating rate, c , giving the differential equation

$$\frac{dF}{dt} = F(a - cS). \quad (2.47)$$

Now the system of Volterra equations given by has no known simple time domain solution, although numerical solution is obviously trivial. However, let us examine the F versus S situation by dividing (2.47) by (2.45). This gives us

$$\frac{dF}{dS} = \frac{F}{\lambda F - k} \frac{a - cS}{S}. \quad (2.48)$$

The solution is easily seen to be

$$F^{-k} e^{\lambda F} = E e^{-cS} S^a, \quad (2.49)$$

with E a constant. Now, let us use trace the path of F versus S . We note first that $F = k/\lambda$ gives an unchanging S population; $S = a/c$ gives an unchanging F population.

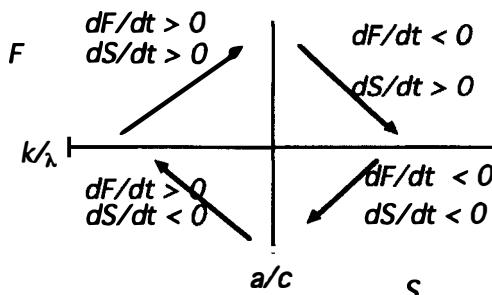


Figure 2.16. Linear Volterra population plot.

The consequences of Figure 2.16 are that the F versus S plot must either be a closed repeating curve or a spiral. We can use (2.49) to eliminate the possibility of a spiral. Let us examine the level curves of F and S corresponding to the common Z values in

$$F^{-k} e^{\lambda F} = E e^{-cS} S^a = Z. \quad (2.50)$$

In Figure 2.17, we sketch the shapes of Z versus F , respectively, and use these values to trace the F versus S curve. We note that since each value of Z corresponds to at most four points on the F versus S curve, a spiral structure is out of the question, so we obtain the kind of closed curve that was consistent with the rough data presented to Volterra. Using Figure 2.17 leading the shark curve by 90 degrees.

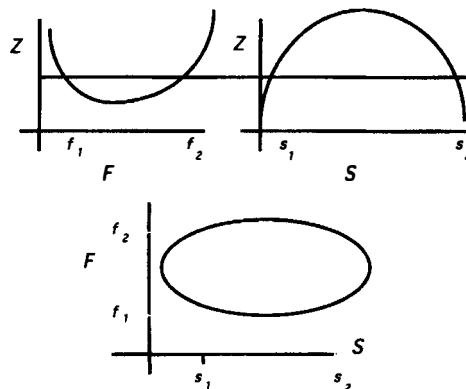


Figure 2.17. Volterra plots.

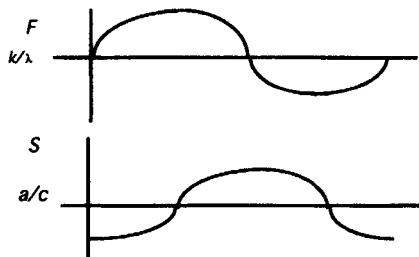


Figure 2.18. Volterra plots.

Let us now turn to a seemingly quite different problem, that of modeling the body's immune response to cancer. Calling the number of cancer cells X , let us postulate the existence of antibodies in the human organism that identify and attempt to destroy cancer cells. Let us call the number of these "immunoentities" Y , and suppose that they are given in X units; that is, one unit of Y annihilates and is annihilated by one cancer cell. Then, we can model the two populations via

$$\frac{dX}{dt} = \lambda + aX - bXY \quad (2.51)$$

and

$$\frac{dY}{dt} = cX - bXY. \quad (2.52)$$

The justification for such a model is as follows. Cancer cells are produced at a constant rate λ , which is a function of environmental factors, inability of the body to make accurate copies of some of the cells when they divide, and so on. a is the growth rate of the cancer cells. b is the rate at which antibodies attack and destroy the cancer cells. c is the rate of response of the antibody population to the presence of cancer cells.

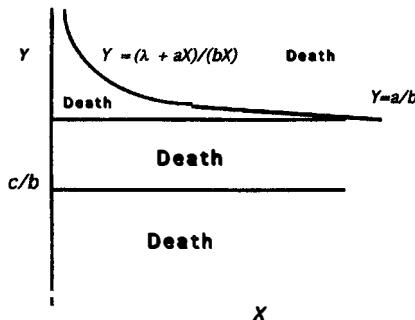


Figure 2.19. Immune system plots.

Although we cannot obtain closed-form solutions for the system given by (2.51) and (2.52), we can sketch a system of curves that will give us some feel as to which individuals will have immune systems that can cope with the oncogenesis process. From (2.52), we note that Y decreases if $dY/dt = cX - bXY < 0$, that is, if $Y > c/b$. If the inequality is reversed, then Y will increase. Similarly, from (2.51), we note that X decreases if $dX/dt = \lambda + aX - bXY < 0$, that is, if $Y > (\lambda + aX)/bX$. Let us examine the consequences of these facts by looking at Figure 2.19. The prognosis here would appear to be very bad. The body cannot fight back the cancer cells and must be overwhelmed.

However, let us examine the more hopeful scenario in Figure 2.20. We note the change if c increases dramatically relative to a . We now have regions where the body will arrive at a stable equilibrium of cancer cells and antibodies. We should also note that in both Figure 2.19 and Figure 2.20 the situation of an individual who starts out with no antibody backup at the beginning of the process is bad.

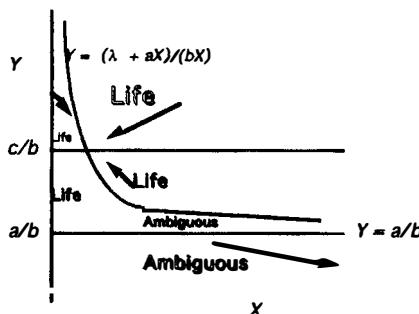


Figure 2.20. Optimistic immune plot.

We can glean other insights from the model. For example, a large enough value of λ can overwhelm any value of c . Thus, no organism can reasonably expect to have the immune response power to overcome all oncogenic shocks, no matter how big. Next, even if X is very large, provided only that we can change the biological situation to increase c dramatically, while suppressing λ , the tumor can be defeated.

The model considered here is obviously not only hugely simplified, but it is purely speculative. We have, at present, no good means of measuring X and Y . But it should be remembered that the model generally precedes the collection of data: generally, data are collected in the light of a model. In the case of Volterra's fish model, partial data were available because the selling of fish was measured for economic reasons. Volterra was, in short, fortunate that he could proceed from a well-developed data set to an explanatory model. This was serendipitous and unusual.

Generally, we waste much if we insist on dealing only with existing data sets and refuse to conjecture on the basis of what may be only anecdotal information. If we are being sufficiently bold, then for every conjecture that subsequently becomes substantiated, we should expect to be wrong a dozen times. Model building is not so much the safe and cozy codification of what we are confident about as it is a means of orderly speculation.

2.5 Pyramid Clubs for Fun and Profit

There are those who hold that the very formalism of the "free market" will produce good—irrespective of the production of any product or service other than the right to participate in the "enterprise" itself. One example of such an enterprise is gambling. Here, the player may understand that he is engaging in an activity in which his long-run expectations for success are dim—the odds are against him. Nevertheless, he will enter the enterprise for fun, excitement, and the chance that, if he only plays the game a small number of times, he will get lucky and beat the odds.

Another example of an enterprise that apparently produces no good or service is that of the pyramid club. Unlike gambling, the pyramid club gives the participants the notion that they almost certainly will "win"; that is, their gain will exceed, by a very significant margin, the cost of their participation. Let us consider a typical club structure. For the cost of \$2000, the member is allowed to recruit up to six new members. For each member he recruits, he receives a commission of \$1000. Furthermore, each of the new members is inducted with the same conditions as those of the member who inducted them. Now for each recruit made by second-level members, the first-level member receives a commission of \$100. This member is allowed to share in these \$100 commissions down through the fifth level. Generally, there is some time limit as to how long the member has to recruit his second-level members—typically a year. Thus, his anticipated return is

$$\begin{aligned}\text{anticipated return} &= \$1000 \times 6 + (6^2 + 6^3 + 6^4 + 6^5) \times \$100 \quad (2.53) \\ &= \$938,400.\end{aligned}$$

It is this apparent certainty of gain that attracts many to pyramid enterprises. Many state governments claim that this hope of gain is hugely unrealistic, and, thus, that pyramid enterprises constitute fraud. We wish to examine this claim.

Let us suppose we consider only those members of society who would become members if asked. Let us say that at any given time those who are already members will be included in the pool Y and those who have not yet joined but would if asked are included in the pool X . If we examine the probability that a member will effect a recruitment in time interval Δt , this appears to be given by

$$P(\text{recruitment in } [t, t + \Delta t]) = \frac{kX}{X + Y} \Delta t. \quad (2.54)$$

Here, k is the yearly rate of recruitment if all persons in the pool were nonmembers (e.g., $k = 6$). Then we have that the expected number of recruits by all members in $[t, t + \Delta t]$ is given by

$$E(\text{number of recruits in } [t, t + \Delta t]) = \frac{kXY}{X + Y} \Delta t. \quad (2.55)$$

We will neglect any exodus from the pool. Also, we neglect entries into the pool. Thus, if we replace the expectation of Y by Y itself, and divide by Δt , and let Δt go to 0, we have

$$\frac{dY}{dt} = \frac{kXY}{X + Y}. \quad (2.56)$$

Let us make the assumption that $X + Y = c$, a constant. Then we have the easily solvable (using partial fractions) equation

$$\frac{dY}{Y(c - Y)} = \frac{k}{c} dt. \quad (2.57)$$

So we have

$$t = \frac{1}{k} \ln \left(\frac{Y}{c - Y} \right) - \frac{1}{k} \ln \left(\frac{Y_0}{c - Y_0} \right). \quad (2.58)$$

Now, when $dY/dt = 0$, there is no further increase of Y . Thus, the equilibrium (and maximum) value of Y is given by

$$Y_e = c. \quad (2.59)$$

For the present example, the maximum value of Y , Y_e , will only be reached at $t = \infty$. But it is relevant to ask how long it will take before Y equals, say, .99c. If we assume that Y_0 equals .0001c, a little computation shows that t (when $Y = .99c$) = 2.3 years.

Now, the rate of recruitment per member per year at any given time is given by

$$\frac{dY/dt}{Y} = \frac{k(c - Y)}{c}. \quad (2.60)$$

At time $t = 2.3$, and thereafter,

$$\frac{dY/dt}{Y} \leq 0.06, \quad \frac{dY/dt}{Y} \leq 0.06. \quad (2.61)$$

Unfortunately, a member who joins at $t = 1.87$ or thereafter must replace the 6 in (2.54) by a number no greater than .06. Thus, the anticipated return to a member entering at this time is rather less than 938,400:

$$\begin{aligned} \text{anticipated return} &\leq \$1000 \times .06 + (.06^2 + .06^3 + .06^4 + .06^5) \times \$100 \\ &= 60.38. \end{aligned} \quad (2.62)$$

The difference between a pyramid structure and a bona fide franchising enterprise is clear. In franchising enterprises in which a reasonable good or service is being distributed, there is a rational expectation of gain to members even if they sell no franchises. Potential members may buy into the enterprise purely on the basis of this expectation. Still, it is clear that a different kind of saturation effect is important. The owner of a fast food restaurant may find that he has opened in an area which already has more such establishments than the pool of potential customers. But a careful marketing analysis will be enormously helpful in avoiding this kind of snafu. The primary saturation effect is not caused by a lack of potential purchasers of fast food restaurants but by an absence of customers. However, there is little doubt that many franchising operations infuse in

potential members the idea that their main profit will be realized by selling distributorships. Indeed, many such operations are *de facto* pyramid operations. Thus, it would appear to be impossible for the government to come up with a nonstifling definition of pyramid clubs which could not be circumvented by simply providing, in addition to the recruiting license, some modest good or service (numbered "collectors' item" bronze paper weights should work nicely). The old maxim of *caveat emptor* would appear to be the best protection for the public.

The model of a pyramid club is an example of epidemic structure, although no transmission of germs is involved. Nor should the term "epidemic" be considered always to have negative connotations. It simply has to do with the ability of one population to recruit, willfully or otherwise, members of another population into its ranks at a self-sustaining rate.

Problems

2.1. We shall begin with what I term a *conversational flowchart*. In reality, this sort of rather informal flowchart, mixing verbal and symbolic means of delineating a time based progression, has long since replaced the pseudo electric circuit charts of years past.

By Δt , we mean a relatively short interval of time. In a combat lasting over a month, it may well be the case that we may safely take each Δt to be a day. Now by $\Delta u(t)$ we will mean the difference between the size of the blue force u between time t and time $t + \Delta t$. Suppose, for example, that the size of the blue force at time $t = 30$ was 4050, and the size of the blue force at time $t + \Delta t$ was 4000. Then, we would have

$$\Delta u(t) = u(t + \Delta t) - u(t) = u(30) = u(31) - u(30) = -50.$$

Similarly, if we knew that $u(t) = 4050$ and that $\Delta u = -50$, then we could easily compute that

$$u(t + \Delta t) = u(t) + \Delta u(t) = 4050 - 50 = 4000.$$

Suppose next that we had some model for the change in u . One such is given in equation (2.10). Namely, we have that the degradation of the blue force is proportional to the size of the red force at that time:

$$\Delta u(t) = u(t + \Delta t) - u(t) = -c_1 v(t) \Delta t.$$

So, then, if $c_1 = .05$, and $v(0) = 21,500$, and $u(0) = 70,000$, then we would have

$$u(1) = u(0) - c_1 v(0) \Delta t = 70,000 - .05(21,500) = 71,925.$$

Of course, the red side will be degraded by the blue force. Let us suppose the model for this degradation is

$$\Delta v(t) = v(t + \Delta t) - v(t) = -c_2 u(t) \Delta t.$$

Then, taking $v(0) = 21,500$ and $c_2 = .01$, we would have

$$v(1) = v(0) - c_2 u(0) \Delta t = 21,500 - .01(73,000) = 20,770.$$

We simply carry on the process in this way for a prescribed number of time intervals (or until one of the force sizes becomes zero). Thus

$u(2) = u(1) - c_1 v(1) \Delta t = 71,925 - .05(20,770) = 70,887$,
and

$$v(2) = v(1) - c_2 u(1) \Delta t = 20,770 - .01(71,925) = 20,061.$$

If you can carry this out for 36 days, congratulations. You have achieved a simulation of the casualty figures of the Battle of Iwo Jima.

Suppose U.S. intelligence had made a mistake, and there were actually 50,000 Japanese troops on the island on day one. Compute the course of the modeled battle.

Next, to make the simulation more accurate, use the fact that the Americans stormed the beaches with 54,000 men on the first day, were reinforced by 6,000 more men on day 3 and 13,000 on day 6. (Here assume a total of 21,500 Japanese troops).

This is the way the modeling was done first time by Engel. It assumes the Japanese force went to zero by day 36 and was 8,500 on day 21. Engel had the advantage of not bothering to look and see whether there was any Japanese force data. Actually, there were 1,500 Japanese troops still fighting on day 21 and well over zero on day 36.

Maintaining the levels of attrition constants, how long would it have taken the Americans to eliminate the Japanese if they (the Americans) had had landed 54,000 troops on the first day and 10,000 daily for the next ten days? (Assume 21,500 Japanese troops on day one).

See what happens when you use the model in (2.24) with $d_1 = d_2 = .0000019$. This is the undirected fire model where the combatants cannot see each other much of the time and are simply firing at sectors where the enemy is entrenched. Now, it could be argued that the Japanese could see the Americans, but that the Japanese were concealed during much of the combat. You might try a hybrid model, where American degredation proceeds via

$$\Delta u(t) = u(t + \Delta t) - u(t) = -c_1 v(t) \Delta t$$

and the Japanese degredation proceeds according to

$$\Delta v(t) = v(t + \Delta t) - v(t) = -d_2 v(t) u(t) \Delta t.$$

You might try and model another battle (such as the Alamo). Or you might try and model another conflict, possibly one that does not involve warfare. Be creative!

2.2. Starting on September 1, 1939, a force of 188 Polish soldiers defended the Westerplatte peninsula against, roughly, 3,750 Germans for seven days. Assuming that all the Poles became casualties, as well as 1,750 Germans, compute the relative effectiveness of a Polish combatant to that of a German one. Assume directed fire against the Germans, undirected against the Poles.

2.3. Pick two opposing forces and use Lanchester theory to describe their performances under a variety of scenarios. For example, you might consider Confederates versus Yankees at both Fredericksburg and Gettysburg.

Carefully delineate your work employing something like the conversational mode flowchart used in **Problem 2.1**.

2.4. Give a data analysis for the Volterra model applied to specific predator and prey populations.

2.5. Develop a model for resistance to antibiotics.

References

- [1] Dupuy, R. E. and Dupuy, T. N. (1970). *The Encyclopedia of Military History*, New York: Harper & Row.
- [2] Engel, J. H. (1954). "A verification of Lanchester's law," *Operations Research*, **12**, 344–358.
- [3] Gabriel, R. A. (1984). *Operation Peace for Galilee*, New York: Hill and Wang.
- [4] Graunt, J. (1662). *Natural and Political Observations on the Bills of Mortality*, London.
- [5] Haberman, R. (1977). *Mathematical Models*, Englewood Cliffs, NJ: Prentice Hall.
- [6] Lanchester, F. W. (1916). "Mathematics in warfare," reprinted in *The World of Mathematics*, Vol. 4, J. R. Newman, ed., New York: Simon and Schuster, 1956, pp. 2138–2157.
- [7] Newcomb, R. F. (1965). *Iwo Jima*, New York: Holt, Rinehart and Winston.
- [8] Thompson, J. R. (1979). "An example of data-poor model validation," in *Decision Information*, C. P. Tsokos and R. M. Thrall, eds., New York: Academic, pp. 405–408.
- [9] Thompson, J. R. (1979). "An argument for fortified defense in western Europe," in *Decision Information*, C. P. Tsokos and R. M. Thrall, eds., New York: Academic, pp. 395–404.

Chapter 3

Epidemics

3.1 Introduction

The author has enjoyed more than three decades of collaboration with members of Houston's Texas Medical Center. Oncological modeling has been a particular interest. But I have spent a great deal of time dealing with contagious diseases in general and AIDS in particular. It seems strange to me that most of my colleagues use the word "epidemic" in conjunction with noncontagious diseases such as cancer and multiple sclerosis, but that is the case. Indeed, I know only a very few people other than myself who work with contagious diseases and restrict the use of "epidemic" to describe diseases of contagion. I gather that the reasons for this kind of twisting of the language have mainly to do with the sociology and political matrix in which disease investigators find themselves. Although in other chapters of this book, we shall deal with such noncontagious diseases as cancer, in this chapter we shall be concerned with epidemics in the classical contagion sense.

Among the contagious diseases of history, leprosy has an important position because it is so often mentioned in the *Bible*. The author has observed that the number of verses in the *Old Testament* concerning leprosy is roughly three times the number of verses dealing with the famous kosher regulations. This disease was typically controlled in ancient cultures by isolating the infectives from the rest of the population. A harsh system to be sure, but one which was more effective than some protocols that have been introduced into Third World Countries by First World medications that have proved less effective than promised.

It is interesting to note that leprosy has a long history in East Texas. The time line does not admit of the possibility that it was vectored from, say, Mexico. The armadillo is one possible candidate for the spread of the

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

disease. As the reduction of the grinding poverty following the War Between the States has advanced, more Texans have given up undercooked armadillo meat for alternatives. However, the threat from this disease should not be underrated.

Please have a look at the short discussion of the shutting down of the (<http://www.abcnews.go.com/sections/us/DailyNews/hansens990329a.html>) Carville, Louisiana, facility. This article is somewhat typical of the "rosy scenario" view of epidemics. There are some rather bizarre claims in the article. For example:

Hansen's Disease is among the least contagious of infectious diseases. More than 95 percent of the population is naturally immune to it. And its easily treated these days.

"It is only contagious in certain stages, and once medicated, patients show no risk to the public," says Dr. Bruce Clements, Carville's clinical director of patient care. "Most people who worked with Hansen's Disease sufferers never contracted the disease."

Note that Dr. Clements does not state that it is extremely rare for people who work with Hansen's Disease sufferers to contract the disease. If 95% were immune to the disease, then, considering that health care workers at Carville have historically taken careful precautions to protect themselves, "most people" would have been replaced by "almost all people." Carville has traditionally paid high wages to workers and provided rather posh living conditions for them. This has not been done without reason. The notion that 95% of the human race is immune to leprosy is probably ridiculous. Considering the high mortality rates from many diseases, an epidemic to which only 5% of the population was susceptible would hardly have attracted such stringent policies for epidemiological control. Moreover, we know that the spouses of lepers who went to Molokai frequently accompanied them. Anecdotal evidence indicates that they generally contracted the disease themselves, as did the first Molokai mission priest Father Damien.

Working in the late 1980s with Professor of Pathology Raymond McBride (McBride ran the major pathology lab of Houston) of the Baylor College of Medicine, I learned that leprosy had been endemic among East Texas farmers as far back as records were kept. This was not a Third World importation, it was endemic in the almost completely Anglo East Texas population. Certainly the relative frequency of the disease has diminished. Why is that? We really do not know (and probably we should try and find out). Conjectures abound. For example, some have opined that the armadillo might be a vector for the disease. Apparently, modern East Texans prize armadillo stew rather less than did their forbears. They also live in air-conditioned houses with indoor toilets and running water.

The \$18 million which it cost the U.S. taxpayers annually to keep the Carville Center open was money well spent. The notion that the 200 new U.S. cases of leprosy per year in the United States can be treated in local

healthcare centers without danger of spread is an example of the triumph of hope over experience.

Let us now move to another disease of antiquity (actually virtually all contagious diseases have been around for thousands of years; there are no new species of animals and no new diseases), namely smallpox. We have good historical evidence that it was this disease that destroyed the inhabitants of the Athens–Piraeus fortifications of Pericles.

Just as we can destroy species of animals, it should be theoretically possible to destroy bacteria and viruses which cause particular diseases. Such is the confidence of the U.S. Centers for Disease Control that smallpox (outside laboratories) has been eradicated that inoculation against the disease in the United States has been long since halted. We really have trivial supplies of the vaccine in the event that someone decided to weaponize some of the virus from laboratories. The risk from the vaccine is infinitesimal.

However, vaccination for whooping cough, a disease much less deadly than smallpox is essentially a requirement for an American child getting admitted into the public school system. And the incidence of autism has skyrocketed in the United States since the whooping cough vaccine became a requirement to get into school. Correlations, although not a proof, are an important part of exploratory data analysis.

3.2 John Snow and the East End of London Cholera Epidemic of 1854

Let us start with an example from 19th century epidemiology. Please read the material at <http://www.cdc.gov/excite/snow.pdf>. You should also read the complete text at http://www.ph.ucla.edu/epi/snow/snowbook_a1.html of 19-century pioneering epidemiologist John Snow (to get the full four installments, you change a1 in the above to a2, a3, a4). We are concerned here with an outbreak of cholera in London. What Snow did was to look at the frequency of cholera cases plotted on a map and use this information to seek for an (<http://www.ph.ucla.edu/epi/snow/map1ea.htm>) apparent cause of the epidemic somewhere in the center of the high-frequency area. We note that Dr. Snow performed his data analysis without the benefit of mathematical formulas. His was a good example of “exploratory data analysis.” EDA, as often as not, relies on graphical procedures. Transmission by drinking water was one of the candidates thought of as possible by Snow. A simple plot showing a hash mark for a death pointed Snow to a water source almost in the center of high death concentration. Truly, it could have been argued (and was) that the pump was simply coincidentally located in the center of the death circle. Snow did a great deal of canvassing work. He noted, for example, that the workers in a brewery close to the pump did not exhibit cholera deaths. But the owner of the brewery revealed that it had its own deep well which was used for the making of their beer (which

the workers consumed in lieu of water). Snow further found instances of persons remote from the pump who nevertheless obtained their water from it (including a lady from the upscale West End of London, who daily had a supply carted to her) and subsequently died of cholera. Snow made a case for removing the pump handle that was based on his plot plus extensive interviewing.

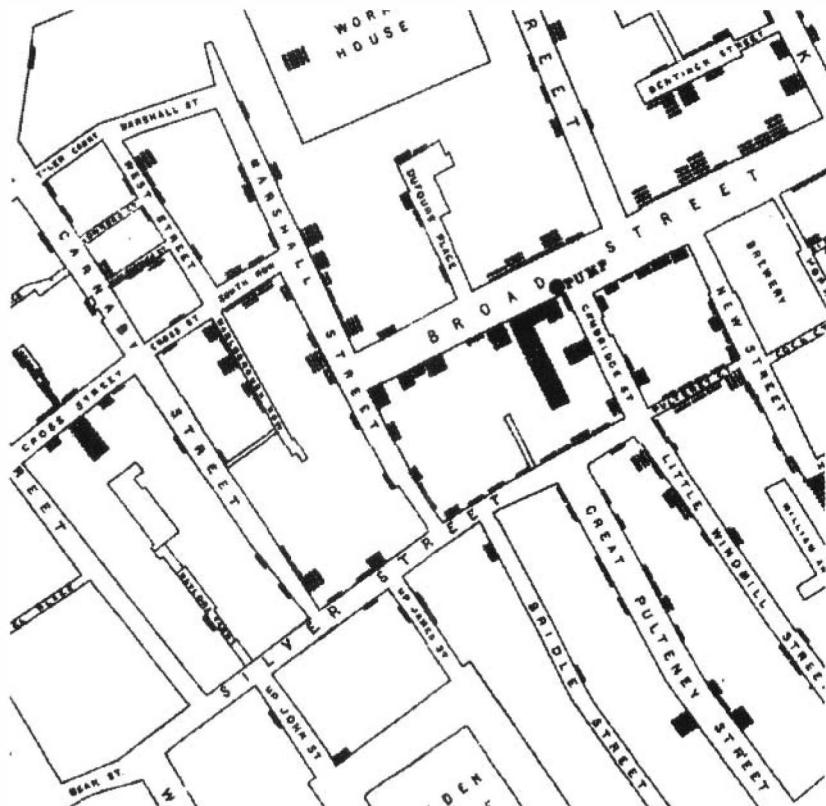


Figure 3.1. Vicinity of the “Broad Street Pump”.

The removal of the pump handle from the now famous “Broad Street Pump” was a matter of inconvenience for many. They had to gather their water from farther away. Here is a classic example of cost/benefit analysis. Snow evidently won his argument with the local parish council arguing that the modest inconvenience was well worth while in light of the fact of a raging epidemic.

Apparently, a shallow well in the area, served by one pump, had been contaminated by cholera feces transported to a cesspool that drained into the well. Removal of the pump handle by Snow caused some inconvenience, but it seems to have saved hundreds of lives. This is a classic situation in the treatment of epidemics: isolation of the susceptible population from active sources of the infection. Notice how Snow by a very simple sociological

intervention (removing the handle of the Broad Street Pump) ended a dangerous epidemic without vaccines or antibiotics. Later on in this chapter, it is argued that a simple sociological intervention could be used against the United States AIDS epidemic, again without the use of the vaccines (unavailable for AIDS) or antiviral agents (unavailable for AIDS).

3.3 Prelude: The Postwar Polio Epidemic

Effective immunizations against many of the killing diseases of the 19th century, plus antibiotics massively used during World War II, gave the promise of the end to life-threatening contagion in the United States. The killers of the future would be those largely associated with the aging process, such as cancer, stroke, and heart attack.

However, in the postwar years, polio, which already had stricken some (including President Roosevelt), became a highly visible scourge in a number of American cities, particularly in the South, particularly among the young. In 1952, over 55,000 cases were reported. Mortality rates in America, due to good care, had by that time dropped to well under 10%. Nonetheless, the spectacle of children confined to wheelchairs or iron lungs was a disturbing one.

This was in the years before the emigration of the middle classes to suburbia, and most schools tended to have representation from a wide range of socioeconomic groups. Incidence rates were the highest in the summers, when the schools were closed. But, at the intuitive level, it was clear that polio was a disease predominantly of school age children, and that there was a fair amount of clustering of cases. Although the causative agent had not been isolated, there was little doubt that it was a virus, that it favored young hosts, that the throat was the likely pathway, and that transmission was greatest in the hot weather.

In such a situation, it might appear that a prudent public health policy would be to discourage summer gatherings of children, particularly in confined indoor settings or in swimming pools. Such an inference might well be put down as a prejudice of causation where none existed. Indeed, this was the era of the kiddie matinee and new municipal swimming facilities given by city governments to their citizens in celebration of a perceived affluence following the War. Some parents did, to the displeasure of their children, attempt to deprive them of matinees and swimming excursions, but such were in the distinct minority. From time to time, city officials would take such steps as shutting down municipal swimming pools, but this was unusual and always temporary. There was a large economic constituency for matinees and swimming pools. The constituency for shutting them down was acting on intuition and without business support. The results were that the movies and pools generally stayed open all summer. The epidemic flourished.

There was a great deal of expectation that "the cavalry will soon ride to the rescue" in the form of an expected vaccine against the disease. In 1955, the Salk vaccine¹ did appear, and new polio cases, for the United States, became a thing of the past. Of course, a residual population of tens of thousands of Americans remained, crippled by polio.

There was very little in the way of a postmortem examination about how effective public health policy had been in managing the American polio epidemic. In fact, there had been essentially no proactive policy at all. But two effective anti-polio vaccines (Salk and then Sabin) seemed to have brought everything right in the end. If there were serious efforts to learn from the mistakes in management of the American polio epidemic, this author has not seen them.

Polio had, apparently, been simply a bump in the road toward a time in which life-threatening contagious diseases in America would be a thing of the past. However, having spent my childhood in Memphis, Tennessee (one of the epicenters of the postwar polio epidemic), that epidemic was something I would never forget. My parents were among the number of those who forbade matinees and swimming pools to their children. But among my childhood friends there were several who died from polio, and many others crippled by it.

3.4 AIDS: A New Epidemic for America

In 1983, I was investigating the common practice of using stochastic models in dealing with various aspects of diseases. When attempting to model the progression of cancer within an individual, a good case could be made for going stochastic. For example, one matter of concern with solid tumors is whether the primary tumor throws off a metastasis before it has been removed surgically. Whether it has or has not will largely determine whether surgical removal of the primary tumor has cured the patient. Such a phenomenon needs to be modeled stochastically.

However, when modeling the progression of a contagious disease through a population, the common current practice of using a stochastic model and then finding, for example, the moment generating function of the number $Y(t)$ of infectives seems unnecessarily complicated, particularly if, at the end of the day, one decides simply to extract $E[Y(t)]$, the expected number of infectives. Moreover, any sociological data, if available, are likely to be in terms of aggregate information, such as the average number of contacts per day.

I had decided to write a paper giving examples where deterministic modeling would probably be appropriate. I selected the AIDS epidemic because

¹In 1999, evidence started to appear that contamination of the Salk vaccine by a monkey virus, not unrelated to HIV, was causing many recipients of the Salk vaccine to develop a variety of cancers, possibly due to a destruction of parts of their immune system.

it was current news, with a few hundred cases reported nationally. Although reporting at the time tended to downplay the seriousness of the epidemic (and, of course, the name was pointedly innocuous, the same as an appetite suppressant of the times), there was a palpable undercurrent of horror in the medical community. It looked like a study that might be important.

Even at the very early stage of an observed United States AIDS epidemic, several matters appeared clear to me:

- The disease favored the homosexual male community and outbreaks seemed most noticeable in areas with sociologically identifiable gay communities.
- The disease was also killing (generally rather quickly) people with acute hemophilia.
- Given the virologist's maxim that there are no new diseases, AIDS, in the United States, had been identified starting around 1980 because of some sociological change. A disease endemic under earlier norms, it had blossomed into an epidemic due to a change in society.

At the time, which was before the HIV virus had been isolated and identified, there was a great deal of commentary both in the popular press and in the medical literature (including that of the Centers for Disease Control) to the effect that AIDS was a new disease. Those statements were not only putatively false, but were also potentially harmful. First of all, from a practical virological standpoint, a new disease might have as a practical implication genetic engineering by a hostile foreign power. This was a time of high tension in the Cold War, and such an allegation had the potential for causing serious ramifications at the level of national defense.

Secondly, treating an unknown disease as a new disease essentially removes the possibility of stopping the epidemic sociologically by simply seeking out and removing (or lessening) the cause(s) that resulted in the endemic being driven over the epidemiological threshold.

For example, if somehow a disease (say, the Lunar Pox) has been introduced from the moon via the return of moon rocks by American astronauts, that is an entirely different matter than, say, a mysterious outbreak of dysentery in St. Louis. For dysentery in St. Louis, we check food and water supplies, and quickly look for "the usual suspects" — unrefrigerated meat, leakage of toxins into the water supply, and so on. Given proper resources, eliminating the epidemic should be straightforward.

For the Lunar Pox, there are no usual suspects. We cannot, by reverting to some sociological *status quo ante*, solve our problem. We can only look for a bacterium or virus and try for a cure or vaccine. The age-old way of eliminating an epidemic by sociological means is difficult — perhaps impossible.

In 1982, it was already clear that the United States public health establishment was essentially treating AIDS as though it were the Lunar Pox.

The epidemic was at levels hardly worthy of the name in Western Europe, but it was growing. Each of the European countries was following classical sociological protocols for dealing with a venereal disease. These all involved some measure of defacilitating contacts between infectives and susceptibles. The French demanded bright lighting in gay "make-out" areas. Periodic arrests of transvestite prostitutes on the Bois de Bologne were widely publicized. The Swedes took much more draconian steps, mild in comparison with those of the Cubans. The Americans took no significant sociological steps at all.

However, as though following the Lunar Pox strategy, the Americans outdid the rest of the world in money thrown at research related to AIDS. Some of this was spent on isolating the unknown virus. However, it was the French, spending pennies to the Americans' dollars, at the Pasteur Institute (financed largely by a legacy from the late Duke and Duchess of Windsor) who first isolated HIV. In the intervening 30 years since isolation of the virus, no effective vaccine or cure has been produced.

3.5 Why an AIDS Epidemic in America?

Although the popular press in the early 1980s talked of AIDS as being a new disease, as noted, prudence and experience indicated that it was not. Just as new species of animals have not been noted during human history, the odds for a sudden appearance (absent genetic engineering) of a new virus are not good. My own discussions with pathologists with some years of experience gave anecdotal cases of young Anglo males who had presented with Kaposi's sarcoma at times going back to early days in the pathologists' careers. This pathology, previously seldom seen in persons of Northern European extraction, now widely associated with AIDS, was at the time simply noted as isolated and unexplained. Indeed, a few years after the discovery of the HIV virus, HIV was discovered in decades old refrigerated human blood samples from both Africa and America.

Although it was clear that AIDS was not a new disease, as an epidemic it had never been recorded. Because some early cases were from the Congo, there was an assumption by many that the disease might have its origins there. Clearly, record keeping in the Congo was not and is not very good. But Belgian colonial troops had been located in that region for many years. Any venereal disease acquired in the Congo should have been vectored into Europe in the 19th century. But no AIDS-like disease had been noted. It would appear, then, that AIDS was not contracted easily as is the case, say, with syphilis. Somehow, the appearance of AIDS as an epidemic in the 1980s, and not previously, might be connected with higher rates of promiscuous sexual activity made possible by the relative affluence of the times.

Then there was the matter of the selective appearance of AIDS in the

American homosexual community. If the disease required virus in some quantity for effective transmission (the swift progression of the disease in hemophiliacs plus the lack of notice of AIDS in earlier times gave clues that such might be the case), then the profiles in Figures 3.2 and 3.3 give some idea why the epidemic seemed to be centered in the American homosexual community. If passive to active transmission is much less likely than active to passive, then clearly the homosexual transmission patterns facilitate the disease more than the heterosexual ones.

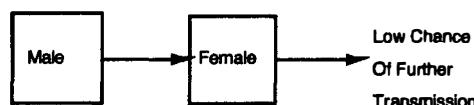


Figure 3.2. Heterosexual transmission of AIDS.



Figure 3.3. Homosexual transmission of AIDS.

One important consideration that seemed to have escaped attention was the appearance of the epidemic in 1980 instead of 10 years earlier. Gay lifestyles had begun to be tolerated by law enforcement authorities in the major urban centers of America by the late 1960s. If homosexuality was the facilitating behavior of the epidemic, then why no epidemic before 1980? Of course, believers in the "new disease" theory could simply claim that the causative agent was not present until around 1980. In the popular history of the early American AIDS epidemic, *And the Band Played On*, Randy Shilts points at a gay flight attendant from Quebec as a candidate for "patient zero." But this "Lunar Pox" theory was not a position that any responsible epidemiologist could take (and, indeed, as pointed out, later investigations revealed HIV samples in human blood going back into the 1940s).

What accounts for the significant time differential between civil tolerance of homosexual behavior prior to 1970 and the appearance of the AIDS epidemic in the 1980s? Were there some other sociological changes that had taken place in the late 1970s that might have driven the endemic over the epidemiological threshold?

It should be noted that in 1983, data were skimpy and incomplete. As is frequently the case with epidemics, decisions need to be made at the early stages when one needs to work on the basis of skimpy data, analogy with

other historical epidemics, and a model constructed on the best information available.

I remember in 1983 thinking back to the earlier American polio epidemic that had produced little in the way of sociological intervention and less in the way of models to explain the progress of the disease. Although polio epidemics had been noted for some years (the first noticed epidemic occurred around the time of World War I in Stockholm), the American public health service had indeed treated it like the “Lunar Pox.” That is, they discarded sociological intervention based on past experience of transmission pathways and relied on the appearance of vaccines at any moment. They had been somewhat lucky, since Salk started testing his vaccine in 1952 (certainly they were luckier than the thousands who had died and the tens of thousands who had been permanently crippled). But basing policy on hope and virological research was a dangerous policy (how dangerous we are still learning as we face the reality of 650,000 Americans dead by 2011 from AIDS).

Although some evangelical clergymen inveighed against the epidemic as divine retribution on homosexuals, the function of epidemiologists is to use their God-given wits to stop epidemics. In 1983, virtually nothing was being done except to wait for virological miracles.

3.5.1 Political Correctness Can Kill

One possible candidate was the turning of a blind eye by authorities to the gay bathhouses that started in the late 1970s. These were places where gays could engage in high frequency anonymous sexual contact. By the late 1970s they were allowed to operate without regulation in the major metropolitan centers of America. My initial intuition was that the key was the total average contact rate among the target population. Was the marginal increase in the contact rate facilitated by the bathhouses sufficient to drive the endemic across the epidemiological threshold? It did not seem likely. Reports were that most gays seldom (many, never) frequented the bathhouses.

In the matter of the present AIDS epidemic in the United States, a great deal of money is being spent on AIDS. However, practically nothing in the way of steps for stopping the transmission of the disease is being done (beyond education in the use of condoms). Indeed, powerful voices in the Congress speak against any sort of government intervention. On April 13, 1982, Congressman Henry Waxman [2] stated in a meeting of his Subcommittee on Health and the Environment, “I intend to fight any effort by anyone at any level to make public health policy regarding Kaposi’s sarcoma or any other disease on the basis of his or her personal prejudices regarding other people’s sexual preferences or life styles.” (It is significant to note that Representative Waxman has been one of the most strident voices in the fight to stop smoking and global warming, considering rigorous measures

acceptable to end these threats to human health.) We do not even have a very good idea as to what fraction of the target population in the United States is HIV positive, and anything approaching mandatory testing is regarded by American political leaders as an unacceptable infringement of civil liberties. In light of Congressman Waxman's warnings, it would have taken brave public health officials to close the gay bathhouses. The Centers for Disease Control have broad discretionary powers and its members have military uniforms to indicate their authority. They have no tenure, however. The Director of the CDC could have closed the bathhouses, but that would have been an act of courage which could have ended his career. It appears odd to say so, but of all the players in the United States AIDS epidemic, Congressman Waxman may be more responsible than any other for what has turned out to be a death tally exceeding any of America's wars, including its most lethal, the American War Between the States (aka the Civil War).

3.6 The Effect of the Gay Bathhouses

But perhaps my intuitions were wrong. Perhaps it was not only the total average contact rate that was important, but a skewing of contact rates, with the presence of a high activity subpopulation (the bathhouse customers) somehow driving the epidemic. It was worth a modeling try.

The model developed in [3] considered the situation in which there are two subpopulations: the majority, less sexually active, and a minority with greater activity than that of the majority. We use the subscript “1” to denote the majority portion of the target (gay) population, and the subscript “2” to denote the minority portion. The latter subpopulation, constituting fraction p of the target population, will be taken to have a contact rate τ times the rate k of the majority subpopulation. The following differential equations model the growth of the number of susceptibles X_i and infectives Y_i in subpopulation i ($i = 1, 2$).

$$\begin{aligned}\frac{dY_1}{dt} &= \frac{k\alpha X_1(Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)} - (\gamma + \mu)Y_1, \\ \frac{dY_2}{dt} &= \frac{k\alpha\tau X_2(Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)} - (\gamma + \mu)Y_2, \\ \frac{dX_1}{dt} &= -\frac{k\alpha X_1(Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)} + (1-p)\lambda - \mu X_1, \\ \frac{dX_2}{dt} &= -\frac{k\alpha\tau X_2(Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)} + p\lambda - \mu X_2.\end{aligned}\tag{3.1}$$

where

k = number of contacts per month,

α = probability of contact causing AIDS,

λ = immigration rate into the population,
 μ = emigration rate from the population,
 γ = marginal emigration rate from the population due
 to sickness and death.

In Thompson [3], it was noted that if we started with 1,000 infectives in a target population with $k\alpha = 0.05$, $\tau = 1$, a susceptible population of 3,000,000 and the best guesses then available ($\mu = 1/(15 \times 12) = 0.00556$, $\gamma = 0.1$, $\lambda = 16,666$) for the other parameters, the disease advanced as shown in Table 3.1.

Table 3.1. Extrapolated AIDS cases: $k\alpha = 0.05$, $\tau = 1$.

Year	Cumulative deaths	Fraction infective	
1	1751	0.00034	
2	2650	0.00018	
3	3112	0.00009	
4	3349	0.00005	
5	3571	0.00002	
10	3594	0.000001	

Next, a situation was considered in which the overall contact rate was the same as in Table 3.1, but it was skewed with the more sexually active subpopulation 2 (of size 10%) having contact rates 16 times those of the less active population.

Table 3.2. Extrapolated AIDS cases: $k\alpha = 0.02$, $\tau = 16, p = 0.10$.

Year	Cumulative deaths	Fraction infective
1	2,184	0.0007
2	6,536	0.0020
3	20,583	0.0067
4	64,157	0.0197
5	170,030	0.0421
10	855,839	0.0229
15	1,056,571	0.0122
20	1,269,362	0.0182

Even though the overall average contact rate in Table 3.1 and Table 3.2 is the same ($k\alpha$)_{overall} = 0.05, the situation is dramatically different in the two

cases. Here, it seemed, was a *prima facie* explanation as to how AIDS was pushed over the threshold to a full-blown epidemic in the United States: a small but sexually very active subpopulation.

I note that nothing more sophisticated than some numerical quadrature was required to obtain the results in these tables. In the ensuing arguments concerning why AIDS became an epidemic in the United States, everything beyond the simple deterministic model (3.1) will be, essentially, frosting on the cake. This was the way things stood in 1984 when I presented the paper at the summer meetings of the Society for Computer Simulation in Vancouver. It hardly created a stir among the mainly pharmacokinetic audience who attended the talk. And, frankly, at the time I did not think too much about it because I supposed that probably even as the paper was being written, the “powers that be” were shutting down the bathhouses. The deaths at the time were numbered in the hundreds, and I did not suppose that things would be allowed to proceed much longer without sociological intervention. Unfortunately, I was mistaken.

In November 1986, the First International Conference on Population Dynamics took place at the University of Mississippi where there were some of the best biomathematical modelers from Europe and the United States. I presented my AIDS results [6], somewhat updated, at a plenary session. By this time, I was already alarmed by the progress of the disease (over 40,000 cases diagnosed and the bathhouses still open). The bottom line of the talk had become more shrill: namely, every month delayed in shutting down the bathhouses in the United States would result in thousands of deaths. The reaction of the audience this time was concern, partly because the prognosis seemed rather chilling, partly because the argument was simple to follow and seemed to lack holes, and partly because it was clear that something was pretty much the matter if things had gone so far off track.

After the talk, the well-known Polish probabilist Robert Bartoszyński, with whom I had carried out a lengthy modeling investigation of breast cancer and melanoma (at the Curie-Sklodowska Institute in Poland and at Rice), took me aside and asked whether I did not feel unsafe making such claims. “Who,” I asked, “will these claims make unhappy”? “The homosexuals,” said Bartoszyński. “No, Robert,” I said, “I am trying to save their lives. It will be the public health establishment who will be offended.”

And so it has been in the intervening years. I have given AIDS talks before audiences with significant gay attendance in San Francisco, Houston, and other locales without any gay person expressing offense. Indeed, in his 1997 book [1], Gabriel Rotello, one of the leaders of the American gay community, not only acknowledges the validity of my model but also constructs a survival plan for gay society in which the bathhouses have no place.

3.7 A More Detailed Look at the Model

A threshold investigation of the two-activity population model (3.1) is appropriate here. Even today, let alone in the mid-1980s, there was no chance that one would have reliable estimates for all the parameters k , α , γ , μ , λ , p , τ . Happily, one of the techniques sometimes available to the modeler is the opportunity to express the problem in such a form that most of the parameters will cancel. For the present case, we will attempt to determine the $k\alpha$ value necessary to sustain the epidemic when the number of infectives is very small. For this epidemic in its early stages one can manage to get a picture of the bathhouse effect using only a few parameters: namely, the proportion p of the target population which is sexually very active and the activity multiplier τ .

For $Y_1 = Y_2 = 0$ the equilibrium values for X_1 and X_2 are $(1-p)(\lambda/\mu)$ and $p(\lambda/\mu)$, respectively. Expanding the right-hand sides of (3.1) in a Maclaurin series, we have (using lower case symbols for the perturbations from 0)

$$\begin{aligned}\frac{dy_1}{dt} &= \left[\frac{k\alpha(1-p)}{1-p+\tau p} - (\gamma + \mu) \right] y_1 + \frac{k\alpha(1-p)\tau}{1-p+\tau p} y_2 \\ \frac{dy_2}{dt} &= \frac{k\alpha\tau p}{1-p+\tau p} y_1 + \left[\frac{k\alpha\tau^2 p}{1-p+\tau p} - (\gamma + \mu) \right] y_2.\end{aligned}$$

Summing then gives

$$\frac{dy_1}{dt} + \frac{dy_2}{dt} = [k\alpha - (\gamma + \mu)] y_1 + [k\alpha\tau - (\gamma + \mu)] y_2.$$

In the early stages of the epidemic,

$$\frac{dy_1/dt}{dy_2/dt} = \frac{(1-p)}{p\tau}.$$

That is to say, the new infectives will be generated proportionately to their relative numerosity in the initial susceptible pool times their relative activity levels. So, assuming a negligible number of initial infectives, we have

$$y_1 = \frac{(1-p)}{p\tau} y_2.$$

Substituting in the expression for $dy_1/dt + dy_2/dt$, we see that for the epidemic to be sustained, we must have

$$k\alpha > \frac{(1+\mu)(1-p+\tau p)}{1-p+p\tau^2} (\gamma + \mu). \quad (3.2)$$

Accordingly we define the *heterogeneous threshold* via

$$k_{\text{het}}\alpha = \frac{(1+\mu)(1-p+\tau p)}{1-p+p\tau^2} (\gamma + \mu).$$

Now, in the homogeneous contact case (i.e., $\tau = 1$), we note that for the epidemic not to be sustained, the condition in equation (3.3) must hold.

$$k\alpha < (\gamma + \mu). \quad (3.3)$$

Accordingly we define the *homogeneous threshold* by

$$k_{\text{hom}}\alpha = (\gamma + \mu).$$

For the heterogeneous contact case with k_{het} , the average contact rate is given by

$$k_{\text{ave}}\alpha = p\tau(k_{\text{het}}\alpha) + (1 - p)(k_{\text{het}}\alpha) = \frac{(1 + \mu)(1 - p + \tau p)}{1 - p + p\tau^2}(\gamma + \mu).$$

Dividing the sustaining value $k_{\text{hom}}\alpha$ by the sustaining value $k_{\text{ave}}\alpha$ for the heterogeneous contact case then produces

$$Q = \frac{1 - p + \tau^2 p}{(1 - p + \tau p)^2}.$$

Notice that we have been able here to reduce the parameters necessary for consideration from seven to two. This is fairly typical for model-based approaches: the dimensionality of the parameter space may be reducible in answering specific questions. Figure 3.4 shows a plot of this “enhancement factor” Q as a function of τ . Note that the addition of heterogeneity to the transmission picture has roughly the same effect as if all members of the target population had more than doubled their contact rate. Remember that the picture has been corrected to discount any increase in the overall contact rate which occurred as a result of adding heterogeneity. In other words, the enhancement factor is totally a result of heterogeneity. It is this heterogeneity effect which I have maintained (since 1984) to be the cause of AIDS getting over the threshold of sustainability in the United States.

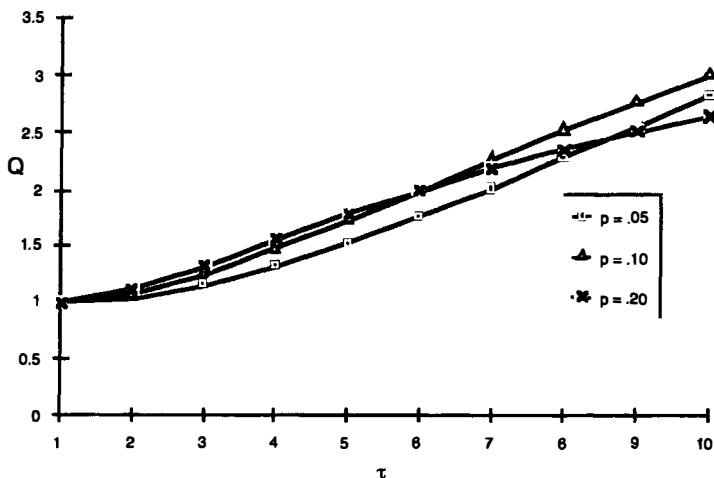


Figure 3.4 Effect of a high activity subpopulation.

If this all still seems counterintuitive, then let us consider the following argument at the level of one individual infective. Suppose, first of all, that the disease is such that one contact changes a susceptible to an infective. Then let us suppose we have an infective who is going to engage in five contacts. What number of susceptibles (assuming equal mixing) will give the highest expected number of conversions of susceptibles to infectives? Note that if the number of susceptibles is small, the expectation will be lessened by the “overkill effect”: i.e., there is the danger that some of the contacts will be “wasted” by being applied to an individual already infected by one of the other five contacts. Clearly, here the optimal value for the size N of the susceptible pool is infinity, for then the expected number of conversions from susceptible to infective $E(\mathcal{I} | N = \infty)$ is five.

Now let us change the situation to one in which two contacts, rather than one, are required to change a susceptible to an infective. We will still assume a total of five contacts. Clearly, if $N = 1$ then the expected number of conversions is $E = 1$; there has been wastage due to overkill. Next, let us assume the number of susceptibles has grown to $N = 2$. Then the probability of two new infectives is given by

$$P(2 | N = 2) = \sum_{j=2}^3 \binom{5}{j} \left(\frac{1}{2}\right)^5 = \frac{20}{32}.$$

The probability of only one new infective is $1 - P(2 | N = 2)$. Thus the expected number of new infectives is

$$E(\mathcal{I} | N = 2) = 2 \left(\frac{20}{32}\right) + 1 \left(\frac{12}{32}\right) = 1.625.$$

Now when there are $N = 3$ susceptibles, the contact configurations leading to two new infectives are of the type $(2, 2, 1)$ and $(3, 2, 0)$. All other configurations will produce only one new infective. So the probability of two new infectives is given by

$$P(2 | N = 3) = \binom{3}{1} \frac{5!}{2!2!1!} \left(\frac{1}{3}\right)^5 + \binom{3}{1} \binom{2}{1} \frac{5!}{3!2!} \left(\frac{1}{3}\right)^5 = \frac{150}{243},$$

and the expected number of new infectives is

$$E(\mathcal{I} | N = 3) = 2 \left(\frac{150}{243}\right) + 1 \left(\frac{93}{243}\right) = 1.617.$$

Additional calculations give $E(\mathcal{I} | N = 4) = 1.469$ and $E(\mathcal{I} | N = 5) = 1.314$. For very large N , $E(\mathcal{I})$ is of order $1/N$. Apparently, for the situation where there are a total of five contacts, the value of the number in the susceptible pool that maximizes the total number of new infectives from the one original infective is $N = 2$, not ∞ . Obviously, we are oversimplifying, since we stop after only the contacts of the original infective. The situation is much more complicated here, since an epidemic is created by the new infectives infecting others and so on. As well, there is the matter of a distribution of the number of contacts required to give the disease. We have in our main model (3.1) avoided the complexities of branching process modeling by going deterministic. The argument above is given to present an intuitive feel as to the facilitating potential of a high contact core in driving a disease over the threshold of sustainability.

In the case of AIDS, the average number of contacts required to break down the immune system sufficiently to cause the person ultimately to get AIDS is much larger than two. The obvious implication is that a great facilitator for the epidemic being sustained is the presence of a subpopulation of susceptibles whose members have many contacts. In the simple example above, we note that even if the total number of contacts were precisely five, from a standpoint of facilitating the epidemic, it would be best to concentrate the contacts into a small pool of susceptibles. In other words, if the total number of contacts is fixed at some level, it is best to start the epidemic by concentrating the contacts within a small subpopulation. Perhaps the analogy to starting a fire, not by dropping a match onto a pile of logs, but rather onto some kindling beneath the logs, is helpful.

3.8 Forays into the Public Policy Arena

The senior Professor of Pathology at the Baylor College of Medicine in the 1980s was Raymond McBride. McBride had been one of the pioneers in immunosuppression for organ transplantation and was the Chief of Pathology Services for the Harris County (Houston) Medical District. Distressed

to see the ravages of AIDS on autopsied victims, he was quite keen to have municipal authorities act to close down the bathhouses. He and I co-authored a front page op-ed piece for the *Houston Chronicle* 4-9-1989 titled "Close Houston's Gay Bathhouses", taking care not to mention the names and addresses of the two major offending establishments lest some vigilante act be taken against them. Hardly a ripple of interest, even though Houston, with less than one-tenth the population of Canada, had more AIDS cases than that entire country. We tried to motivate members of the City Council. When interviewed by a reporter, the office of the Councilman in whose district these two bathhouses were situated shrugged the whole matter off by asking, "What's a bathhouse"? I served on the American Statistical Association's Ad Hoc Committee on AIDS from its inception until its demise. But our mandate was never allowed to extend to modeling. Only the methodology of data analysis was permitted. Nor were we allowed, as a committee, to compare America's AIDS incidence with that from other countries.

The situation was not unlike that of the earlier polio epidemic. There were specific interests for not addressing the bathhouse issue, but there was only a nonspecific general interest for addressing it.

Although I myself had no experience with the blood-testing issue, it should be noted that early on in the epidemic, long before the discovery of HIV, it was known that over 90% of the persons with AIDS tested positive to antibodies against Hepatitis-B. For many months, the major blood collecting agencies in the United States resisted employing the surrogate Hepatitis test for contaminated blood. The result was rampant death among hemophiliacs and significant AIDS infections among persons requiring large amounts of blood products for surgery.

The statistician/economist/sociologist Vilfredo Pareto remarked that Aristotle had made one mistake when he presented to the world the system of logical thinking. The mistake was Aristotle's assumption that once humankind understood logical consistency, actions, including public policy, would be made on the basis of reason. Pareto noted that the historical record showed otherwise. The more important the decision, Pareto noted, the less likely was logical inference based on facts. This a significant concern in decision making. So, it has unfortunately been with policy concerning AIDS.

3.9 Modeling the Mature Epidemic

In the United States, the AIDS epidemic crossed the threshold of viability long ago. Consequently, we should investigate the dynamics of the mature epidemic. Unfortunately, we then lose the ability to disregard five of the seven parameters and must content ourselves with picking reasonable values for those parameters. A detailed analysis is given in Thompson and Go [7].

In the following, we will make certain ballpark assumptions about some of the underlying parameters. Suppose the contact rate before the possible bathhouse closings is given by

$$(k\alpha)_{\text{overall}} = (1 - p + \tau p)(\gamma + \mu). \quad (3.4)$$

This represents an average contact rate for the two-activity model. We shall take $\mu = 1/(180 \text{ months})$ and $\lambda = 16,666 \text{ per month}$. (We are assuming a target population, absent the epidemic, of roughly 3,000,000.) For a given fraction π of infectives in the target population, we ask what is the ratio of contact rates causing elimination of the epidemic for the closings case divided by that without closings.

Figure 3.5 shows the ratio of contact rates (with closings relative to without closings) as a function of π for $p = 0.1$ and $\gamma = \frac{1}{60}$. It would appear that as long as the proportion of infectives π is no greater than 40% of the target population, there would be a significant benefit from bathhouse closings. The benefit decreases once we get to 40%. However, because of the fact that there appears to be a continuing influx of new entrants into the susceptible pool, there is good reason to close these establishments. Generally, restoring the sociological *status quo ante* is an effective means of stopping an epidemic; often this is difficult to achieve. Closing the bathhouses continues to be an appropriate action, even though a less effective one than if it had been taken early on in the history of the epidemic.

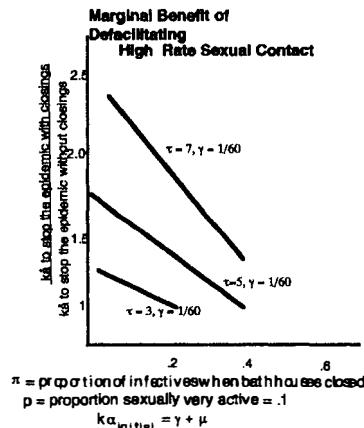


Figure 3.5. Effect of bathhouse closings in a mature epidemic.

How many Americans know that the USA has more AIDS cases than the rest of the First World combined? The highest American death rate in any of our wars was the War Between the States. AIDS in America has already killed more than the 600,000 combat dead from that War.

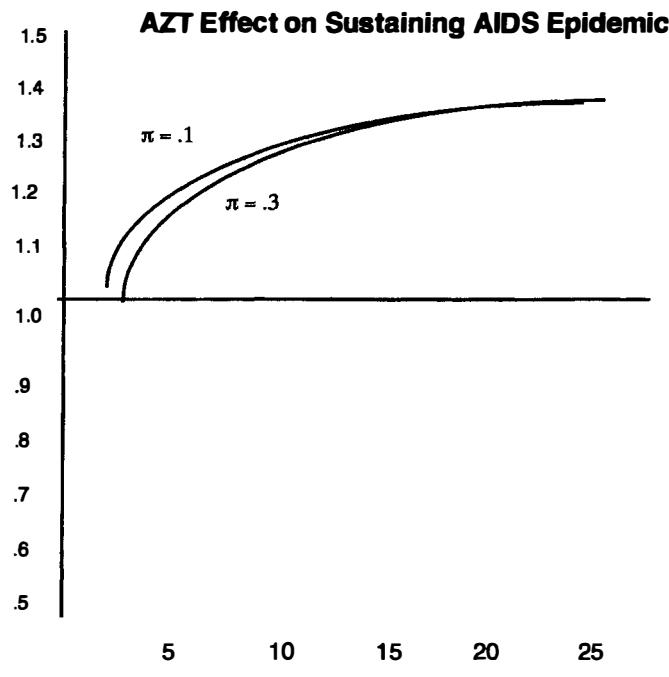


Figure 3.6. AZT and proportion of infectives.

Next, we look at the possible effects on the AIDS epidemic of administering a drug, such as AZT, to the entire infective population. Obviously, infectives who die shortly after contracting a contagious disease represent less of an enhancement to the viability of an epidemic than those who live a long time in the infective state. In the case of AIDS, it is probably unreasonable to assume that those who, by the use of medication, increase their T cell count to an extent where apparently normal health has returned, will decide to assume a chaste life style for the rest of their lives. We shall assume that the drug increases life expectancy by two years. Figure 3.6 demonstrates the change in the percent infective if the drug also increases the period of infectivity by two years for various proportions π of infective at the time that the drug is administered. The curves plot the ratio of the proportion infective using AZT to the proportion infective if AZT is not used (with $\gamma = 1/60$) and they asymptote to $1.4 = 84/60$, as should be the case. The greater pool of infectives in the target population can, under certain circumstances, create a kind of "Typhoid Mary" effect, where long-lived infectives wander around spreading the disease. Clearly, it should be the policy of health care professionals to help extend the time of quality life for each patient treated. However, it is hardly responsible to fail to realize that, by so doing, in the case of AIDS, there is an obligation of the treated

infective to take steps to ensure that he does not transmit the disease to susceptibles. To the extent that this is not the case, the highly laudable use of AZT to improve the length and quality of life for AIDS victims is probably increasing the number of deaths from AIDS.

3.10 AIDS as a Facilitator of Other Epidemics

In 1994 Webster West [11] completed a doctoral dissertation attempting to see to what extent AIDS could enhance the spread of tuberculosis in America. As we are primarily concerned here with the spread of AIDS itself, we shall not dwell very long on the tuberculosis adjuvancy issue. The reader is referred to relevant papers elsewhere [12,13].

West did discover that if one used stochastic process models and then took the mean trace, one obtained the same results as those obtained simply by using deterministic differential equation models. In the United States, since the Second World War at least, tuberculosis has been a cause of death mainly of the elderly (for example, Mrs. Eleanor Roosevelt died of it). Tuberculosis is carried by the air, and its epidemiological progression is enhanced by infected persons who are well enough to walk around in elevators, offices, and so on. When tuberculosis is confined to elderly persons, essentially not moving freely about, it is largely self-contained. But HIV infected persons are generally young, and generally employed, at least before the later stages of full-blown AIDS.

West discovered that the result of AIDS facilitating tuberculosis was likely to be only a few hundred additional deaths per year. His model further revealed that modest resources expended in the early treatment of persons infected with tuberculosis could bring even these relatively modest numbers down.

3.11 Comparisons with First World Countries

As noted in Section 3.4, the position of other developed countries toward defacilitating contacts between infectives and susceptibles was quite different from that in the United States. In a very real sense, these other countries can be used as a "control" when examining the epidemic in the United States. Good data for new cases did not become easier and easier to obtain as the epidemic progressed. Whereas in the earlier time span of the epidemic fairly good data for all First World countries could be obtained via "gopher" sites, increasingly it became more and more disconnected as data bases supposedly moved to the Internet. The reality was that the information on the gopher sites stayed in place but was not brought up to date, whereas data on the Internet appeared temporally disconnected. Great patience was required to follow a group of countries over a period of time, and because of holes in the data, it was not at all clear whether

anything but snippet comparisons could be made. I published one of these at a conference in 1989 [6], but the data available to me at the time gave only suggestions of what was happening. There seemed to be something important going on that went to the issue of the United States being a source of infection for other First World countries.

I kept sending out queries to the Centers for Disease Control and the World Health Organization (WHO), but without much success. Finally, in early 1998, Ms. Rachel Mackenzie of the WHO contacted me and provided me, not with a URL, but with the data itself, which was in the hands of the Working Group on Global HIV/AIDS, and STD Surveillance which is a joint Working Group between WHO and UNAIDS. I wish to acknowledge my gratitude to Ms. Mackenzie and her colleagues for allowing me to use their database.

Figure 3.7 shows the staggering differences in cumulative number of AIDS cases between the United States and France, Denmark, Netherlands, Canada, and the UK. The pool of infectives in the USA dwarfs those of the other First World countries. Whenever I would bring up the enormous differential between the AIDS rate in the United States and those in Europe, my European colleagues would generally attribute all this to a time lag effect. Somehow the United States had a head start on AIDS, but in time the European countries would catch up. If other First World countries were lagging the USA, then one would expect some sort of variation in new AIDS cases such as that depicted in Figure 3.8. However, Figure 3.9 demonstrates that the time lagging hypothesis is not supported by the data. No other First World country is catching up to the USA. Moreover, a downturn in new case rates is observable in all the countries shown.

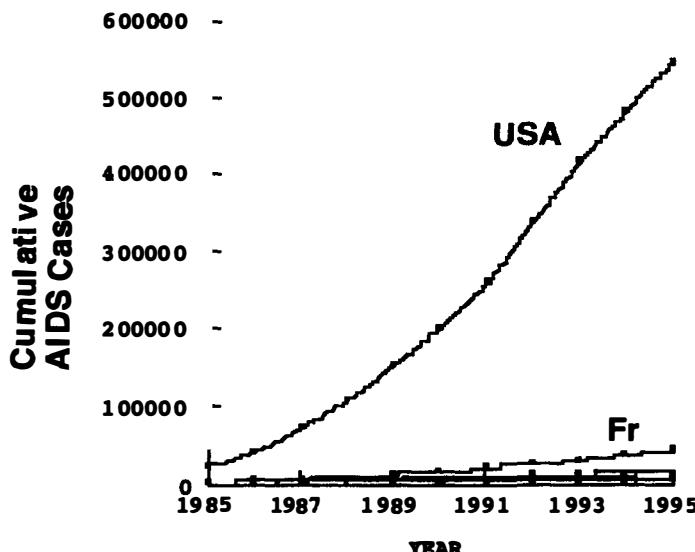


Figure 3.7. Cumulative AIDS cases 1985-1995.

If the USA Simply Leads the Rest of the First World in AIDS

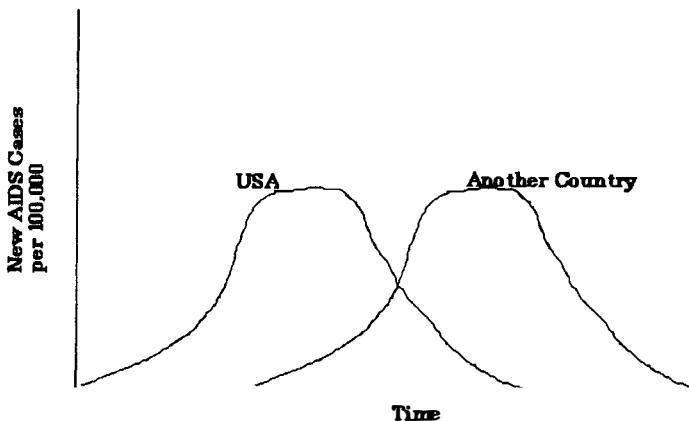


Figure 3.8. A time lagged scenario.

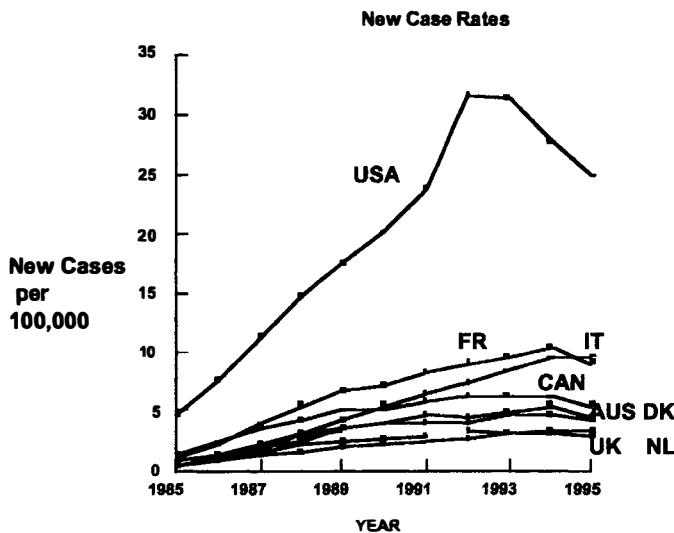


Figure 3.9. New case rates by country.

Additional insight is provided by Figure 3.10 in which we divide the annual incidence of AIDS per 100,000 in the USA by that for various other First World countries. Note the relative constancy of the new case ratio across the years for each country when compared to the United States. Thus, for the United Kingdom, it is around 9, for Denmark 6, etc. It is a matter of note that this relative constancy of new case rates is maintained over the period examined (eleven years). In a similar comparison, Figure

3.11 shows that the cumulative cases per 100,000 of AIDS in the United States divided by that for other First World countries gives essentially the same values observed for the new case rates in Figure 3.10.

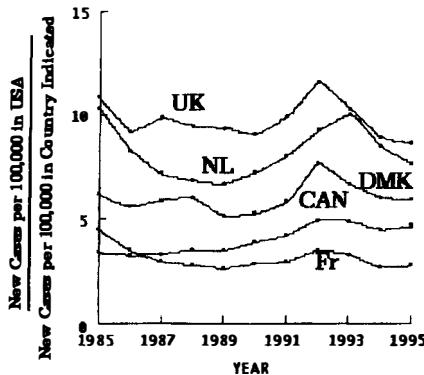


Figure 3.10. Comparative new case rates.

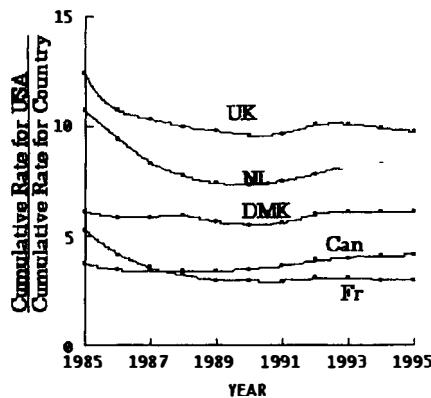


Figure 3.11. Comparative cumulative case incidence.

To investigate further, let us consider a piecewise in time exponential model for the number of AIDS cases, say in Country A:

$$\frac{dy_A}{dt} = k_A(t)y_A. \quad (3.5)$$

Figure 3.12 gives estimates for the rates k on a year-by-year basis using

$$k_A(t) \approx \frac{\text{new cases per year}}{\text{cumulative cases}}.$$

Note the apparent near equality of rates for the countries considered. To show this more clearly, Figure 3.13 displays the ratio of the annual estimated

piecewise national rates divided by the annual estimated rate of the United States.

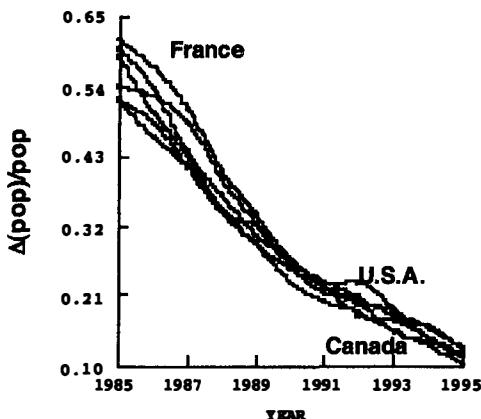


Figure 3.12. $k_{Country}(t)$ values.

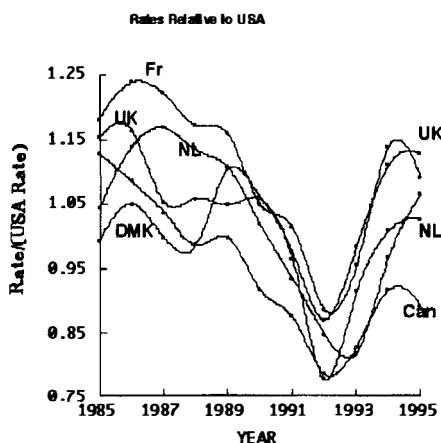


Figure 3.13. Ratios of piecewise rate estimates.

It is a matter of some interest that the k values are essentially the same for each of the countries shown in any given year. How shall we explain a situation where one country has a much greater incidence of new cases, year by year, yet the rate of increase for all countries is the same? For example, by mid-1997, the United Kingdom had a cumulative total of 15,081 cases compared to 612,078 for the United States. This ratio is 40.59 whereas the ratio of populations is only 4.33. This gives us a comparative incidence proportion of 9.37. However, at the same time, Canada had a cumulative AIDS total of 15,101. The United States population is 9.27 times that of Canada, so the comparative incidence proportion for the United States versus Canada in mid-1997 was 4.37. The comparative incidence of the

United States vis-a-vis the United Kingdom is over twice that of the United States vis-a-vis Canada. Yet, in all three countries the rate of growth of AIDS cases is nearly the same. This rate changes from year to year, from around 0.54 in 1985 to roughly 0.12 in 1995. Yet it is very nearly the same for each country in any given year. One could therefore predict the number of new cases in France in a given year, just about as well knowing the case history of the United States instead of that in France. The correlation of new cases for the United States with that for each of the other countries considered is extremely high, generally around 0.96. It is hard to explain this by an appeal to some sort of magical synchronicity, particularly since we have the fact that though the growth rates of AIDS in the countries are roughly the same for any given year, the new case relative incidence per 100,000 for the United States is several times that of any of the other countries.

Recall from Section 3.5 the conjecture made in the mid-80s that it was the bathhouses which caused the stand-alone epidemic in the United States. But, as we have seen, the bathhouse phenomenon really does not exist in the rest of the First World. How is it, then, that there are stand-alone AIDS epidemics in each of these countries? I do not believe there are stand-alone AIDS epidemics in these countries.

To model this situation, let us suppose there is a country, say Country Zero, in which the sociology favors a stand-alone AIDS epidemic. From other First World countries there is extensive travel to and from Country Zero, as indicated by Figure 3.14. If AIDS, with its very low infectivity rates, breaks out in Country Zero, then naturally the disease will spread to the other countries. But if the infectivity level is sufficiently low, then the maintenance of an apparent epidemic in each of the countries will be dependent on continuing visits to and from Country Zero.

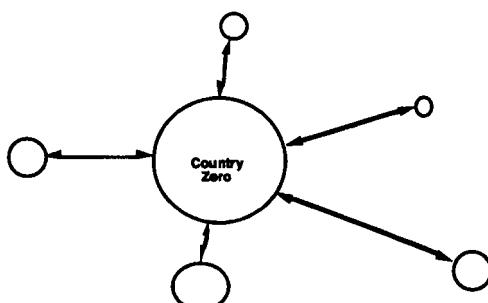


Figure 3.14. Country Zero.

Now let us suppose the fraction of infectives is low in country j . Thus, we shall assume that the susceptible pool is roughly constant. Let x_j be the number of infectives in country j and let z be the number of infectives in Country Zero. Let us suppose we have the empirical fact that, both for

Country Zero and the other countries, we can use the same β_t in the growth models

$$\frac{dz}{dt} = \beta_t z \quad (3.6)$$

$$\frac{dx_j}{dt} = \beta_t x_j. \quad (3.7)$$

Let the population of country j be given by N_j and that of Country Zero be given by N_Z . Suppose the new case rate in Country Zero divided by that for country j is relatively constant over time:

$$\frac{z/N_Z}{x_j/N_j} = c_j. \quad (3.8)$$

Let us suppose that, at any given time, the transmission of the disease in a country is proportional to both the number of infectives in the country and the number of infectives in Country Zero. Then from (3.7) and (3.8)

$$\frac{dx_j}{dt} = \alpha_{jt} x_j + \eta_{jt} z = \left(\alpha_{jt} + \frac{N_Z}{N_j} c_j \eta_{jt} \right) x_j = \beta_t x_j, \quad (3.9)$$

where α_{jt} and η_{jt} are the transmission rates into country j from that country's infectives and Country Zero's infectives, respectively. We are assuming that infectives from other countries will have relatively little effect on the increase of infectives in Country Zero. Thus, for a short time span, (3.6) gives

$$z(t) \approx z(0) e^{\beta_t t},$$

and (3.9) is roughly

$$\frac{dx_j}{dt} = \alpha_{jt} x_j + \eta_{jt} z(0) e^{\beta_t t}.$$

Now, we note that the epidemic in a country can be sustained even if α_{jt} is negative, provided the transmission from the Country Zero infectives is sufficiently high. If we wish to look at the comparative effect of Country Zero transmission on country j vis-a-vis country i , we have

$$\eta_{jt} = \frac{c_i}{c_j} \frac{N_j}{N_i} \eta_{it} + \frac{\alpha_{it} - \alpha_{jt}}{c_j} \frac{N_j}{N_Z}.$$

If for two countries i and j we have $\alpha_{it} = \alpha_{jt}$, then

$$\eta_{jt} = \frac{c_i}{c_j} \frac{N_j}{N_i} \eta_{it}.$$

Using (3.8) this can be expressed as

$$\frac{x_j}{x_i} = \frac{\eta_{jt}}{\eta_{it}}.$$

If η_{jt} doubles, then according to the model, the number of infectives in country j doubles.

Let us see what the situation would be in Canada if, as a stand alone, the epidemic is just at the edge of sustainability: i.e., $\alpha_{\text{Can},t} = 0$. Then, going back to a universal β_t for all countries including Country Zero (America), we have from (3.9) (using the c_{Can} value for 1995, 4.14).

$$\begin{aligned}\eta_{\text{Can},t} &= \frac{N_{\text{Can}}}{N_{\text{USA}}} \frac{1}{c_{\text{Can}}} \beta_t \\ &= \frac{26,832,000}{248,709,873} \frac{1}{4.14} \beta_t \\ &= 0.026 \beta_t.\end{aligned}$$

Thus, according to the model, activity rates from USA infectives roughly 2.6% of that experienced in the United States could sustain a Canadian epidemic at a comparative incidence ratio of around 4 to 1, United States to Canada. (If someone would conjecture that it is rather the Canadian infectives who are causing the epidemic in the United States, that would require the activity rate of Canadian infectives with American susceptibles to be $1/0.026 = 38.5$ times that of Canadian infectives with Canadian susceptibles.) If this activity rate would double to 5.2%, then the Canadian total infectives would double, but the rate $(1/x_{\text{Can}}) dx_{\text{Can}}/dt$ would still grow at rate β_t . Similar calculations show that

$$\begin{aligned}\eta_{\text{Fr},t} &= 0.076 \beta_t, \\ \eta_{\text{UK},t} &= 0.024 \beta_t, \\ \eta_{\text{DK},t} &= 0.0034 \beta_t, \\ \eta_{\text{NL},t} &= 0.0075 \beta_t.\end{aligned}$$

In summary, we have observed some surprises and have tried to come up with plausible explanations for those surprises. The relative incidence of AIDS for various First World countries when compared to that of the United States appears, for each country, to be relatively constant over time and this incidence appears to be roughly the same for cumulative ratios and for ratios of new cases. The rate of growth β_t for AIDS changes year by year, but it seems to be nearly the same for all the First World countries considered (Figure 3.12), including the United States. The bathhouse phenomenon is generally not present in First World countries other than the United States. Yet AIDS has a continuing small (compared to that of the United States), though significant, presence in First World countries other than the United States. The new case (piecewise exponential) rate there tracks that of the United States rather closely, country by country. We have shown that a model where a term for "travel" from and to the United States is dominant does show one way in which these surprises can be explained. Some years ago [3–10], I pointed out that the American gay community was made unsafe by the presence of a small subpopulation which visited

the bathhouses, even though the large majority of gays, as individuals, might not frequent these establishments. The present analysis gives some indication that the high AIDS incidence in the United States should be a matter of concern to other First World countries as long as travel to and from the USA continues at the brisk rates seen since the early 1980s.

Developing a model requires risk taking. The model, if it is to be useful, will be developed almost always without anything approaching a full data set. We could always find, as the fuller story comes in, that we were wrong. Then, in the case of epidemiology, we might find that by the time we publish our results, the virologists will have come up with a vaccine, perhaps rendering our model interesting but less than relevant. Most perilous of all, however, is to neglect the construction of a model.

3.12 Conclusions: A Modeler's Portfolio

This chapter has given an overview of around 25 years of my work on the AIDS epidemic. I did not treat this work as an academic exercise. Rather, by public talks, articles in the popular press, service on the ASA AIDS Committee, and meetings with public officials, I tried to change the public policy on the bathhouses, without effect. So it is correct to say that I have not been successful in influencing public policy as I had wished. I well recall, by the late 1980s certainly, that things were not going as I had wished.

I never had the experience of somebody getting up at a professional meeting and poking holes in my AIDS model. I would get comments like, "Well, we see that you have shown a plausible way that the epidemic got started. But that does us little good in providing a plan of action now that the epidemic is well under way." Of course, this statement is not correct, for two reasons. First of all, I have addressed what the effect of closing the bathhouses would be during the mature epidemic. Secondly, effective restoration of the *status quo ante* will, almost always, reverse the course of an epidemic. In the case of polio, for example, closing of the public swimming pools and the suburban cinemas would have greatly defacilitated the epidemic, even after it was well under way.

To my shock, some colleagues took me aside to say that AIDS might be a very good thing, since it was discouraging a lifestyle of which neither these colleagues nor I approved. I always responded that our obligation in health care was to improve the lives of all persons, whether we liked their lifestyles or not. Moreover, I noted that a continuing entry of young males into the sociologically defined gay communities showed that the discouragement induced by the dreadful deaths generally associated with AIDS was not working the way they supposed. For example, in Houston, most of the leadership of the gay community had died off by the early 1990s. The death toll in Houston was staggering, more than in all Canada which has

over ten times Houston's population. And yet, the people who died were replaced by a new wave of infectives.

Perhaps most significantly of all, I would hear amazement that my modeling research was receiving any government support since there seemed to be little statistical interest in such public policy consequential modeling. Vast sums had been spent, for example, in support of the design of procedures whereby blood samples could be anonymously dumped into a pool with that of, say, nine other individuals and this exercise repeated many times in such a way to determine the fraction of AIDS infectives in the United States, while ensuring the privacy of those tested. But modeling the progression of the epidemic was not receiving much NIH or PHS support. I was fortunate indeed that the Army Research Office has allowed me to work on modeling problems generally.

The notion of becoming some sort of full-time activist for modification of government policy toward defacilitating the epidemic was tempting. Some hold that, like an entrepreneur with a good idea for a product, the researcher should put all his/her energy into one enterprise at a time. Certainly, to save the hundreds of thousands of lives which have been needlessly lost to AIDS, such single-minded fanaticism would have been more than justified. However, based on the considerable effort that I had expended, it seemed to me that public policy was not going to be changed. If there had been some sort of focused attack on my AIDS model, then I might simply have hoped that a better explanation or a more complete model might win the day. But I had received the worst possible response: "We see your model, find no mistakes in it, and concede that it squares with the data, but it must be flawed because it does not square with policy."

So I continued my general career policy, which is somewhat similar to that of an investment portfolio. The basis of portfolio theory is that putting all of one's assets in one stock, even one with enormous expected return, is generally not a good idea. One is much better advised to use the weak law of large numbers and put one's capital in several enterprises of reasonably good expectation of return, so that the variability of the return of the overall portfolio will be brought down to much better levels than those associated with a single stock. It seems to me that this is a good idea for modeling researchers in allocating their intellectual assets.

During the period since the start of my work on AIDS, I founded the Department of Statistics at Rice, which now has 18 core faculty, 8 of them Fellows of the ASA. Again, during this period, I wrote eight books (AIDS figured in only three of these and only as chapters). I produced seven doctoral students during the interval, only one of these writing on AIDS. I managed to obtain United Nations funding to start a Quality Control Task Force in Poland following the fall of Russian domination of that country. I developed computer intensive strategies for simulation based estimation and continuous resampling, largely in connection with modeling work in cancer. I did a modest amount of consulting, saving in the process one or

two companies from bankruptcy. I started the development of anti-efficient market theory models which work fine as stochastic simulations, but cannot be handled in closed form. And so on. If AIDS was part of my professional "portfolio," it accounted for only, say, ten percent of the investment. I could have increased my efforts, but it became fairly clear that this was a battle which I could not win.

As I have so far been unable to find political support for closing down bathhouses in America, it could be argued that the AIDS modeling part of the portfolio was not productive. I disagree. Our business as modelers is, first of all, to understand the essentials of the process we are modeling. Only rarely, and generally in relatively simple situations, such as changing the quality control policy of a corporation, should we expect to be able to say, "There; I have fixed it."

The optimism concerning a quick discovery of an AIDS cure has dimmed. No doubt, one will be found at some time in the future. However, after tens of billions of dollars already expended without a cure or vaccine, it is unwise to continue on our present route of muddling through until a miracle occurs. By this time, so many hundreds of thousands of American lives have been wasted by not shutting down high contact facilitating establishments that changing policy could leave open a myriad of litigious possibilities. The families of the dead or dying might have good reason to ask why such policies were not taken 30 years ago. In the early 1980s I noted that AIDS might well kill more Americans than those killed in our bloodiest military conflict, the War between the States, around 600,000. It has already done so. Nobel laureate Joseph Stiglitz has argued, America's recent Middle Eastern Wars have cost the USA \$3 trillion. I doubt that the United States AIDS epidemic will cost less. The loss of over 600,000 lives and the productivity of those lives is huge.

Why did America allow itself to be drawn into the dreadful conflict in the Middle East? Certainly, one might opine, as I do, that it was due to powerful lobbies. But, regarding AIDS, the CDC's failing to take the most elementary epidemiological step of simply closing down the gay bathhouses as all other First World countries ultimately did, is beyond my comprehension. AIDS has no lobby. But political correctness does. I conjecture that Congressman Waxman's strictures simply frightened our public health officials into inactivity.

Modelers are not generally members of the political/economic power structure, which Pareto termed the "circle of the elites." We cannot ourselves hope to change public policy. But it is certainly our business to develop models that increase understanding of some system or other which appears to need fixing. We should follow the path of Chaucer's poor Clerk of Oxford: "...gladly would he learn and gladly teach."

Following the American polio epidemic of the postwar years, no modeler appears to have attempted to describe what went wrong with its management. Had that been done, perhaps a totally different response might have

taken place when AIDS came on the scene. At the very least, I hope that my modeling of AIDS will have some impact on public policy concerning the next plague when it comes, and come it surely will.

Problems

3.1. Using data from WHO, I have obtained estimates for the piecewise growth rates from equation (3.5) as shown in Table 3.5. Construct a bootstrap test to test the hypothesis that, year by year, the kinetic constant is the same for the United States as for the other countries shown.

Table 3.5. Estimates of Kinetic Constants.

Year	U.S.A.	U.K.	Canada	Denmark	France	Netherlands
1985	0.518	0.597	0.584	0.513	0.611	0.540
1986	0.459	0.535	0.498	0.482	0.569	0.523
1987	0.413	0.434	0.428	0.411	0.504	0.482
1988	0.347	0.367	0.341	0.342	0.406	0.393
1989	0.290	0.304	0.289	0.320	0.336	0.321
1990	0.251	0.265	0.230	0.266	0.263	0.256
1991	0.229	0.224	0.201	0.221	0.232	0.214
1992	0.232	0.202	0.182	0.180	0.205	0.196
1993	0.187	0.184	0.155	0.171	0.179	0.153
1994	0.143	0.158	0.131	0.144	0.163	0.138
1995	0.113	0.127	0.101	0.116	0.123	0.120

3.2. The combinatorics of finding the expected number of infectives created in the early days of an epidemic can quickly grow tedious. Moreover, it is very easy to make mistakes. Resampling gives us an easy way out. If there are n contacts to be spread among N individuals in a short period of time (say, the time of infectivity of the infectives), we may repeatedly take integer samples from 1 to N and count the fraction of times integers are repeated n or more times. Using this approach, if a total of 10 contacts are to be made, find the size of the susceptible pool which gives the largest number of expected infections, given that at least three contacts are required to convert a susceptible into an infective.

3.3. In Table 3.6 we show the ratio of the cumulative incidence of AIDS per 100,000 population for the United States divided by that for the United Kingdom, Canada, Denmark, France, and the Netherlands. Construct a resampling based test of the hypothesis that these ratios are constant over the 10 year period considered.

Table 3.6. Ratios of U.S.A. AIDS Incidences to Those of Other Countries

Year	U.K.	Canada	Denmark	France	Netherlands
1985	12.427	3.735	10.693	6.092	5.254
1986	10.695	3.468	9.432	5.831	4.193
1987	10.300	3.379	8.318	5.842	3.541
1988	9.982	3.405	7.727	5.888	3.219
1989	9.784	3.408	7.389	5.635	3.010
1990	9.597	3.505	7.346	5.521	2.960
1991	9.669	3.636	7.489	5.581	2.949
1992	10.048	3.871	7.833	5.954	3.050
1993	10.088	4.023	8.163	6.075	3.083
1994	9.904	4.080	8.208	6.067	3.012
1995	9.744	4.138	8.140	6.048	2.977

3.4. The assumption of a sexually very active subpopulation is, of course, not the only way to bring AIDS to epidemic levels. Redo Table 12.3 but make $\gamma = .01$. This scenario has increased the sexually active period of an AIDS infective from 10 months to 100 months.

3.5. Computers have become so fast, storage so plentiful, that we are tempted to dispense with differential equation aggregates and work directly with the underlying axioms. Such an approach was suggested in Chapter 3. Let us consider one not quite atomistic approach. Create a population of 300 susceptibles, 30 of whom have an activity level τ times that of the dominant population. Suppose that one (high activity) infective is introduced into the population. Keep track of all the members of the population as susceptible individuals S , “retired” susceptible individuals R , infective individuals I , and dead individuals D . See what τ needs to be to sustain the epidemic with high probability. At the beginning of the time interval $[t, t + \Delta t]$,

$$P(\text{new susceptible appears in } [t, t + \Delta t]) = \lambda \Delta t.$$

If such a person appears, we add him to the number of susceptibles, according to the proportion of .10 for high activity, .90 for low activity.

Then a susceptible may, for whatever reason, remove himself from the pool of risk.

$$P(\text{susceptible } J, \text{ “retires” in } [t, t + \Delta t]) = \mu \Delta t.$$

If this happens, we remove him from the infective pool and add him to the retired pool.

Next, an infective may die (or be so sick as to be inactive):

$$P(\text{an infective dies in } [t, t + \Delta t]) = \gamma \Delta t.$$

If this happens, we remove him from the pool of infectives and add him to the list of the dead.

Then, for each susceptible person,

$$P(\text{low-activity susceptible, converts to infective in } [t, t + \Delta t)) \\ = \frac{k\alpha\Delta t X_1 (Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)}.$$

If such a change is made, we add the individual to the pool of low-activity infectives removing him from the pool of low-activity susceptibles. Similarly,

$$P(\text{high-activity susceptible, converts to infective in } [t, t + \Delta t)) \\ = \frac{\tau k \alpha \Delta t X_1 (Y_1 + \tau Y_2)}{X_1 + Y_1 + \tau(Y_2 + X_2)}.$$

If such a conversion takes place, we remove the person from the pool of high-activity susceptibles, adding him to the pool of high-activity infectives. where

- τ = multiple of number of contacts of low-activity population for high-activity population;
- k = number of contacts per month;
- α = probability of contact causing AIDS;
- λ = immigration rate into sexually active gay population;
- μ = emigration rate from sexually active gay population;
- γ = marginal emigration rate from sexually active gay population due to sickness and death;
- X_1 = number of low-activity susceptibles;
- X_2 = number of high-activity susceptibles;
- Y_1 = number of low-activity infectives;
- Y_2 = number of high-activity infectives.

This problem may well indicate the reason that “higher order” languages are frequently not the choice for nontrivial simulations, which are generally DO-LOOP intensive. The running time for this program in FORTRAN or C is a tiny fraction of that required when the program is written in MATLAB or R. There is, naturally, no particular reason why this need be true. There is no reason why DO-LOOPS cannot be accommodated if only the compiler be written to do so.

3.6. There are many processes of the empirical birth-and-death variety related to those for epidemics. For example, there is the whole topic of simulating warfare. Suppose we have two sides, the Red and the Blue. Then we may ([4], pp. 55–71), if there are n subforces of Red, and m

subforces of Blue, write down the heterogeneous force Lanchester equations

$$\frac{du_j}{dt} = - \sum_{i=1}^n k_{ij} c_{1ij} v_i$$

$$\frac{dv_i}{dt} = - \sum_{j=1}^m l_{ji} c_{2ji} u_j,$$

where k_{ij} represents the allocation (between 0 and 1) such that $\sum_{j=1}^m k_{ij} \leq 1$ of the i th Red subforce's firepower against the j th Blue subforce. Also, c_{1ij} represents the attrition coefficient of the i th Red subforce against the j th Blue subforce; and similarly for l_{ji} and c_{2ji} . Write down the stochastic laws that one might use to simulate this system at the unit level (e.g., one company of Red tanks against two companies of Blue infantry). Such procedures for combat attrition were used since von Reiswitz introduced them in 1820 (with dice tosses and patient officers sitting around game boards). Interestingly, such games can easily be computerized, and their concordance with historical reality is excellent.

References

- [1] Rotello, G. (1997). *Sexual Ecology: AIDS and the Destiny of Gay Men*. New York: Dutton, 85–89.
- [2] Shilts, R. (1987). *And the Band Played On: Politics, People, and the AIDS Epidemic*. New York: St. Martin's Press, 144.
- [3] Thompson, J. R. (1984). “Deterministic versus stochastic modeling in neoplasia,” *Proceedings of the 1984 Computer Simulation Conference*. New York: North-Holland, 822–825.
- [4] Thompson, J. R. (1989). *Empirical Model Building*. New York: John Wiley & Sons, 79–91.
- [5] Thompson, J. R. and Tapia, R. A. (1990). *Nonparametric Function Estimation, Modeling and Simulation*. Philadelphia: SIAM, 233–243.
- [6] Thompson, J. R. (1989). “AIDS: the mismanagement of an epidemic,” *Comput. Math. Applic.*, 18, 965–972.
- [7] Thompson, J. R. and Go, K. (1991). “AIDS: modeling the mature epidemic in the American gay community,” in *Mathematical Population Dy-*

namics, O. Arino, D. Axelrod and M. Kimmel, eds. New York: Marcel Dekker, 371–382.

[8] Thompson, J. R.(1998). “The United States AIDS epidemic in First World context,” in *Advances in Mathematical Population Dynamics: Molecules, Cells and Man*, O. Arino, D. Axelrod, and M. Kimmel, eds. Singapore: World Scientific Publishing Company, 345–354.

[9] Thompson, J. R.(1999). “Understanding the AIDS epidemic: A modeler’s odyssey,” in *Applied Mathematical Modeling*, D. Shier and T. Wallenius, eds., New York: CRC Press, 41–69.

[10] Thompson, J. R.(2000). “Is the United States Country Zero for the First-World AIDS epidemic?” *The Journal of Theoretical Biology*, pp. 621–628.

[11] West, R. W. (1994). *Modeling the Potential Impact of HIV on the Spread of Tuberculosis in the United States*. Doctoral Dissertation. Houston: Rice University

[12] West, R. W. and Thompson, J. R. (1998). “Modeling the impact of HIV on the spread of tuberculosis in the United States,” *Math. Biosci.*, **143**, 35–60.

[13] West, R.W. and Thompson, J. R. (1998). “Models for the simple epidemic,” *Mathematical Biosciences*, **141**, 29–39.

Chapter 4

Bootstrapping

4.1 Introduction

The monumental work of Peter Bruce in the development of Julian Simon's *Resampling Stats* [7] can be downloaded from Statistics101.net [6] as free-ware provided you have version 1.4 or later of the Java Runtime Environment (JRE). You can download the latest version from www.java.sun.com.

The *bootstrap* is also available in most of the current computer languages. I have chosen to write bootstrapping programs in *Resampling Stats* because of its intuitive relation to actual problems.

Charles Darwin was the originator of the theory of evolution by natural selection. According to this theory, animals and plants which have superior survival characteristics are more likely to live to procreate than those with inferior survival characteristics. Sometimes, these superior survival characteristics are passed on to the next generation. And thus, over millenia, animals and plants are produced with superior survival characteristics. This theory, like any other, should be viewed critically and in the light of data.

Darwin carried out some experiments in which he tried to test the hypothesis that cross-fertilized corn plants produced higher stalks than self-fertilized ones. He had this data analyzed by his cousin Francis Galton, one of the founders of modern statistics (both Darwin and Galton were knighted by Queen Victoria for their scientific work). It turns out that Galton's analysis (which supported Darwin's conjecture) was seriously flawed. The data set was analyzed many years later by Ronald Fisher [5] (knighted by Queen Elizabeth II), and this time the analysis was much better. But Fisher made the assumption that the stalk heights observed would follow a particular probability distribution, the *normal* or *Gaussian*. One of the goals in this

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

book is to minimize prior assumptions on the distributions of data sets.

In the case of Darwin's corn data, we have several possible questions to be addressed. Naturally, although Darwin was looking specifically at a certain variety of a specific grain plant, the implicit question is more general: Is cross-fertilization good or bad? The common wisdom of most cultures is that it is likely a good thing. And utilizing almost any kind of analysis on the height based Darwin corn data, we arrive at an answer consistent with the common wisdom.

In the case of Darwin and Galton and Fisher, we see that the surrogate for "goodness" is stalk height at a fixed time after sprouting. It could have been otherwise. Darwin might have used other criteria (e.g., grain yield, resistance to disease and drought, flavor of grain, etc.). We cannot include every possible criterion. The sociologist/anthropologist Ashley Montague noted that by a number of criteria, sponges were superior to human beings. The selection of a criterion to be measured and analyzed almost always requires a certain amount of subjectivity.

In the case of the Darwin corn data, the null hypothesis selected is that there is no difference in stalk height between the cross-fertilized and self-fertilized plants. But when it comes to the alternative hypothesis the specificity is more vague. For example, we could select an alternative hypothesis model (à la Darwin, Galton, and Fisher) that cross-fertilization increases stalk height with variation from this rule being due to unexplained factors. Or, we could hypothesize that, on the average, cross-fertilization leads to increased stalk height. Or, we might opine that the median stalk height of a group of cross-fertilized plants tends to be greater than that of a group of self-fertilized plants. Each of these alternative hypotheses is different (although the first implies the next two). In the case of the Darwin data, each of the hypotheses seem to be supported by the data.

Selection of hypotheses is not easy, and the literature is replete with examples of studies where inappropriate hypotheses led to ridiculous or pointless conclusions. But an ivory tower disdain of specifying hypotheses leads, as a practical matter, to the radical position of Ashley Montague, where Shakespeare is no better than a sponge. That is more multiculturalism than is consistent with progress, scientific or otherwise.

4.2 Bootstrapping Analysis of Darwin's Data

In Table 4.1, we show the data essentially as presented by Darwin to Galton. We note here a "pot effect" and, possibly, an "adjacency within pot effect." At any rate, Darwin presented the data to Galton paired. Naturally, Darwin had made every attempt to equalize soil and water conditions across pots. It might well seem to us, in retrospect, that such equalization should be readily obtainable and that we might simply pool the data into 15 cross-fertilized and 15 self-fertilized plants.

Table 4.1

Case	Pot	Y Crossed X Self-Fertilized	Difference
1	I	23.500	6.125
2	I	12.000	-8.375
3	I	21.000	1.000
4	II	22.000	2.000
5	II	19.124	0.749
6	II	21.500	2.875
7	III	22.125	3.500
8	III	20.375	5.125
9	III	18.250	1.750
10	III	21.625	3.625
11	III	23.250	7.000
12	IV	21.000	3.000
13	IV	22.125	9.375
14	IV	23.000	7.500
15	IV	12.000	-6.000
		Sum	39.25

The sum of the differences for the 15 cases is rounded to 39.25. The evidence points toward a positive difference between the cross-fertilized and the self-fertilized, but we need some way of assessing how confident we can be that the difference is significantly greater than what one would expect from random choice. Let us assume that the height of a stalk is a result of three factors:

1. Cross-fertilized versus self-fertilized average effects, μ_{CF} and μ_{SF}
2. Pot effects ($\mu_I, \mu_{II}, \mu_{III}$, or μ_{IV}).
3. Random variation aka *noise* ϵ which is usually assumed to average out to zero.

So, if we were looking at cross-fertilized plant number 2, we might say that

$$CF(2) = \mu_{CF} + \mu_I + \epsilon.$$

Let us consider the differences:

$$d_2 = Y_2 - X_2 \tag{4.1}$$

$$= \mu_{CF} + \mu_I + \epsilon - (\mu_{SF} + \mu_I + \epsilon) \tag{4.2}$$

$$= \mu_{CF} - \mu_{SF} + \text{zero average noise} \tag{4.2}$$

and similarly for the other differences.¹ Then, if we look at the 15 differences, we can obtain

$$\bar{d} = \frac{1}{15} \sum_{j=1}^{15} d_j. \quad (4.3)$$

Now, let us carry out the following strategy. Pool the cross-fertilized and self-fertilized plants in each pot. So, Table 4.1 becomes

Table 4.2		
Pot	1	
I	23.500	17.375
I	12.000	20.375
I	21.000	20.000
II	22.000	20.000
II	19.124	18.375
II	21.500	18.625
III	22.125	18.625
III	20.375	15.250
III	18.250	16.500
III	21.625	18.000
III	23.250	16.250
IV	21.000	18.000
IV	22.125	12.750
IV	23.000	15.500
IV	12.000	18.000

1. We then sample three plants from Pot I and treat these as being “cross-fertilized.”
2. We sample three plants from Pot I and treat these as being “self fertilized.”
3. We pair the plants and take their differences.
4. We carry out similar operations for Pots II, III and IV.
5. We sum all the 15 differences so obtained.
6. We repeat the above 5000 times.

A program (zearandom.sta) in the *Resampling Stats* language to achieve this is given in Figure 4.1 below. (Recall that the program also runs on the freeware version at statistics101.net.)

¹Each ϵ is a different noise term. When we subtract two ϵ s, we get a zero average noise term which is as likely to be positive as negative.

'ze random.sta

'Within four pots, the cross-fertilized and self-fertilized
'plants are pooled within pots. So, from Pot I,
'we obtain a set of six stalk heights. We select randomly
'with replacement three of these and
'designate them as "cross-fertilized."
'We select randomly and with replacement
' three more plants from
'Pot I and designate them as "self-fertilized.
'We proceed similarly for Pots II, III, and IV. We then take the
'difference between the 15 "cross-fertilized" and "self-fertilized"
'plants. We repeat the above 5000 times,
' sort the 5000 differences
'so obtained and make a histogram of them.
' This gives us a picture
'of how the differences should look in the original data if there
'were no difference between cross-fertilized and self-fertilized
'heights.

```
maxsize default 10000
copy ( 23.500 12.000 21.000)A1
copy(17.375 20.375 20.000)B1
copy(22.000 19.124 21.500)A2
copy (20.000 18.375 18.625)B2
copy(22.125 20.375 18.250 21.625 23.250)A3
copy(18.625 15.250 16.500 18.000 16.250)B3
copy(21.000 22.125 23.000 12.000)A4
copy(18.000 12.750 15.500 18.000)B4
concat A1 B1 Pot1
concat A2 B2 Pot2
concat A3 B3 Pot3
concat A4 B4 Pot4
repeat 5000
sample 3 Pot1 P1
sample 3 Pot1 P2
sample 3 Pot2 P3
sample 3 Pot2 P4
sample 5 Pot3 P5
sample 5 Pot3 P6
sample 4 Pot4 P7
sample 4 Pot4 P8
subtract P1 P2 PP1
subtract P3 P4 PP2
subtract P5 P6 PP3
```

```

subtract P7 P8 PP4
concat PP1 PP2 PP3 PP4 POOL
sum POOL PSUM
Score PSUM F
Sort F S
Mean F K
Variance F L
END
histogram F

```

Figure 4.1. Pooled within pot (zearandom.sta.).

If the null hypothesis is true, then

$$\mu_{CF} = \mu_{SF}.$$

We recall that in the actual data set, the sum of the differences was 39.25. How many of the 5,000 differences where we treat all observations as having the same μ exceed 39.25 or are less than -39.25 ?

4.3 A Bootstrap Approximation to Fisher's Nonparametric Test

When this was done we obtained the histogram in Figure 4.2.

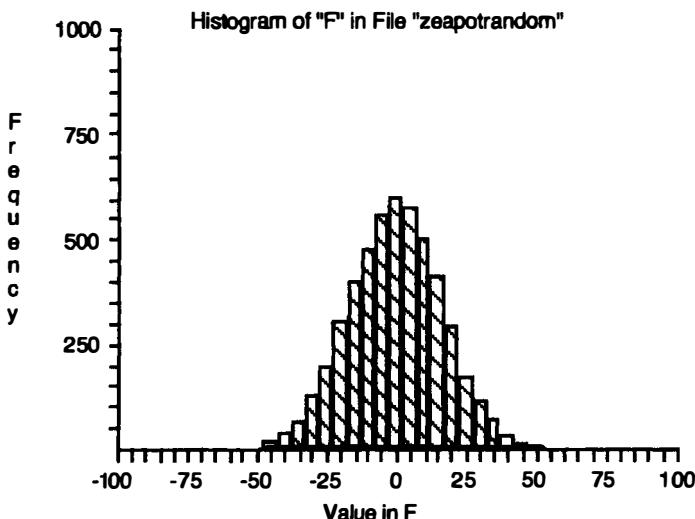


Figure 4.2. Resampled sum of differences with random allocation.

In only 56 of these was the sum of differences greater than 39.25 or less than -39.25 . Thus, our resampling test would reject the null hypothesis

at the $2 \times 56/5000 = .0224$ level of significance. Naturally, if we carried out our resampling plan again, we would obtain somewhat different results, for bootstrapping is based on random sampling. But you should not be surprised as you repeat the experiment several times, that you keep getting results which reject the null hypothesis. The evidence is overwhelmingly against it.

Bootstrapping is a wonderful mental exercise, for generally we can devise several different resampling strategies to test a null hypothesis. We have just gone through one such argument with the program given in Figure 4.1. Let us consider another quite natural way to proceed.

Now, it is a well-known fact that experimental investigators can have a tendency to present data to the statistician so as to promote the best chance for a significant argument in favor of significance of the data in support of an hypothesis. We note that the pairings in each pot might have been made in some way that this would be achieved. Better, then, to consider looking at differences in which the pairings are achieved randomly. That is to say, we do not know how Darwin elected to pair his observations within pot. But we have reason to suppose that the conditions for growing were rather uniform throughout each pot. Accordingly, we use the following resampling strategy (see *Resampling Stats* code in Figure 4.3):

Resampling Test Based on Within Pot Information

1. Sample (with replacement) three of the crossed plants in Pot I and then sample (with replacement) three of the self-fertilized plants in Pot I.
2. Compute the differences between crossed and self-fertilized plants in each run where pairings are done randomly.
3. Sum the differences.
4. Carry out similar resamplings from each of the four pots.
5. Compute the sum of differences.
6. Repeat 5000 times.

'zeapots.sta

```
' Sample (with replacement) three of the crossed plants in Pot I.
' Then sample (with replacement) three of the self-fertilized plants
' in 'Pot I.
' Compute the differences between crossed and self-fertilized plants
' in each run where pairings are done randomly.
' Sum the differences.
' Carry out similar resamplings from each of the four pots.
' Compute the sum of differences.
' Repeat 5000 times.
```

```

copy ( 23.500 12.000 21.000)A1
copy(17.375 20.375 20.000)B1
copy(22.000 19.124 21.500)A2
copy (20.000 18.375 18.625)B2
copy(22.125 20.375 18.250 21.625 23.250)A3
copy(18.625 18.625 15.250 16.500 18.000 16.250)B3
copy(21.000 22.125 23.000 12.000)A4
copy(18.000 12.750 15.500 18.000)B4
repeat 5000
sample 3 A1 C1
sample 3 B1 D1
sample 3 A2 C2
sample 3 B2 D2
Sample 5 A3 C3
sample 5 B3 D3
sample 4 A4 C4
sample 4 B4 D4
concat C1 C2 C3 C4 C
concat D1 D2 D3 D4 D
Subtract C D G
Sum G E
Score E F
END
Sort F FF
histogram FF

```

Figure 4.3. Within pots bootstrap program zeapots.sta.

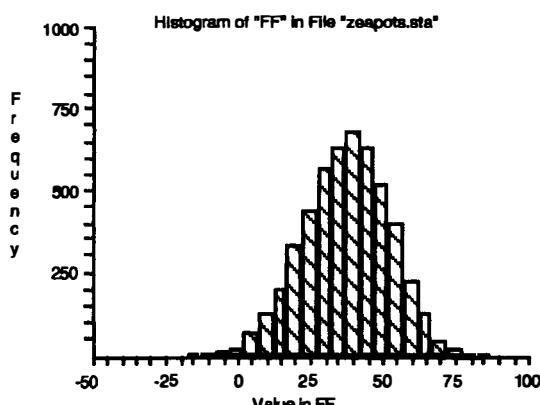


Figure 4.4. Resampled sum of differences with resampling within pots.

Only 6 of the 5000 simulations gave a sum less than zero. The assumption of equality of stalk heights would be rejected at the $2 \times 18/5000 = 0.007$ level. The mean of the difference distributions computed in this way is 39.48.

4.4 A Resampling-Based Sign Test

Let us suppose that one of the self-fertilized plants had been much taller. Suppose, for example, that the 12.75 value in Pot IV had been inflated by misrecording to 32.75. Then, running zeapots.sta with the indicated change, we find the histogram in Figure 4.5, and we note that 874 (or over 17%) of the resampled means are less than zero. The test is no longer significant.

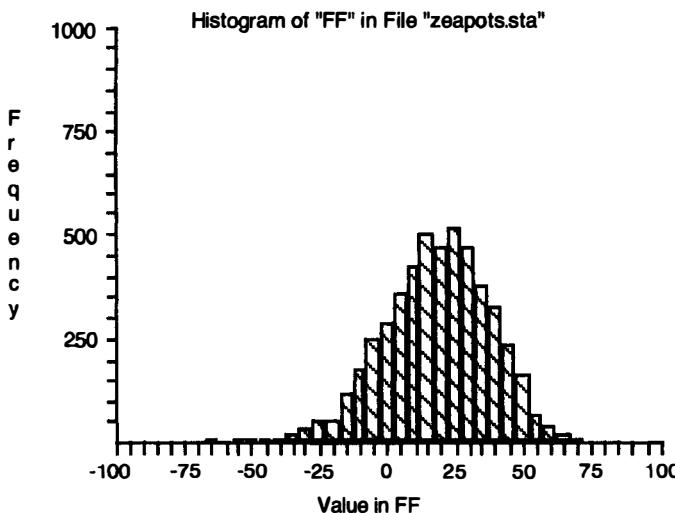


Figure 4.5. Within pots bootstrap with one inflated recording.

Suppose we disregard the sizes of the differences, relying completely on the signs of the differences. Every time the distance between a cross-fertilized plant and a self-fertilized plant in the same pot is positive, we score +1. Every time the difference is negative we score -1. Clearly, if there is no intrinsic difference in the heights of cross-fertilized and self-fertilized plants, then a plus and a minus are equally likely.

We will now use a program which is essentially the same as zeapots.sta except we shall code any positive difference as +1 and any negative difference as -1. Then, we note that if there really is no intrinsic difference, roughly half the sum of the 15 scores should be less than zero, half greater than zero.

```

' Sample (with replacement) three
' of the crossed plants in Pot I.
' Then sample (with replacement) three
' of the self-fertilized plants
' in Pot I.
' Compute the differences between crossed
' and self-fertilized plants
' in each run where pairings
' are done randomly.
' Score +1 for each positive difference
' and -1 for each negative difference.
' Sum the differences.
' Carry out similar resamplings from each of the four pots.
' Compute the sum of differences.
' Repeat 5000 times.
' Compute the histogram of the sums.
maxsize default 14000
copy ( 23.500 12.000 21.000)A1
copy(17.375 20.375 20.000)B1
copy(22.000 19.124 21.500)A2
copy (20.000 18.375 18.625)B2
copy(22.125 20.375 18.250 21.625 23.250)A3
copy(18.625 15.250 16.500 18.000 16.250)B3
copy(21.000 22.125 23.000 12.000)A4
copy(18.000 22.750 15.500 18.000)B4
repeat 5000
sample 3 A1 C1
sample 3 B1 D1
sample 3 A2 C2
sample 3 B2 D2
Sample 5 A3 C3
sample 5 B3 D3
sample 4 A4 C4
sample 4 B4 D4
concat C1 C2 C3 C4 C
concat D1 D2 D3 D4 D
Subtract C D E
Count E;0 G
Score G F
Sort F S
Mean F K
Variance F L
END
histogram F

```

Figure 4.6. Sign test zeasign.sta.

The resulting histogram is shown in Figure 4.7. Not one of the 5000 sum of differences is less than zero. Clearly, then, the null hypothesis which says there is no intrinsic difference in stalk heights due to cross-fertilized as opposed to self-fertilized is not supported by the resampling sign test applied to the data in Table 4.1, even when one of the self-fertilized observations has been inflated by 20.

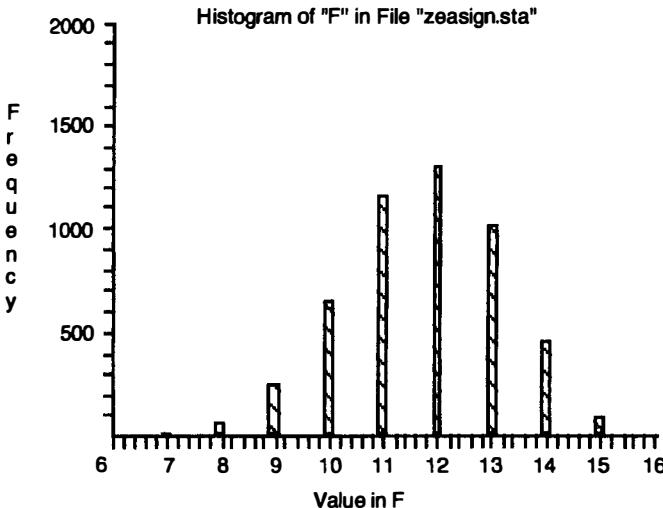


Figure 4.7. Bootstrapped sign test with one inflated observation.

4.5 A Bootstrapping Approach for Confidence Intervals

What can we say about the difference between the heights of cross-fertilized and self-fertilized *zea mays* plants? We have measured 15 differences and found that the total difference is 39.25. That means that our estimate of $d = \hat{\mu}_{CF} - \hat{\mu}_{SF}$ is given by

$$\hat{d} = \hat{\mu}_{CF} - \hat{\mu}_{SF} = \frac{39.25}{15} = 2.6167.$$

How sure are we that d is precisely equal to 2.6167? We would not be wise to claim that d is precisely equal to 2.6167, but we can make some statement of the sort “With 95% confidence we state that the true value of d lies between d_l and d_r .” To find this interval, we might first take our 15 observed differences and take a random sample (with replacement) of size 15, then compute the mean of the sample. We do this 5,000 times, rank the means so obtained from smallest to largest. Counting up from the smallest

to ranked mean 125 and down from the largest to ranked mean 4,875 we obtain our 95% confidence interval $15 \times \hat{d}$.

To obtain the 95% confidence interval on d , we need to take these values and divide by 15.

A program for doing this is given in `zeaconfint.sta` below.

```
'zeaconfint.sta
'We would like to be able to make a
'statement of the sort that says that
'if we repeated our corn stalk experiment
'many times, then with 95% the sum of the 15 differences will lie
'between x and y. We achieve this goal
'by resampling randomly from the 15 differences
'with replacement 10,000 times, then counting up
'to the 250th observation and down to the
'9,750th observation. There is a function in RS
'which automatically achieves this: Percentile
```

```
copy (6.125 -8.375 1.000 2.000 0.749 2.875 3.500 5.125 1.750
3.625 7.000 3.000 9.375 7.500 -6.000)A
```

```
repeat 10000
sample 15 A B
sum B C
score C G
End
histogram G
Percentile G(2.5 97.5)Z print Z
```

Figure 4.8. Zeaconfint.sta for histogram of $15 \times \hat{d}$.

Here, we obtain for $d_{.025}$ and $d_{.975}$, $2.000/15 = 0.1333$ and $71.500/15 = 4.767$, respectively. We can therefore express our informed opinion that with 95% probability, the true value of $d = \mu_{CF} - \mu_{SF}$ is between 0.1333 and 4.767.

4.6 Solving Ill-Structured Problems

The question posed by Darwin to his cousin Francis Galton was relatively well posed. In the real world, we have to make decisions based on poorly designed experiments and with ill-defined goals.

Table 4.3. Standardized Reading Scores.

Group A	Group B	Group C
3.5	2.5	-1.4
4.8	2.5	4.9
-4.4	1.7	5.3
-3.2	3.2	1.4
1.9	5.1	5.3
0.0	3.0	0.4
1.3	5.4	4.8
3.6	-1.1	8.2
4.0	1.5	9.8
0.5	1.4	5.4
0.3	0.5	-3.6
-2.8	1.2	3.8
-0.4	3.8	3.0
-2.5	1.6	6.0
5.4	5.4	5.4
0.0	0.2	3.1
-3.0	2.9	4.9
3.1	3.5	2.6
3.6	3.8	1.5
-0.2	2.4	9.9
-1.1	-0.2	4.3
-3.9		1.4
2.1		-4.1
5.7		4.9
1.7		
3.0		
-2.0		
0.7		
0.4		
-0.3		

Let us consider the following problem from the arena of public education. There is frequently concern over the methodology used for improving students' reading skills. Particularly in suburban school districts, there is an attempt to improve upon the standard methodology. In this instance, the statistician is confronted with data from three different ways of improving the reading skills of fourth graders. In one school, there were three different classes, each taught by different methodologies. The first method, call it **A**, is the standard that has been used for five years. The second method, call it **B**, differs only from the standard one in that a different textbook is used. The third method, call it **C**, is markedly different from the first two. In **C**, there is intervention in the form of extensive instruction, in small student groups, by doctoral students in education from a nearby state university.

Whereas **A** and **B** have essentially the same cost, method **C** would, absent the intervention for free by five doctoral students and their advisor, be very costly indeed.

At the end of the year, a standardized test is given to students from all three groups, and their scores are measured as departures from the scores of students in the past 5 years taking the same test. The scores are given in Table 4.3.

We have 30 tested students from **A**, 21 from **B**, and 24 from **C**. Of course, we can point out to the Board of Education that our task has not been made easier by the fact that there are three different teachers involved and the "teacher effect" is generally important. But that will not do us much good. The Board points out that it would have been practically impossible, under the limitations of size of the school, to have eliminated the teacher effect. We can, supposedly, take some comfort from the fact that the principal points out that the teachers, all similar in age and background, are uniformly "excellent." We are also assured that, in the interests of multicultural breadth, students are randomly mixed into classes from year to year. In any event, the data are as presented, and it is pointless to dwell too much on what might be done in the best of all possible worlds. Telling a client that "this job is just hopeless the way you have presented it to me" is not very good for keeping the consultancy going. Moreover, such an attitude is generally overly pessimistic.

One intuitive measure of the effectiveness of such programs is the mean score for each of the methods: 0.73, 2.4, and 3.64, respectively. The standard deviations are 2.81, 1.78, and 3.53, respectively.

The first thing we note is that the improvements, if any, are modest. The principal replies that such incrementalism is the way that improvements are made. She wishes to know whether or not the improvements for each of the two new methods, **B** and **C** are real or may simply be disregarded as due to chance. In Figure 4.9, we show a histogram of bootstrapped means using a resampling size of 10,000.

Now, the mindset of the bootstrapper is to consider the data to represent all reality, all possibilities that can ever occur. So, as a first step, we could compare means of samples (with replacement) of size 21 from **B** with means of samples (with replacement) of size 30 from the current pedagogy, **A**. Then, we similarly compare the new costlier pedagogy **C** with that of the current pedagogy **A**. Finally, we similarly compare the new costlier pedagogy **C** with that of the new less costly pedagogy **B**. The program for achieving these three figures is given in the program booktest.sta

'booktest.sta

'Here, we have three sets of test data:

'30 from the class using the old paradigm,

'21 from the class using new cheaper

'paradigm,

'24 from the class using the new
 'costly paradigm.
 'Since the data are rather numerous,
 'we write them as text files and
 'read them into the program.
 ' We take 10,000 random samples of
 'size 30 from the old paradigm,
 'computing the mean;
 'size 21 from the new cheaper paradigm,
 'computing the mean;
 'size 24 from the new costlier paradigm,
 'computing the mean.
 'For each of the 10,000 runs we compute
 'difference between mean of new cheap
 'and mean of old paradigm;
 'difference between mean of new costly
 'and mean of old;
 'difference between mean of new costlier
 'and mean of new cheaper paradigm.
 'We compute the fraction of times mean of
 'new cheaper paradigm minus mean of old paradigm
 ' is greater than 0 and call this sigcpold.
 'We compute the fraction of times mean of
 'new costlier paradigm minus mean of old
 'paradigm is greater than 0 and call
 'this sigcsold.
 'We compute the fraction of times mean of
 'new costlier paradigm minus mean of new
 'costlier paradigm is greater than 0 and
 'call this sigcscp.

```

read file "tabs" xold xcheap xcostly
repeat 10000
sample 30 xold oldsamp
sample 21 xcheap chpsamp
sample 24 xcostly costsamp
mean oldsamp oldmean
mean costsamp costmean
mean chpsamp chpmean
subtract chpmean oldmean dchpold
score dchpold zdchpold
subtract costmean chpmean dcoschp
score dcoschp zcoschp
subtract costmean oldmean dexold
score dexold zexold
  
```

```

end
histogram zdchpold
count zdchpold;0 sigcount
divide sigcount 10000 sigcpold
print sigcpold
count zexold;0 sigcount
divide sigcount 10000 sigcsold
print sigcsold
count zcoschp;0 sigcount
divide sigcount 10000 sigcscp
print sigcscp

```

Figure 4.9. Test for differences of the three paradigms.

In Figure 4.10 we show the histogram of differences between the two procedures using 10,000 resamplings. Here, denoting the resampled sample mean by μ , we note that the average performance of a resampled class using methodology **B** is greater than that of methodology **A** in over 91% of the runs. This gives us a bootstrapped “significance level” of .086, that is, the chances are only 8.6% that, if the methodologies were equally effective based on class average scores on the standardized test, that we would have seen a performance difference as large or larger than that which we have observed. Although significance levels are traditionally taken to be .05 or .01, significance levels should actually be adjusted to the reality of the situation. The cost of changing textbooks is marginal. The Board of Education might be well advised, therefore, to consider moving to methodology **B**.

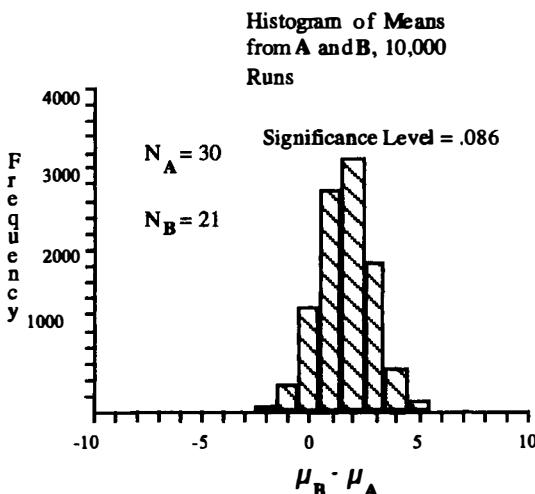


Figure 4.10. Means histogram of B versus A.

Next, let us investigate the resampling result when testing means of the costly procedure, C, versus that of the old standard, A. We demonstrate these results in Figure 4.11.

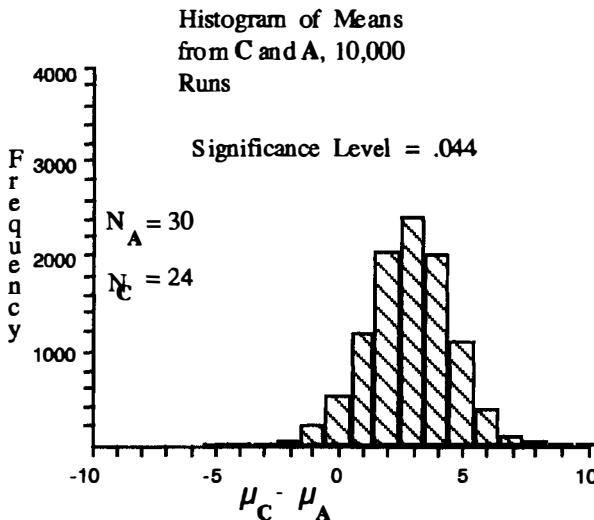


Figure 4.11. Means histogram of C versus A.

There seems little question as to the superiority of C to A. We see a bootstrap significance level of .044. On the other hand, we recall that C was an experimental, labor-intensive protocol that would be difficult to implement. Perhaps we should raise the question as to how much better it is than the cheap B protocol. We give the resampled comparisons of mean scores from C and B in Figure 4.12. If the two procedures had equal efficacy, then we would have observed a difference as great as we have observed 22.2% of the time.

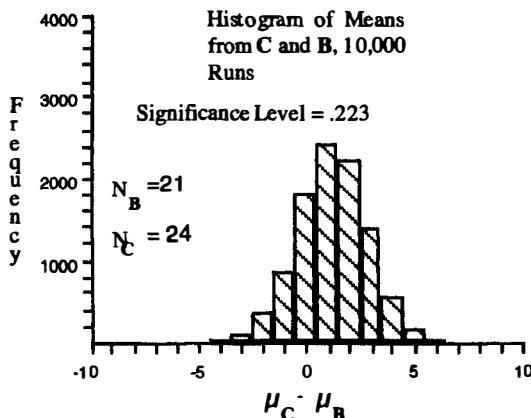


Figure 4.12. Means histogram of C versus B.

Our analysis does not give the Board of Education an unambiguous call. But it would appear that the cost of going to plan **B** (i.e., changing to the new textbook) may be the way to go.

Problems

4.1. An IQ test is administered to students from two high schools in the same city (Table 4.4). Using a bootstrap procedure, comment upon the conjecture that both groups **A** and **B** have the same underlying IQ.

Table 4.4. Intelligence Quotient Scores.

Group A	Group B
116.7	112.3
98.0	120.2
117.3	120.6
97.2	101.1
119.3	85.9
73.4	90.7
110.4	98.7
88.4	125.3
123.8	84.8
74.3	103.2
144.9	117.2
97.6	121.7
66.7	100.0
114.1	101.1
142.7	128.9
87.1	90.6
109.9	92.8
77.8	113.4
74.9	143.5
77.8	120.1
91.1	100.0
86.3	103.2
119.2	112.8
104.5	125.2
95.1	127.3
106.9	127.9
84.6	147.9
99.3	
96.9	
77.6	

4.2. For the two schools selected, the conjecture is made that there is a significant correlation between IQ and family income (measured in thousands of dollars). Use the data in Table 4.5 to obtain a bootstrap procedure for testing the conjecture.

Table 4.5. Intelligence Quotient Scores.

GroupA	Income	GroupB	Income
116.7	15.1	112.3	21.4
98.0	20.3	120.2	66.2
117.3	25.7	120.6	45.1
97.2	56.3	101.1	23.1
119.3	45.2	85.9	19.1
73.4	70.2	90.7	22.1
110.4	19.1	98.7	21.1
88.4	14.2	125.3	45.2
123.8	72.4	84.8	11.1
74.3	14.2	103.2	74.1
144.9	97.3	117.2	44.1
97.6	36.0	121.7	97.2
66.7	13.2	100.0	23.1
114.1	36.1	101.1	19.3
142.7	19.1	128.9	35.6
87.1	44.7	90.6	22.1
109.9	55.1	92.8	13.1
77.8	72.1	113.4	23.8
74.9	15.1	143.5	101.3
77.8	13.9	120.1	87.1
91.1	19.1	100.0	44.4
86.3	56.2	103.2	28.1
119.2	34.1	112.8	10.5
104.5	45.1	125.2	36.7
95.1	24.8	127.3	12.3
106.9	16.2	127.9	28.1
84.6	23.1	147.9	15.2
99.3	18.1		
96.9	39.9		
77.6	15.2		

4.3. Now, even before estimating μ_B and μ_C in Table 4.3, we see that adding on the graduate student teaching assistants appears to have effect. However, it is possible that this appearance is simply the result of randomness. Moreover, the adding on of the teaching assistants is not without cost. Orthodox Bayesians are generally reluctant to construct significance tests. Construct a resampling procedure to determine if it is realistic to assume that the add-on of graduate student teaching assistants is of no benefit.

References

- [1] Block, J. (1960). "On the number of significant findings to be expected by chance," *Psychometrika*, **25**, 158–168.
- [2] Box, G. E. P. and Draper, N. R. (1987). *Empirical Model-Building and Response Surfaces*. New York: John Wiley & Sons, 84–90.
- [3] Dwass, M. (1957). "Modified randomization sets for nonparametric hypotheses," *Annals of Mathematical Statistics*, **29**, 181–187.
- [4] Efron, B. (1979). "Bootstrap methods: another look at the jackknife," *Annals of Statistics*, **8**, 1–26.
- [5] Fisher, R. A. (1935). *The Design of Experiments*. Edinburgh: Oliver & Boyd, 29–49.
- [6] Grosberg, J. (2010). *Introduction to Resampling Statistics Using Statistics101.net*.
- [7] Simon, J. L. (1990). *Resampling Stats*. Arlington, Va.: Resampling Stats, Inc.

Chapter 5

Monte Carlo Solutions of Differential Equations

5.1 Introduction

In a real sense, computer science was started by the Polish Enigma Code breaker Marian Rejewski, working with his colleagues Jerzy Rozycki. and Henryk Zygalski. A true empirical model builder, Rejewski used several exploratory techniques to reduce the number of the Kriegsmarine's code combinations from an impossible 10^{92} to a manageable 10^6 . The German Enigma machine had no electric power. Rejewski and his associates built a decoding device which did. Their programming operating system might be regarded as a precursor to UNIX. And they completed their prototype in 1932, a year before Hitler came to power in Germany. The Polish government shared their decryption device with its English and French "allies" before the start of the Second World War. Without the decryption, there is little doubt England would have been starved into submission by the German submarine fleet. Successful in keeping their possession of an Enigma breaking device secret, the English were even more successful in keeping the secret of who had built it.

As cryptanalysis is somewhat exotic, we can say that when it comes to the equations of nuclear reactors and other equations of applied engineering, the Hungarian John von Neumann could be said to be the "founder of computer science." Unlike the Enigma decoder, von Neumann's computer worked very much like the computers we still use today. von Neumann conceived of and built the first serious digital computer as a device for handling simulation algorithms that he had formulated for dealing with problems in

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

nuclear engineering.¹ Ideally, if we are dealing with problems of heat transfer, neutron flux, and so on, in regular and symmetrical regions, the classical nineteenth and early twentieth century differential-integral-difference equation formulations can be utilized. However, if the regions are complicated, if indeed we are concerned about a maze of pipes, cooling vessels, rods, and so on, the closed-form solutions are not available. This means that many person-years would be required to come up with all the approximation-theoretic quadrature calculations to ensure that a satisfactory plant will result if the plans are implemented. von Neumann noticed that if large numbers of simple repetitive computations could be readily performed by machine, a method could be devised which would serve as an alternative to quadrature.

In reality, the quadrature issue, which Monte Carlo was largely developed to address, is rather unimportant compared to the much more important issue of direct simulation. To make a distinction between Monte Carlo and simulation, let us consider the following two paradigms shown in Figure 5.1. In the upper flowchart, we note a traditional means of coping with the numerical results of a model. We start out with axioms at the micro level which are generally easily understood. For example, one such axiom might be that a gas particle starts at a particular point and moves step by step in three-space according to specified laws until it collides with a wall. Dealing with each specific gas molecule out of a total of, say, 10^{12} molecules is a hopeless task. Thus, investigators in the nineteenth century quite naturally and correctly were led to means for summary information about the gas molecules. That is to say, they had to content themselves with differential-integral-equation models as average representations of the effect of trillions of molecules.

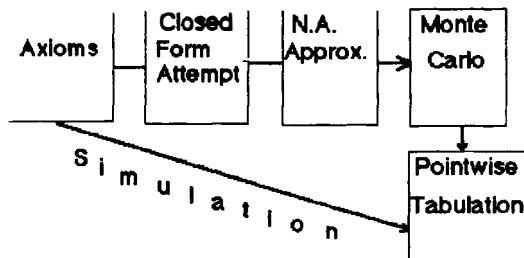


Figure 5.1. Two ways of problem solving.

In Figure 5.1 the upper path gives the paradigm for solving such problems based on precomputer age models. We start with axioms which are accepted by most investigators in the field. These are transformed into a differential-integral-difference-equation type of summary model. Then,

¹The discussion in this chapter largely follows [3].

a generally *pro forma* attempt is made to arrive at a closed-form solution, that is, a representation which can be holistically comprehended by an observer and which lends itself to precise numerical evaluation of the dependent variables as we change the parameters of the model and the independent variables. This attempt is generally unsuccessful and leads only to some nonholistic quadrature-like setup for numerical evaluation of the independent variables. If the dimensionality of the quadrature is greater than 2, the user moves rather quickly to a random quadrature Monte Carlo approach. What would have happened had computers been developed a century before they were? Would differential-integral equation modeling be the backbone of so much of physical science the way it still is today? It is an open question.

The fact is that we now have the computer speed to use the algorithm in the lower part of the diagram. We can now frequently dispense with the traditional approach by one which goes directly from the microaxioms to pointwise evaluation of the dependent variables. The technique for making this "great leap forward" is, in principle, simplicity itself.

Simulation carries out that which would earlier have been thought to be impossible, namely, to follow the progress of the particles, the cells, whatever. We do not do this for all the particles, but for a representative sample. We still do not have the computer speed to deal with 10^{10} particles; but we can readily deal with, say, 10^4 or 10^5 . For many purposes, such a size is more than sufficient to yield acceptable accuracy. Among the advantages of a simulation approach is principally that it enables us to eliminate time-consuming and artificial approximation-theoretic activities and spend our time in more useful pursuits.

More importantly, simulation enables us to deal with problems which are so complex in their "closed-form" manifestation that they are presently attacked only in *ad hoc* fashion. For example, econometric approaches are frequently linear, not because such approaches are supported by microeconomic theory, but because the complexities of dealing with the nonlinear consequences of the microeconomic theory are so overwhelming. Similarly, in mathematical oncology, the use of linear models is motivated by the failure of the natural branching process models to lead to numerically approximateable closed forms.

We have long since passed the point where computers can enable us to change fundamentally the ways we pose and solve problems. We have had the hardware capabilities for a long time to implement all the techniques covered in this chapter. But the proliferation of fast computing to the desktop will encourage private developers to develop simulation-based procedures for a large and growing market of users who need to get from specific problems to useful solutions in the shortest time possible. We now have the ability to use the computer, not as a fast calculator, but as a device which changes fundamentally the process of going from the microaxioms to the macrorealization.

5.2 Gambler's Ruin

There is an old temptation in applied mathematics to pose new problems, whenever possible, in classical “toy problem” formulation. One such is that of “gambler’s ruin” [1]. We consider two gentlemen gamblers, A and B, who start to gamble in a zero-sum game with stakes x and $b - x$, respectively. At each round, each gambler puts up a stake of h dollars. The probability that A wins a round is p , while the probability that B wins a round is $q = 1 - p$. We wish to compute the probability that A ultimately wins the game. Let us define $v(x, t)$ to be the probability that A wins the game starting with capital x on or before the t th round. Similarly, $u(x, t)$ is the probability that B wins the game with his stake of $b - x$ on or before the t th round. Let $w(x, t)$ be the probability the game has not terminated by the t th round.

Each of the three variables u , v , and w is bounded below by zero and above by one. Moreover, u and v are nondecreasing in t . w is nonincreasing in t . Thus, we can take limits of each of these as t goes to infinity. We shall call these limits $v(x)$, $u(x)$, and $w(x)$, respectively.

Although we shall briefly digress to get the closed-form solution to gambler’s ruin, such a solution is really unimportant for our differential-integral simulation purposes. It will be the fundamental recursion in (5.1), which will be the basis for practically everything we do in this section.

$$v(x, t) = pv(x + h, t - \lambda) + qv(x - h, t - \lambda). \quad (5.1)$$

That is, the probability A wins the game on or before the t th round is given by the probability that he wins the first round and then ultimately wins the game with his new stake of $x + h$ in $t - \lambda$ rounds plus the probability he loses the first round and then wins the game in $t - \lambda$ rounds with his new stake of $x - h$. Here we have used a time increment of λ .

Taking limits in (5.1), we have

$$(p + q)v(x) = pv(x + h) + qv(x - h). \quad (5.2)$$

Rewriting (5.2), we have

$$p[v(x + h) - v(x)] = q[v(x) - v(x - h)]. \quad (5.3)$$

Let us make the further simplifying assumption that $b = Nh$. Then

$$v(\{n + 1\}h) - v(nh) = (q/p)[v(nh) - v(\{n - 1\}h)]. \quad (5.4)$$

Notes that $v(0) = 0$ and $v(Nh) = 1$. So writing (5.4) *in extenso*, we have

$$\begin{aligned} v(nh) - v(\{n - 1\}h) &= (q/p)[v(\{n - 1\}h) - v(\{n - 2\}h)] \\ v(\{n - 1\}h) - v(\{n - 2\}h) &= (q/p)[v(\{n - 2\}h) - v(\{n - 3\}h)] \\ &\dots = \dots \end{aligned}$$

$$\begin{aligned}
 \dots &= \dots \\
 \dots &= \dots \\
 v(2h) - v(h) &= q/p[v(h) - v(0)] = (q/p)v(h).
 \end{aligned} \tag{5.5}$$

Substituting up the ladder, we have

$$v(nh) - v(\{n-1\}h) = (q/p)^{n-1}v(h). \tag{5.6}$$

Substituting (5.5) in the extenso version of (5.4), we have

$$\begin{aligned}
 v(nh) - v(\{n-1\}h) &= (q/p)^{n-1}v(h) \\
 v(\{n-1\}h) - v(\{n-2\}h) &= (q/p)^{n-2}v(h) \\
 \dots &= \dots \\
 \dots &= \dots \\
 \dots &= \dots \\
 v(h) &= v(h).
 \end{aligned} \tag{5.7}$$

Adding, we have

$$v(nh) = \left(1 + \frac{q}{p} + \left(\frac{q}{p}\right)^2 + \dots + \left(\frac{q}{p}\right)^{n-1}\right)v(h). \tag{5.8}$$

Recalling that $v(Nh) = 1$, we have

$$1 = \left(1 + \frac{q}{p} + \left(\frac{q}{p}\right)^2 + \dots + \left(\frac{q}{p}\right)^{N-1}\right)v(h) \tag{5.9}$$

Thus, we have

$$v(nh) = \frac{1 + q/p + (q/p)^2 + \dots + (q/p)^{n-1}}{1 + q/p + (q/p)^2 + \dots + (q/p)^{N-1}} \tag{5.10}$$

For $p = q = .5$, this gives,

$$v(nh) = \frac{n}{N}; \text{ i.e., } v(x) = \frac{x}{b}. \tag{5.11}$$

Otherwise, multiplying (5.10) by $[1 - p/q]/[1 - p/q]$, we have

$$\begin{aligned}
 v(nh) &= \frac{1 - (q/p)^n}{1 - (q/p)^N}, \quad \text{i.e.,} \\
 v(x) &= \frac{1 - (q/p)^{x/h}}{1 - (q/p)^{b/h}}.
 \end{aligned} \tag{5.12}$$

Now, by symmetry,

$$u(x) = \frac{(q/x)^{x/h} - (q/h)^{b/h}}{1 - (q/p)^{b/h}} \tag{5.13}$$

From (5.12) and (5.13), we have

$$v(x) + u(x) = 1. \quad (5.14)$$

Consequently, $w(x) = 0$; that is, the game must terminate with probability 1. Thus, we can use a simulation to come up with reasonable estimates of the probability of A ultimately winning the game. A flowchart of such a simulation is given in Figure 5.2.

We note that this simulation gives us a ready means of estimating a rough 95% confidence interval for $v(x)$, namely

$$v(x) = \frac{W}{M} \pm \frac{2\sqrt{W(1 - W/M)}}{M}. \quad (5.15)$$

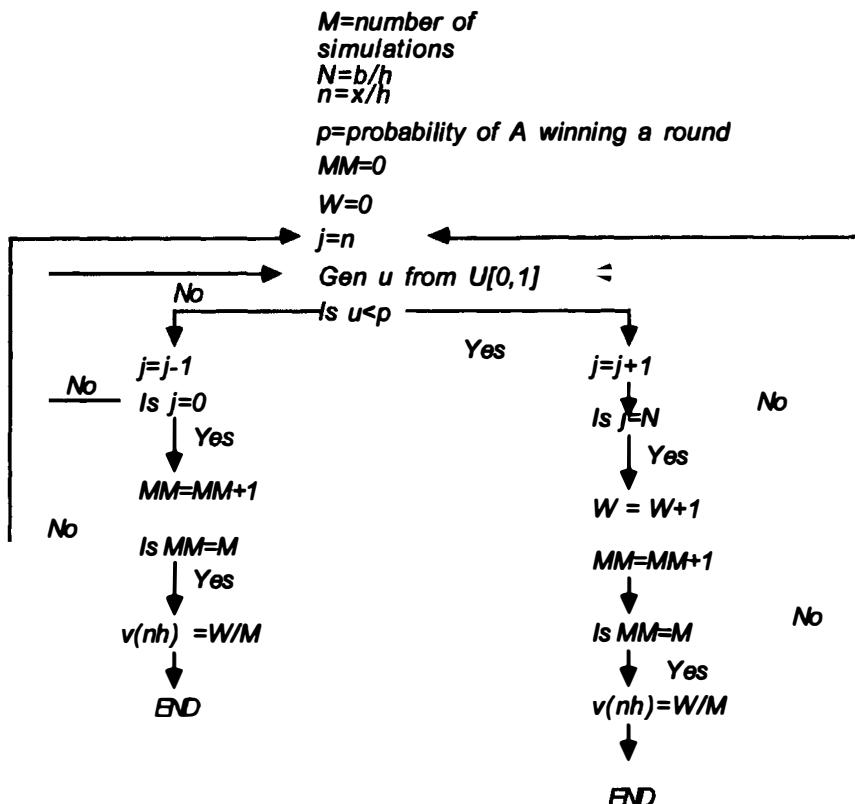


Figure 5.2. Gambler's ruin.

5.3 Solution of Simple Differential Equations

Since we have shown a closed-form solution for the gambler's ruin problem, it would be ridiculous for us to use a simulation to solve it. It is by means of an analogy of real-world problems to the general equation (5.1) that simulation becomes useful. Rewriting (5.3), we have

$$p\Delta v(x) = q\Delta v(x - h), \quad (5.16)$$

where

$$\Delta v(x) = [v(x + h) - v(x)]/h.$$

Subtracting $q\Delta v(x)$ from both sides of (5.16), we have

$$(p - q)\Delta v(x) = q[\Delta v(x - h) - \Delta v(x)], \quad (5.17)$$

or

$$\Delta^2 v(x) + \frac{p - q}{qh} \Delta v(x) = 0, \quad (5.18)$$

where

$$\Delta^2 v(x) = \frac{\Delta v(x) - \Delta v(x - h)}{h}.$$

For h sufficiently small, this is an approximation to

$$\frac{d^2 v}{dx^2} + 2\beta \frac{dv}{dx} = 0, \quad (5.19)$$

where

$$\frac{p - q}{qh} = 2\beta.$$

Now, suppose that we are given the boundary conditions of (5.19), $v(0) = 0$ and $v(b) = 1$. Then our flowchart in Figure 5.2 gives us a ready means of approximating the solution to (5.19). We simply set $p = (2\beta h + 1)/(2\beta h + 2)$, taking care to see that h is sufficiently small, if β be negative, to have p positive. To make sure that we have chosen h sufficiently small that the simulation is a good approximation to the differential equation, typically we use simulations with successively smaller h until we see little change in $v(x) \approx W/N$.

Suppose that the boundary conditions are less accommodating, for example, suppose that $v(0)$ and $v(b)$ take arbitrary values. A moment's reflection shows that

$$v(x) \approx \frac{W}{N}v(b) + \left(1 - \frac{W}{N}\right)v(0). \quad (5.20)$$

As a closed-form solution of (5.19) is readily available. We need not consider simulation for this particular problem. But suppose that we generalize (5.19) to the case where β depends on x :

$$\frac{d^2 v}{dx^2} + 2\beta(x) \frac{dv}{dx} = 0. \quad (5.21)$$

Again, we use our flowchart in Figure 5.2, except that at each step we change p via

$$p(x) = \frac{2\beta(x)h + 1}{2\beta(x)h + 2}. \quad (5.22)$$

Once again, $v(x) \approx W/Nv(b) + (1 - W/N)v(0)$. And once again, it is an easy matter to come up with an internal measure of accuracy via (5.15).

It is possible to effect numerous computational efficiencies. For example, we need not start afresh for each new grid value of x . For each pass through the flowchart, we can note all grid points visited during the pass and increase the counter of wins at each of these if the pass terminates at b , the number of losses if the pass terminates at 0.

5.4 Solution of the Fokker–Planck Equation

It is important to note that the simulation used to solve (5.19) actually corresponds, in many cases, to the microaxioms of which the differential equation (5.19) is a summary. This is very much the case for the Fokker–Planck equation which we consider below. Let us suppose that we do not eliminate time in (5.1). We will define

$$\begin{aligned} v(x, t, 0; h, \lambda) &= P[\text{particle starting at } x \text{ will be absorbed at 0} \\ &\quad \text{on or before } t = m\lambda]; \\ v(x, t, b; h, \lambda) &= P[\text{particle starting at } x \text{ will be absorbed at } b \\ &\quad \text{on or before } t = m\lambda]; \\ V(x, t; h, \lambda) &= V(0, t)v(x, t, 0; h, \lambda) + V(b, t)v(x, t, b; h, \lambda). \end{aligned}$$

We define

$$\begin{aligned} \Delta_t V(x, t; h, \lambda) &= \frac{V(x, t + \lambda; h, \lambda) - V(x, t; h, \lambda)}{\lambda}, \\ \Delta_x V(x, t; h, \lambda) &= \frac{V(x + h, t; h, \lambda) - V(x, t; h, \lambda)}{h}, \\ \Delta_{xx} V(x, t; h, \lambda) &= \frac{\Delta_x V(x, t; h, \lambda) - \Delta_x V(x - h, t; h, \lambda)}{h}. \end{aligned}$$

That is, the expected payoff is given by the probability a particle is absorbed at the left at time t , multiplied by the boundary award $V(0, t)$ plus the probability the particle is absorbed at the right at time t times the boundary award $V(b, t)$.

Now, our basic relation in (5.1) still holds, so we have

$$V(x, t + \lambda; h, \lambda) = p(x)V(x + h, t; h, \lambda) + q(x)V(x - h, t; h, \lambda). \quad (5.23)$$

Subtracting $V(x, t; h, \lambda)$ from both sides of (5.23), we have

$$\begin{aligned}
\lambda \Delta_t V(x, t; h, l) &= p(x)[V(x + h, t; h, \lambda) - V(x, t; h, \lambda)] \\
&\quad + q(x)[V(x - h, t; h, \lambda) - V(x, t; h, \lambda)] \\
&= p(x)[V(x + h, t; h, \lambda) - V(x, t; h, \lambda)] \\
&\quad - q(x)[V(x, t; h, \lambda) - V(x - h, t; h, \lambda)] \\
&= hp(x)\Delta_x V(x, t; h, \lambda) - hq(x)\Delta_x V(x - h, t; h, \lambda) \\
&= h[p(x) - q(x)]\Delta_x V(x, t; h, \lambda) \\
&\quad + hq(x)[\Delta_x V(x, t; h, \lambda) - \Delta_x V(x - h, t; h, \lambda)] \\
&= h[p(x) - q(x)]\Delta_x V(x, t; h, \lambda) + h^2 q(x)\Delta_{xx} V(x, t; h, \lambda).
\end{aligned} \tag{5.24}$$

Letting $p(x) = [\beta(x) + 2h\alpha(x)]/[2\beta(x) + 2h\alpha(x)]$ and $q(x) = 1 - p(x)$, we have

$$\lambda \Delta_t V(x, t; h, \lambda) = 2h^2 \frac{\alpha(x)}{2\beta(x) + 2h\alpha(x)} + h^2 \frac{\beta(x)}{2\beta(x) + 2h\alpha(x)} \Delta_{xx} V(x, t; h, \lambda). \tag{5.25}$$

Next, taking h very small with $\lambda/h^2 = \mu$, we have

$$\mu \Delta_t V(x, t; h, \lambda) = \frac{\alpha(x)}{\beta(x)} \Delta_x V(x, t; h, \lambda) + \frac{1}{2} \Delta_{xx} V(x, t; h, \lambda). \tag{5.26}$$

So the simulation, which proceeds directly from the microaxioms, yields in the limit as the infinitesimals go to zero a practical pointwise evaluator of the usual Fokker–Planck equation:

$$2\mu \frac{\partial V}{\partial t} = 2 \frac{\alpha(x)}{\beta(x)} \frac{\partial V}{\partial x} + \frac{\partial^2 V}{\partial x^2}. \tag{5.27}$$

The Fokker–Planck equation is generally not solvable in closed form. Note that we have given a simulation-based approach for solving (5.27), but more importantly we have given a practical means for arriving at the consequences of the original axioms which brought about Fokker–Planck in the first place. So we again raise the intriguing possibility that had computers been available 100 years ago, Fokker and Planck might have simply represented their model in microaxiomatic format instead of giving a differential equation summary thereof. Again, our algorithm is essentially the flowchart in Figure 5.2 with a time counter added on.

5.5 The Dirichlet Problem

Next, we consider another common differential equation model of physics, that of Dirichlet. In \mathcal{R}_k , let there be given a bounded connected region S Γ (Figure 5.3). Let there be given a function $F(x)$ satisfying the equation

of Laplace inside S :

$$\sum_{j=1}^k \frac{\partial^2 \phi_j}{\partial x_j^2} = 0. \quad (5.28)$$

The values of ϕ are given explicitly at every boundary point by the piecewise continuous function $f(Q)$; that is,

$$\phi(x)|_{\Gamma} = f(Q). \quad (5.29)$$

For most boundaries and boundary functions, the determination of ϕ analytically is not known. The usual numerical approximation approach can require a fair amount of setup work, particularly if the dimensionality is 3 or greater. We exhibit below a simulation technique which is, in fact, an actualization of the microaxioms which frequently give rise to (5.28). Although our discussion is limited to \mathcal{R}_2 , the generalization to \mathcal{R}_k is quite straightforward. Let us superimpose over S a square grid of length h on a side. The points of intersection inside S nearest Γ will be referred to as *boundary nodes*. All other nodes inside S shall be referred to as *internal nodes*.

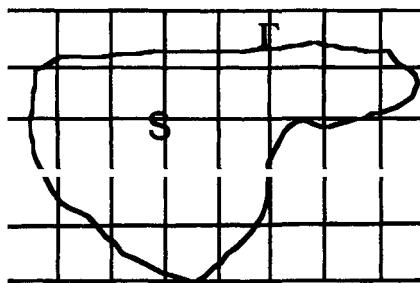


Figure 5.3. The Dirichlet problem.

In Figure 5.4, we consider an internal node with coordinates (x, y) in relation to its four immediate neighbors.

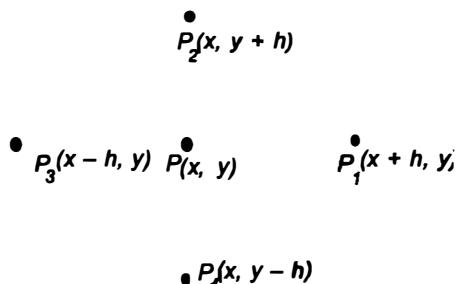


Figure 5.4. Random walk grid for Dirichlet problem.

Now

$$\partial\phi/\partial x|_{x,y} \approx \frac{\phi(x+h/2, y) - \phi(x-h/2, y)}{h} \quad (5.30)$$

and

$$\begin{aligned} \frac{\partial^2\phi}{\partial x^2} &\approx \frac{(\partial\phi/\partial x)|_{x+h/2,y} - (\partial\phi/\partial x)|_{x-h/2,y}}{h} \\ &\approx \frac{\phi(x+h, y) + \phi(x-h, y) - 2\phi(x, y)}{h^2}. \end{aligned} \quad (5.31)$$

Similarly,

$$\frac{\partial^2\phi}{\partial y^2} \approx \frac{\phi(x, y+h) + \phi(x, y-h) - 2\phi(x, y)}{h^2}. \quad (5.32)$$

Equation (5.28) then gives

$$0 = \frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2} \approx \frac{\phi(P_1) + \phi(P_2) + \phi(P_3) + \phi(P_4) - 4\phi(P)}{h^2}. \quad (5.33)$$

So

$$\phi(P) \approx \frac{\phi(P_1) + \phi(P_2) + \phi(P_3) + \phi(P_4)}{4}. \quad (5.34)$$

Equation (5.34) gives us a ready means of a simulation solution to the Dirichlet problem. Starting at the internal node (x, y) , we randomly walk to one of the four adjacent points with equal probabilities. We continue the process until we reach a boundary node, say Q_i . After N walks to the boundary from starting point (x, y) , our estimate of $\phi(x, y)$ is given simply by

$$\phi(x, y) \approx \frac{\sum_{i=1}^N n_i f(Q_i)}{N}, \quad (5.35)$$

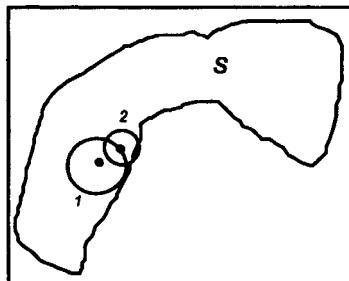


Figure 5.5. Quick steps to the boundary.

where n_i is the number of walks terminating at boundary node Q_i and the summation is taken over all boundary nodes.

In Figure 5.3, if we wish to show ϕ contours throughout S , we can take advantage of computational efficiencies. For example, as we walk from

(x, y) to the boundary, we will traverse $(x + h, y)$ numerous times. By incorporating the walks which traverse $(x + h, y)$ even though $(x + h, y)$ is not our starting point, we can increase the total number of walks used in the evaluation of $\phi(x + h, y)$.

Let us now consider in Figure 5.5 a technique that is particularly useful if we need to evaluate ϕ at only one point in S . Since it can easily be shown that $\phi(x, y)$, the solution to Laplace's equation inside S , is equal to the average of all values taken on a circle centered at (x, y) and lying inside S , we can draw around (x, y) the largest circle lying inside S and then select a point uniformly on the boundary of that circle, use it as the center of a new circle, and select a point at random on that circle. We continue the process until we arrive at a boundary node. Then, again after N walks,

$$\phi(x, y) \approx \frac{\sum_{i=1}^N n_i f(Q_i)}{N}. \quad (5.36)$$

This method works, using hyperspheres, for any dimension k . Again, it should be emphasized that in many cases the simulation is a direct implementation of the microaxioms that gave rise to Laplace's equation.

5.6 Solution of General Elliptic Differential Equations

Next, we consider a general elliptic differential equation in a region S in two-dimensional space. Again, the values on the boundary Γ are given by $f(\cdot)$, which is piecewise continuous on Γ . Inside S ,

$$\beta_{11} \frac{\partial^2 \phi}{\partial x^2} + 2\beta_{12} \frac{\partial^2 \phi}{\partial x \partial y} + \beta_{22} \frac{\partial^2 \phi}{\partial y^2} + 2\alpha_1 \frac{\partial \phi}{\partial x} + 2\alpha_2 \frac{\partial \phi}{\partial y} = 0, \quad (5.37)$$

where $\beta_{11} > 0$, $\beta_{22} > 0$, and $\beta_{11}\beta_{22} - \beta_{12}^2 > 0$. We consider the difference equation corresponding to (5.37), namely:

$$\beta_{11} \Delta_{xx} \phi + 2\beta_{12} \Delta_{xy} \phi + \beta_{22} \Delta_{yy} \phi + 2\alpha_1 \Delta_x \phi + 2\alpha_2 \Delta_y \phi = 0. \quad (5.38)$$

As convenient approximations to the finite differences, we use

$$\begin{aligned} \Delta_{xy} &= \frac{\phi(x + h, y + h) - \phi(x, y + h) - \phi(x + h, y) + \phi(x, y)}{h^2} \\ \Delta_{xx} &= \frac{\phi(x + h, y) + \phi(x - h, y) - 2\phi(x, y)}{h^2} \\ \Delta_{yy} &= \frac{\phi(x, y + h) + \phi(x, y - h) - 2\phi(x, y)}{h^2} \\ \Delta_x &= \frac{\phi(x + h, y) - \phi(x, y)}{h} \\ \Delta_y &= \frac{\phi(x, y + h) - \phi(x, y)}{h}. \end{aligned} \quad (5.39)$$

These differences involve five points around (x, y) . Now, we shall develop a random walk realization of (5.37), which we write out explicitly in Figure 5.6,

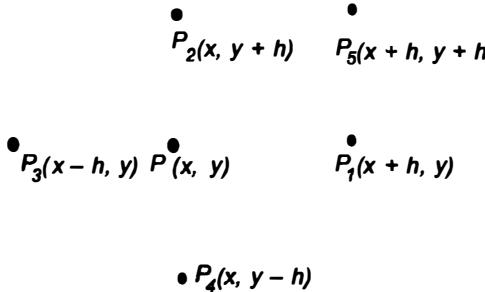


Figure 5.6. Grid for elliptic equation random walk.

$$\begin{aligned}
 & \beta_{11} \frac{\phi(x+h, y) + \phi(x-h, y) - 2\phi(x, y)}{h^2} \\
 & + 2\beta_{12} \frac{\phi(x+h, y+h) - \phi(x, y+h) - \phi(x+h, y) + \phi(x, y)}{h^2} \\
 & + \beta_{22} \frac{\phi(x, y+h) + \phi(x, y-h) - 2\phi(x, y)}{h^2} \\
 & + 2\alpha_1 \frac{\phi(x+h, y) - \phi(x, y)}{h} + 2\alpha_2 \frac{\phi(x, y+h) - \phi(x, y)}{h} = 0. \quad (5.40)
 \end{aligned}$$

We rearrange the terms in (5.40) to give

$$\begin{aligned}
 & \phi(x+h, y)(\beta_{11} + 2\alpha_1 h - 2\beta_{12}) + \phi(x, y+h)(\beta_{22} + 2\alpha_2 h - 2\beta_{12}) \\
 & + \phi(x-h, y)\beta_{11} + \phi(x, y-h)\beta_{22} + \phi(x+h, y+h)2\beta_{12} \\
 & = \phi(x, y)[2\beta_{11} - 2\beta_{12} + 2\beta_{22} + 2(\alpha_1 + \alpha_2 h)]. \quad (5.41)
 \end{aligned}$$

Letting $D = [2\beta_{11} - 2\beta_{12} + 2\beta_{22} + 2(\alpha_1 + \alpha_2 h)]$, we have

$$\begin{aligned}
 & \phi(x+h, y)p_1 + \phi(x, y+h)p_2 + \phi(x-h, y)p_3 + \phi(x, y-h)p_4 \\
 & + \phi(x+h, y+h)p_5 = \phi(x, y), \quad (5.42)
 \end{aligned}$$

with

$$\begin{aligned}
 p_1 &= \frac{\beta_{11} + 2\alpha_1 h - 2\beta_{12}}{D}; & p_2 &= \frac{\beta_{22} + 2\alpha_2 h - 2\beta_{12}}{D}; \\
 p_3 &= \frac{\beta_{11}}{D} i s p_4 = \frac{1\beta_{22}}{D}; & p_5 &= \frac{2\beta_{12}}{D}. \quad (5.43)
 \end{aligned}$$

Note that in the formulation above, we must exercise some care to assure that the probabilities are non-negative. By using the indicated probabilities, we walk randomly to the boundary repeatedly and use the estimate

$$\phi(x, y) = \frac{\sum_{i=1}^N n_i f(Q_i)}{N} . \quad (5.44)$$

5.7 Conclusions

The examples provided in this chapter are given to give the reader a feel as to the practical implementation of simulation-based algorithms as alternatives to the usual numerical approximation techniques. A certain amount of practice quickly brings the user to a point where he or she can write simulation algorithms in days to problems that would require the numerical analyst months to approach.

Other advantages of the simulation approach could be given. For example, since walks to boundaries are so simple to execute, it is easy to conceptualize the utilization of parallel processors to speed up the computations with a minimum of handshaking between the CPUs. But the main advantage of simulation is its ability to enable the user to bypass the traditional path in Figure 5.1 and go directly from the microaxioms to their macro consequences. Our algorithm for solving the Fokker–Planck problem and our algorithm for solving the Dirichlet problem are not simply analogues of the classical differential-equation formulations of these systems. *They are, in fact, descriptions of the axioms that typically give rise to these problems.* Here we note that the classical differential-equation formulation of many of the systems of physics and chemistry proceed from the axioms that form the simulation algorithm. It was simply the case, in a precomputer age, that the differential-integral equation formulations appeared to give the best hope of approximate evaluation via series expansions and the like.

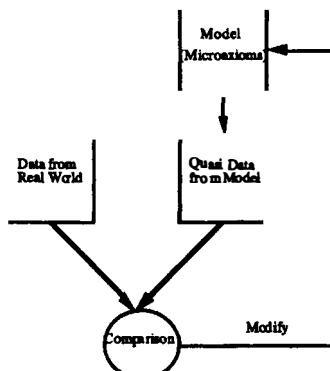


Figure 5.7. The idealized simulation paradigm.

Moving further along, we now have the possibility of implementing the algorithm in Figure 5.7 as a basic model-building paradigm. In the path shown, computer simulation may obviate the necessity of modeling via, say, differential equations and then using the computer as a means of approximating solutions, either by numerical techniques or by Monte Carlo to the axiomitized model. The computer then ceases to be a fast calculator, and shifts into actually simulating the process under consideration. Such a path represents a real paradigm shift and will begin the realization of the computer as a key part of the scientific method.

Problems

5.1. Consider the differential equation

$$\frac{d^2v}{dx^2} + 2x^{2.5} \frac{dv}{dx} = 0,$$

where $v(0) = 1$ and $v(1) = 2$. Use the flowchart in Figure 5.2 to obtain estimates of $v(.2)$, $v(.4)$, $v(.6)$, and $v(.8)$, together with error bounds for these quantities.

5.2. Program a quadrature-type differential equation solver for the following differential equation, where $x(0) = 0$ and $x(1) = 1$:

$$\frac{d^2x}{dt^2} + \beta(t) \frac{dx}{dt} = 0.$$

Compare its performance with a simulation-based approach using Figure 5.2 for the following candidates for β :

- (a) $\beta = 4$
- (b) $\beta = t$
- (c) $\beta = \sin(4\pi t/t + 1)$.

5.3. Consider the differential equation on the unit x interval.

$$\frac{\partial V}{\partial t} = 2 \frac{\partial V}{\partial x} + \frac{\partial^2 V}{\partial x^2},$$

where $V(x, t) = 200x \exp(-t)$. Again, use the flowchart in Figure 5.2 to obtain estimates of $V(.2, t)$, $V(.4, t)$, $V(.6, t)$, and $V(.8, t)$, for $t = 0, 1, 2, 3$.

5.4. Draw contours for T in the interior of the plate shown in Figure 5.8,

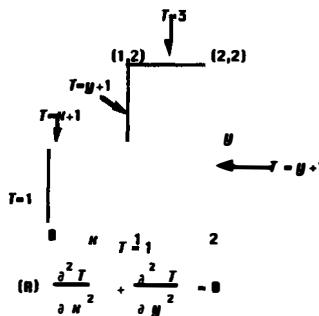


Figure 5.8. Dirichlet problem for a plate.

using increments of 0.2 in both x and y based on the Gambler's Ruin flowchart in Figure 5.2, given that Laplace's equation (A) is satisfied in the interior with the boundary conditions shown. Here, we have a substantial demonstration of how inefficient it would be to run simulation from each interior point (x, y) independently of walks from other points which "step on" (x, y) .

5.5. Consider the equation

$$(2+y)\frac{\partial^2 \phi}{\partial x^2} + x\frac{\partial^2 \phi}{\partial x \partial y} + (2-y)\frac{\partial^2 \phi}{\partial y^2} - \frac{\partial \phi}{\partial x} = 0,$$

which is satisfied inside the unit circle shown in Figure 5.9, with the boundary condition indicated. Use the standard gambler's ruin flowchart in Figure 5.2 to obtain estimates of the ϕ contours inside the unit circle.

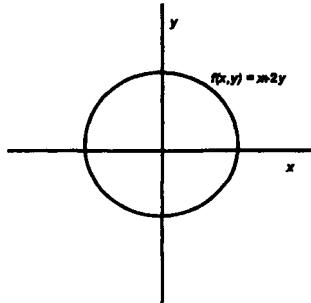


Figure 5.9. Elliptic equation on the circle.

References

- [1] Lattes, R. (1969). *Methods of Resolution for Selected Boundary Problems in Mathematical Physics*. New York: Gordon and Breach.
- [2] Shreider, Y.A. (1966). *The Monte Carlo Method: The Method of Statistical Trials*. New York: Pergamon Press.
- [3] Thompson, J.R. (1989). *Empirical Model Building*. New York: John Wiley & Sons, 95–108.

Chapter 6

SIMDAT, SIMEST, and Pseudoreality

6.1 Introduction

Many of us have had the experience of wishing we had 10 times the data at hand. Many of us have had the experience of trying to estimate the parameters of a model in a situation where we found it mathematically infeasible to write down a likelihood function to maximize. Many of us have needed to look at a higher-dimensional data set by trying to use our rather limited three (or four)-dimensional perceptions of reality.

I recall some years ago being on the doctoral committee of an electrical engineering student who had some time-indexed data, where the sampling intervals were so wide that he was unable to detect features at the frequencies where his interest lay. His solution (of which he was inordinately proud) was to create a spline curve-fitting algorithm which would magically transform his discrete data into continuous data. By one fell swoop he had rendered Nyquist irrelevant. Although second readers should generally keep their silence, I had to point out that he had assumed away high-frequency components by his approach.

Having hosted at Rice in the late 1970s one of the early short courses on exploratory data analysis, I recall one of the two very distinguished instructors in the course making the statement that “EDA (exploratory data analysis) frees us from the straightjacket of models, allowing the data to speak to us unfettered by preconceived notions.” During the balance of the course, whenever a new data set was displayed from one of the dual projectors, a rude psychometrician in the audience would ostentatiously cup his hand to his left ear as though to hear the data better. The psy-

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

chometrician knew full well that data are perceived via a model, implicit or explicit. The strength of EDA is that it has provided us with the most effective analog-digital computer interface currently available, with the human visual system providing the analog portion of the system. Certainly, Tukey's exploratory data analysis [27] is one of the most important data analytical tools of the the last 50 years. But it has tended to tie us to a three-dimensional perception, which is unlikely to be appropriate when dealing with data of high dimensionality.

The same criticism may be made of those of us who have tried to push graphical displays of nonparametric density estimates into high-dimensional situations in the hope that spinning, coloring, and so on, might somehow force the data into a realm where our visual systems would suffice to extract the essence of the system generating the data. For many years now, I have been artistically impressed by those who would take data sets, color them, spin them, project them, and/or time lapse them. But, in retrospect, it is hard to think of many examples where these fun type activities contributed very much to an understanding of the basic mechanism which formed the data set in the first place. Unfortunately, it seems as though many computer intensive studies, in the case of EDA, nonparametric density estimation, nonparametric regression, and so on, in the hands of many users, have more or less degenerated into essentially formalist activities, that is, activities in which we are encouraged not so much to appreciate what data analytical insights the algorithms contribute, but rather to appreciate the algorithms *sui generis*, as intrinsically wonderful.

In the matter of both density estimation and nonparametric regression, the bulk of the computer intensive work continues to emphasize the one-dimensional situation, that in which simple-minded methods (e.g., histograms, hand-fit curves, etc.) have been used for a long time with about as much success as the newer techniques. Is it not interesting to observe that the histogram is still much the most used of the one-dimensional nonparametric density estimators, and that one-dimensional curve fits are, at their most effective, psychometric activities, where one tries to automate what a human observer might do freehand? In the case where several independent variables are contemplated in nonparametric regression, rather unrealistic assumptions about additivity tend to be implemented in order to make the formal mathematics come out tractably.

In the case of one-dimensional nonparametric density estimation, a fair argument can be made that Rosenblatt obtained much of the practical knowledge we have today in his 1956 paper [18]. In the case of multivariate nonparametric density estimation, we have barely scratched the surface. Those who favor pushing three-dimensional graphical arguments into higher dimensions have made definite progress (see, e.g., Scott and Thompson [19] and Scott [20]). But others take the approach that the way to proceed is profoundly nonvisual. Boswell [4], Elliott and Thompson [10], as well as Thompson and Tapia [25] take the position that it is reasonable

to start by seeking for centers of high density using algebraic, nongraphical, algorithms. The case where the density is heavily concentrated on curved manifolds in high-dimensional space is a natural marriage of nonparametric regression and nonparametric density estimation and has not yet received the attention one might have hoped.

The case for nonparametric regression has been eloquently and extensively advocated by Hastie and Tibshirani [15]. Other important contributions include those of Cleveland [6], Cox [7], Eubank [11], Härdle [13], and Hart and Wehrly [14]. However, I fear that the nonparametric regressor is swimming against a sea of intrinsic troubles. First of all, extrapolation continues to be both the major thing we need for application, and something very difficult to achieve, absent an investigation of the underlying model which generates the data. For interpolation purposes, it would appear that we really do not have anything more promising than locally averaged smoothers. The important case here is, of course, the one with a reasonable number of independent variables. Unfortunately, much of the work continues to be for the one-independent-variable case, where the job could be done "with a rusty nail."

In the cases mentioned, EDA, nonparametric regression, nonparametric density estimation, algorithms have been and are being developed which seek to affect data analysis positively by use of the high-speed digital computer. I have to opine that it seems most of this "machine in search of a problem" approach has not yet been particularly successful.

6.2 The Bootstrap: A Dirac Comb Density Estimator

There are statistical areas in which the use of the computer has borne considerable fruit. In Chapter 4, we gave a nontheoretical example laden treatment of the bootstrap of Efron [8], clearly one of the most influential algorithms of the last 30 years or so. To motivate the bootstrap, we follow the discussions in Taylor and Thompson [21], [22], and Thompson and Tapia [25]. Let us first consider the Dirac comb density estimator associated with a one-dimensional data set $\{x_i\}_{i=1}^n$. The Dirac comb density estimator is given by

$$\hat{f}_\delta(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - x_i). \quad (6.1)$$

We may represent $\delta(x)$ as

$$\delta(x) = \lim_{\tau \rightarrow 0} \frac{1}{\sqrt{2\pi}\tau} e^{-x^2/2\tau^2}. \quad (6.2)$$

Waving our hands briskly, $\delta(x)$ can be viewed as a density function that is zero everywhere except at the data points. At each of these points,

the density has mass $1/n$. Nonparametric density estimation is frequently regarded as a subset of smoothing techniques in statistics. $\hat{f}_\delta(x)$ would seem to be infinitely rough and decidedly nonsmooth. Moreover, philosophically, it is strongly nominalist, for it says that all that can happen in any other experiment is a repeat of the data points already observed, each occurring with probability $1/n$. In other words, the data are considered to be all of reality.

For many purposes, however, $\hat{f}_\delta(x)$ is useful. For example, the mean of $\hat{f}_\delta(x)$ is simply

$$\mu_\delta = \int_{-\infty}^{\infty} x \hat{f}_\delta(x) = \frac{1}{n} \sum_{i=1}^n x_i = \bar{x}. \quad (6.3)$$

The sample variance can be represented by

$$\sigma_\delta^2 = \int_{-\infty}^{\infty} (x - \bar{x})^2 \hat{f}_\delta(x) dx = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 = s^2. \quad (6.4)$$

So, if we are interested in discrete characterizations of the data (such as a few lower order moments), a Dirac comb may work quite satisfactorily. The Dirac comb density estimator may be easily extended to higher dimensions. Indeed, such an extension is easier than is the case for other estimators, for it is based on masses at points; and a point in, say, 20-space, is still a point (hence, of zero dimension). Thus, if we have a sample of size n from a density of dimension p , $\hat{f}_\delta(x)$ becomes

$$\hat{f}_\delta(X) = \frac{1}{n} \sum_{i=1}^n \delta(X - X_i), \quad (6.5)$$

where

$$\delta(X) = \lim_{\tau \rightarrow 0} \left(\frac{1}{\sqrt{2\pi}\tau} \right)^p \exp \left(-\frac{\sum_{j=1}^p x_j^2}{2\tau^2} \right), \quad (6.6)$$

with x_j being the j component of X .

For the two-dimensional case, we might wish to develop a 95% confidence interval for the correlation coefficient,

$$\rho = \frac{\text{Cov}(x, y)}{\sigma_x \sigma_y}. \quad (6.7)$$

Now, if we had a sample of size n : $\{x_i, y_i\}_{i=1}^n$, we could construct

$$\hat{f}_\delta(x, y) = \frac{1}{n} \sum_{i=1}^n \delta((x, y) - (x_i, y_i)). \quad (6.8)$$

Next, we construct 10,000 resamplings (with replacement) of size n . That means, for each of the 10,000 resamplings we draw samples from the n data points (with replacement) of size n . For each of the resamplings we compute the sample correlation:

$$r_j = \frac{\sum_{i=1}^n (x_{ji} - \bar{x}_j)(y_{ji} - \bar{y}_j)}{\sqrt{\sum_{i=1}^n (x_{ji} - \bar{x}_j)^2 \sum_{i=1}^n (y_{ji} - \bar{y}_j)^2}}. \quad (6.9)$$

Next, we can rank the sample correlations from smallest to largest. A 95% confidence interval estimate is given by

$$r_{(250)} < \rho < r_{(9750)}. \quad (6.10)$$

One can view the bootstrap as being based on such a Dirac-comb estimator. Although it is clear that such a procedure may have use for estimating the lower moments of some interesting parameters, we should never lose sight of the fact that it is, after all, based on the profoundly discontinuous Dirac-comb estimator \hat{f}_δ . The smoothed bootstrap [9] operates very much like the bootstrap itself, except that to each resampled point one adds, say, a normal variate with small variance. Essentially one samples from a fuzzy Dirac-comb nonparametric density estimator.

Now, for better or for worse, the reality is that much of statistics is concerned with such tasks as estimating a few moments. When we know the underlying density function (and history shows that people get away with assuming that the world is Gaussian more often than might be supposed), then knowledge of a few moments actually gives a continuous description of the underlying system [the first and second moments of a Gaussian (normal) distribution completely characterize the density function everywhere].

However, if the world truly were Gaussian, then we could drop the entire subject of nonparametrics (and most computer-intensive statistical analyses). Let us consider an example where the data really are Gaussian, but the use of the Dirac-comb nonparametric density estimator serves us poorly. For example, suppose we have a sample of size 100 of firings at a bull's-eye of radius 5 centimeters. If the distribution of the shots is circular normal with mean the center of the bull's-eye and deviation 1 meter, then with a probability in excess of .88, none of the shots will hit inside the bull's-eye. Then any Dirac-comb resampling procedure will tell us the bull's-eye is a safe place if we get a base sample (as we probably will) with no shots in the bullseye. Such a problem with the bootstrap motivated SIMDAT, the nonparametric density estimator based resampling algorithm of Taylor and Thompson [21, 22, 25].

6.3 SIMDAT: A Smooth Resampling Algorithm

We note that any realization of a bootstrap simulation most likely will be different from the original sample. Some sample points will disappear. Others will be repeated multiple times. Indeed, the concatenation of a bootstrap followed by a bootstrap based on that bootstrapped simulation, and so on, will lead ultimately to a simulated sample which consists of a single sample point. This is hardly desirable. It might be hoped that a single resampling would be of such a character that we would be almost indifferent as to whether we had this simulation or the original data set. But, of course, it would be dangerous to wander too far from the original sample. A resampling of a resampling of a resampling, and so on, is not nearly as desirable as resamples that always point directly to the original sample.

The bootstrap is clearly a powerful algorithm for many purposes. However, given the ubiquity of fast computing, it would usually be preferred to use resampling schemes based on better nonparametric density estimators than the Dirac comb. One such would be the 1976 algorithm of Guerra, Tapia and Thompson [12], where one obtains a smooth of the empirical cdf and samples from that. This algorithm has been employed for some time as the RNGCT subroutine of Visual Numerics (formerly IMSL). The disadvantage of the algorithm is that it was only written for the one-dimensional case, and that the estimator of the cdf must be explicitly obtained.

One candidate for a nonparametric density estimator to be used for simulation purposes would be

$$\hat{f}_\delta(X) = \frac{1}{n} \sum_{i=1}^n K(X - X_i, \Sigma_i), \quad (6.11)$$

where $K(X, \Sigma_i)$ is a normal distribution centered at zero with locally estimated covariance matrix Σ_i .

Such an estimator, despite its advantages, would seem to be very difficult to construct. However, let us recall what it is we seek: not a nonparametric density estimator, but a random sample from such an estimator. So, perhaps, we can go directly from the actual sample to the pseudosample. Of course, this is precisely what the bootstrap estimator does, with the frequently unfortunate properties associated with a Dirac comb. Fortunately, it is possible to go from the sample directly to the pseudosample in such a way that the resulting estimator behaves very much like that of the normal kernel approach above. This is what the SIMDAT algorithm does.

6.3.1 The SIMDAT Algorithm

Assume given a data set of size n from a p -dimensional variable X , $\{X_i\}_{i=1}^n$. Assume that we have already rescaled our data set so that the marginal sample variances in each vector component are the same. For a given integer m , we can find, for each of the n data points, the $m - 1$ nearest neighbors. These will be stored in an array of size $n \times (m - 1)$.

Suppose we wish to generate a pseudosample of size N . Note that there is no reason to suppose that n and N need be the same (as is the case generally with the bootstrap). To start the algorithm, we sample one of the n data points with probability $1/n$ (just as with the bootstrap). Then, we recall its $m - 1$ nearest neighbors from memory, and compute the mean of the resulting set of m points:

$$\bar{X} = \frac{1}{m} \sum_{i=1}^m X_i. \quad (6.12)$$

Next, we subtract from each of the data points the local mean \bar{X} , thus achieving zero averages of the transformed cloud:

$$\{X'_j\} = \{X_j - \bar{X}\}_{j=1}^m. \quad (6.13)$$

Although we go through the computations of sample means and centering about them here as though they were a part of the simulation process, the operation will be done once only, just as with the determination of the $m - 1$ nearest neighbors of each data point. The $\{X'_j\}$ values as well as the \bar{X} values will be stored in an array of dimension $n \times (m + 1)$.

Next, we generate a random sample of size m from the one-dimensional uniform distribution:

$$U \left(\frac{1}{m} - \sqrt{\frac{3(m-1)}{m^2}}, \frac{1}{m} + \sqrt{\frac{3(m-1)}{m^2}} \right). \quad (6.14)$$

We now generate our centered pseudodata point X' , via

$$X' = \sum_{l=1}^m u_l X'_l. \quad (6.15)$$

Finally, we add back on \bar{X} to obtain our pseudodata point X :

$$X = X' + \bar{X}. \quad (6.16)$$

These, then, are the nuts and bolts of SIMDAT. The major setup cost is the determination of interpoint distances. The tabulation is for each of the n data points, a list indicating the $m - 1$ nearest points. Once the resulting matrix has been obtained, subsequent generation of any desired amount of pseudodata is very rapid.

6.3.2 An Empirical Justification of SIMDAT

As m and n get large, the procedure gives results very much like those of the normal kernel approach mentioned earlier. To see why this is so, we consider the sampled vector X_l and its $m - 1$ nearest neighbors:

$$\{X_l\}_{l=1}^m = \begin{bmatrix} x_{1l} \\ x_{2l} \\ \vdots \\ x_{kl} \end{bmatrix}_{l=1,\dots,m}. \quad (6.17)$$

Let us treat this collection of m points as being from a distribution with mean vector μ and covariance matrix Σ . Now, if $\{u_l\}_{l=1}^m$ is an independent sample from the uniform distribution in (6.14), then

$$E(u_l) = \frac{1}{m}; \quad \text{Var}(u_l) = \frac{m-1}{m^2}; \quad \text{Cov}(u_i, u_j) = 0, \text{ for } i \neq j. \quad (6.18)$$

Then we form the linear combination

$$Z = \sum_{l=1}^m u_l X_l. \quad (6.19)$$

We note that for the r th component of the vector Z , $z_r = u_1 x_{r1} + u_2 x_{r2} + \dots + u_m x_{rm}$,

$$E(z_r) = \mu_r, \quad (6.20)$$

$$\text{Var}(z_r) = \sigma_r^2 + \frac{m-1}{m} \mu_r^2, \quad (6.21)$$

and

$$\text{Cov}(z_r, z_s) = \sigma_{rs} + \frac{m-1}{m} \mu_r \mu_s. \quad (6.22)$$

We observe that if the mean vector of X were $(0, 0, \dots, 0)$, then the mean vector and covariance matrix of Z would be the same as that of X , that is, $E(z_r) = 0$, $\text{Var}(z_r) = \sigma_r^2$, and $\text{Cov}(z_r, z_s) = \sigma_{rs}$. Naturally, by translation to the local sample mean of the nearest-neighbor cloud, we will not quite have achieved this result. But we will come very close to the generation of an observation from the truncated distribution that generated the points in the nearest-neighbor cloud.

For m moderately large, by the central limit theorem, SIMDAT comes close to sampling from n normal distributions with the mean and covariance matrices corresponding to those of the n, m nearest-neighbor clouds.

If we were seeking rules for consistency of the nonparametric density estimator corresponding to SIMDAT, we could use the formula of Mack and Rosenblatt [16] for nearest-neighbor nonparametric density estimators:

$$m = Cn^{4/(p+4)}. \quad (6.23)$$

Actually, as a practical matter, such formulas have little practical relevance, since C is usually not available. Furthermore, we ought to remember that our goal is not to obtain a nonparametric density estimator, but rather, to generate a data set which appears like that of the data set before us. Let us suppose that we err on the side of making m far too small, namely, $m = 1$. That would yield simply the bootstrap. Suppose that we err on the side of making m far too large, namely, $m = n$. That would yield an estimator which roughly sampled from a multivariate normal distribution with the mean vector and covariance matrix computed from the data. In Figure 6.1, we show a sample of size 85 from a mixture of three normal distributions with the weights indicated, and a pseudodata set of size 85 generated by SIMDAT with $m = 5$. We note that the emulation of the data is reasonably good. In Figure 6.2 we go through the same exercise, but with $m = 15$. There, effects of a modest oversmoothing are noted. In general, if the data set is very large, say of size 1000 or greater, good results are obtained with $m \approx .02n$. For smaller values of n , m values in the $.05n$ range appear to work well. A version of SIMDAT in the S language, written by E.N. Atkinson, is available under the name "gendat" from <http://lib.stat.cmu.edu/S/gendat>.

$$\mu_1 = \begin{bmatrix} -1 \\ -2 \end{bmatrix}; \mu_2 = \begin{bmatrix} -2 \\ 3 \end{bmatrix}; \mu_3 = \begin{bmatrix} 2 \\ \frac{3}{2} \end{bmatrix}, \quad \Sigma_1 = \begin{bmatrix} 1 & \frac{1}{2} \\ -\frac{1}{2} & 1 \end{bmatrix}; \Sigma_2 = \begin{bmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & 1 \end{bmatrix}; \Sigma_3 = \begin{bmatrix} 1 & \frac{1}{10} \\ \frac{1}{10} & 1 \end{bmatrix}.$$

$$\text{Mixed Distribution} = \frac{1}{2} N_1 + \frac{1}{3} N_2 + \frac{1}{6} N_3.$$

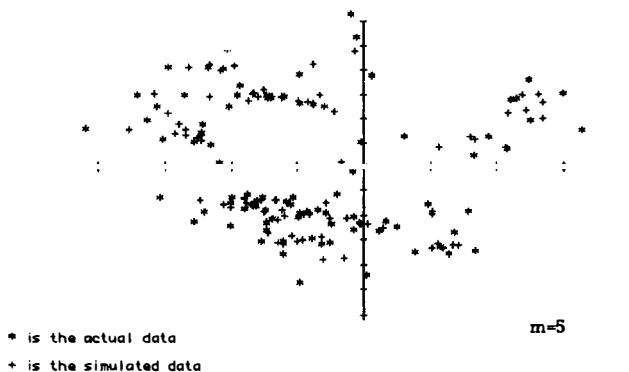


Figure 6.1. Undersmoothed SIMDAT.

$$\mu_1 = \begin{bmatrix} -1 \\ -2 \end{bmatrix}, \mu_2 = \begin{bmatrix} -2 \\ 3 \end{bmatrix}, \mu_3 = \begin{bmatrix} 2 \\ \frac{3}{2} \end{bmatrix}, \quad \Sigma_1 = \begin{bmatrix} 1 & \frac{1}{2} \\ -\frac{1}{2} & 1 \end{bmatrix}, \quad \Sigma_2 = \begin{bmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & 1 \end{bmatrix}, \quad \Sigma_3 = \begin{bmatrix} 1 & \frac{1}{10} \\ \frac{1}{10} & 1 \end{bmatrix}.$$

$$\text{Mixed Distribution} = \frac{1}{2} N_1 + \frac{1}{3} N_2 + \frac{1}{6} N_3.$$

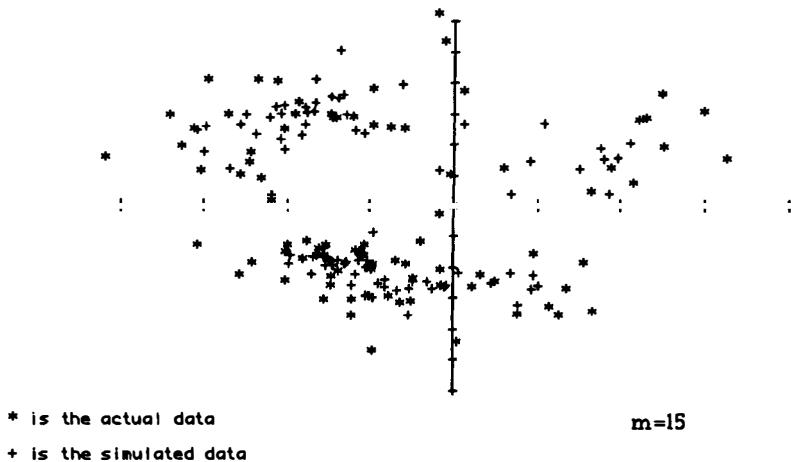


Figure 6.2. Oversmoothed SIMDAT.

For really large data sets, the user may wish to use Fortran or C instead of S. The savings for using the more primitive languages, as opposed to S or R, may be 100-fold.

So far, we have considered basically model-free techniques for examining data. There are, of course, many situations where exploration of a new data set may preclude an early conjecture as to a likely model of the mechanism generating the data. In my opinion, such procedures should usually be first steps in modeling a process. But, unfortunately, they frequently are as far as one goes. In essence, the nonparametric techniques use the power of the computer to bypass altogether the need for the modeling step. Such an approach is likely to be useful mainly as an interpolative device. When the dimensionality of a data set becomes high, say five or more, this *ad hocery* is likely to prove dangerous, since we may be confronted with a number of widely separated modes, with deserts in between. Dealing with such data sets, nonparametrically, away from the modes, is an extrapolation problem, and using the standard smoothed interpolation routines can bring one quickly to disaster.

6.4 SIMEST: An Oncological Example

The power of the computer as an aid to modeling does not get the attention it deserves. Part of the reason is that the human modeling approach tends to be analog rather than digital. Analog computers were replaced by digital computers 40 years ago. Most statisticians remain fascinated by the graphical capabilities of the digital computer. The exploratory data analysis route tends to attempt to replace modeling by visual displays which are then interpreted, in a more-or-less instinctive fashion, by an observer. Statisticians who proceed in this way are functioning somewhat like prototypical cyborgs. After over four decades of seeing data spun, colored, and graphed in a myriad of ways, I have to admit to being disappointed when comparing the promise of EDA with its reality. Its influence amongst academic statisticians has been enormous. Visualization is clearly one of the major areas in the statistical literature. But the inferences drawn from these visualizations in the real world are, relatively speaking, not so numerous. Moreover, when visualization-based inferences are drawn, they tend to give results one might have obtained by classical techniques.

Of course, as in the case of using the computer as a nonparametric smoother, some uses are better than others. In the 1980s a group of Bayesian statisticians convinced one of our leading research universities that the reason statistical analysis had produced marginal results in such areas as oncology had been the traditional dominance of frequentists in biometry. The advent of high-speed computing brought forth the possibility that the insights of physicians could be appropriately blended into priors leading to breakthrough posteriors. Here, we were told that the computer would enable us to carry out another one, two, or three dimensions of quadrature, thus enabling prior information to be infused into the process. But, since the desired prior information really was not available (and may never be in the form required), the computer just enabled people to spin their wheels faster.

It is extremely unfortunate that some are so multicultural in their outlook that they rearrange their research agenda in order to accommodate themselves to our analog-challenged friends, the digital computers. Perhaps the greatest disappointment is to see the modeling aspect of our analog friends, the human beings, being disregarded in favor of using them as gestaltic image processors. This really will not do. We need to rearrange the agenda so that the human beings can gain the maximal assistance from the computers in making inferences from data. That is the purpose of SIMEST.

There is an old adage to the effect that quantitative change carried far enough may produce qualitative change. The fact is that we now have computers so fast and cheap that we can proceed (almost) as though computation were free and instantaneous (with infinite accessible memory thrown in as well). This should change, fundamentally, the way we approach data analysis in the light of models.

There are now numerous examples in several fields where SIMEST has been used to obtain estimates of the parameters characterizing a market-related applied stochastic process (see, e.g., Bridges, Ensor, and Thompson [5]). Below we consider an oncological application to motivate and to explicate SIMEST. We shall first show a traditional model-based data analysis, note the serious (generally insurmountable) difficulties involved, and then give a simulation-based, highly computer-intensive way to get what we require to understand the process and act on that understanding.

6.4.1 An Exploratory Prelude

In the late 1970s, my colleague Barry W. Brown, of the University of Texas M.D. Anderson Cancer Center, and I had started to investigate some conjectures concerning reasons for the relatively poor performance of oncology in the American "War on Cancer." Huge amounts of resources had been spent with less encouraging results than one might have hoped. It was my view that part of the reason might be that the basic orthodoxy for cancer progression was, somehow, flawed.

This basic orthodoxy can be summarized briefly as follows:

At some time, for some reason, a single cell goes wild. It, and its progeny, multiply at rates greater than that required for replacement. The tumor thus formed grows more or less exponentially. From time to time, a cell may break off (metastasize) from the tumor and start up a new tumor at some distance from the primary (original) tumor. The objective of treatment is to find and excise the primary before it has had a chance to form metastases. If this is done, then the surgeon (or radiologist) will have "gotten it all" and the patient is cured. If metastases are formed before the primary is removed, then a cure is unlikely, but the life of the patient may be extended and ameliorated by aggressive administration of chemotherapeutic agents which will kill tumor cells more vigorously than normal cells. Unfortunately, since the agents do attack normal cells as well, a cure of metastasized cancer is unlikely, since the patient's body cannot sustain the dosage required to kill all the cancer cells.

For some cancers, breast cancer, for example, long-term cure rates had not improved very much for many years.

6.4.2 Models and Algorithms

One conjecture, consistent with a roughly constant intensity of display of secondary tumors, is that a patient with a tumor of a particular type is not displaying breakaway colonies only, but also new primary tumors resulting from suppression of a patient's immune system to attack tumors of

a particular type. We can formulate axioms at the micro level which will incorporate the mechanism of new primaries. Such an axiomatization has been formulated by Bartoszyński, Brown, and Thompson [3]. The first five axioms are consistent with the classical view as to metastatic progression. Hypothesis 6 is the mechanism we introduce to explain the nonincreasing intensity function of secondary tumor display.

Hypothesis 1. For any patient, each tumor originates from a single cell and grows at exponential rate α .

Hypothesis 2. The probability that the primary tumor will be detected and removed in $[t, t + \Delta t)$ is given by $bY_0(t)\Delta t + o(\Delta t)$, and until the removal of the primary, the probability of a metastasis in $[t, t + \Delta t)$ is $aY_0(t)\Delta t + o(\Delta t)$, where $Y_0(t)$ is the size of the primary tumor at time t .

Hypothesis 3. For patients with no discovery of secondary tumors in the time of observation, S , put $m_1(t) = Y_1(t) + Y_2(t) + \dots$, where $Y_i(t)$ is the size of the i th originating tumor. After removal of the primary, the probability of a metastasis in $[t, t + \Delta t)$ equals $a m_1(t) + o(\Delta t)$, and the probability of detection of a new tumor in $[t, t + \Delta t)$, is $b m_1(t) + o(\Delta t)$.

Hypothesis 4. For patients who do display a secondary tumor, after removal of the primary and before removal of Y_1 , the probability of detection of a tumor in $[t, t + \Delta t)$ equals $bY_1(t) + o(\Delta t)$, while the probability of detection of a metastasis is $aY_1(t) + o(\Delta t)$.

Hypothesis 5. For patients who do display a secondary tumor, the probability of a metastasis in $[t, t + \Delta t)$ is $a m_2(t)\Delta t + o(\Delta t)$, while the probability of detection of a tumor is $b m_2(t)\Delta t + o(\Delta t)$, where $m_2(t) = Y_2(t) + \dots$

Hypothesis 6. The probability of a systemic occurrence of a tumor in $[t, t + \Delta t)$ equals $\lambda\Delta t + o(\Delta t)$, independent of the prior history of the patient.

Essentially, we shall attempt to develop the likelihood function for this model so that we can find the values of a, b, α , and λ which maximize the likelihood of the data set observed. It turns out that this is a formidable task indeed. The SIMEST algorithm which we develop later gives a quick alternative to finding the likelihood function. However, to give the reader some feel as to the complexity associated with model aggregation from seemingly innocent axioms, we shall give some of the details of getting the likelihood function. First of all, it turns out that in order to have any hope of obtaining a reasonable approximation to the likelihood function, we will have to make some further simplifying assumptions. We shall refer to the period prior to detection of the primary as Phase 0. Phase 1 is the period from detection of the primary to S' , the first time of detection of a secondary tumor. For those patients without a secondary tumor, Phase 1 is the time of observation, S . Phase 2 is the time, if any, between S' and S . Now for the two simplifying axioms. T_0 is defined to be the (unobservable) time between the origination of the primary and the time when it is detected

and removed (at time $t = 0$). T_1 and T_2 are the times until detection and removal of the first and second of the subsequent tumors (times to be counted from $t = 0$). We shall let X be the total mass of all tumors other than the primary at $t = 0$.

Hypothesis 7. For patients who do not display a secondary tumor, growth of the primary tumor, and of all tumors in Phase 1, is deterministically exponential with the growth of all other tumors treated as a pure birth process.

Hypothesis 8. For patients who display a secondary tumor, the growth of the following tumors is treated as deterministic: in Phase 0, tumors $Y_0(t)$ and $Y_1(t)$; in Phase 1, tumor $Y_1(t)$ and all tumors which originated in Phase 0; in Phase 2, all tumors. The growth of remaining tumors in Phases 0 and 1 is treated as a pure birth process.

We now define

$$H(s; t, z) = \exp\left\{\frac{az}{\alpha}e^{\alpha t}(e^s - 1)\log[1 + (e^{-\alpha t} - 1)e^{-s}]\right. \\ \left. + \frac{\lambda}{\alpha}s - \frac{\lambda}{\alpha}\log[1 + e^{\alpha t}(e^s - 1)]\right\} \quad (6.24)$$

and

$$p(t; z) = bze^{\alpha t} \exp\left[-\frac{bz}{\alpha}(e^{\alpha t} - 1)\right]. \quad (6.25)$$

Further, we shall define

$$w(y) = \lambda \left[\int_0^y e^{-\nu(u)} du - y \right], \quad (6.26)$$

where $\nu(u)$ is determined from

$$u = \int_0^\nu (a + b + \alpha s - ae^{-s})^{-1} ds. \quad (6.27)$$

Then, we can establish the following propositions, and from these, the likelihood function:

$$p(T_0 > \tau) = \exp\left[-b \int_0^\tau e^{\alpha t} dt\right] = \exp\left[-\frac{b}{\alpha}(e^{\alpha \tau} - 1)\right]. \quad (6.28)$$

For patients who do not display a secondary tumor, we have

$$P(T_1 > S | X = x) = \exp[-x\nu(S) + w(S)]. \quad (6.29)$$

For patients who develop metastases, we have

$$P(T_1 > S) = P(\text{no secondary tumor in } (0, S)) \\ = \int_0^\infty e^{w(s)} p(t; 1) H(\nu(s); t, 1) dt. \quad (6.30)$$

Similarly, for patients who do display a secondary tumor, we have

$$\begin{aligned}
 P(T_1 = S', T_2 > S) = & \int_0^\infty \int_0^t e^{w(S-S')} p(t; 1) p(S'; e^{\alpha u}) (\lambda + a e^{\alpha(t-u)}) \\
 & \times \exp \left[-\lambda(t-u) - \frac{a}{\alpha} (e^{\alpha(t-u)} - 1) \right] H(\nu(S-S'); S', e^{\alpha u}) \\
 & \quad \times H(\nu(S-S') e^{\alpha S'}; u, e^{\alpha(t-u)}) du dt \\
 & + \int_0^\infty \int_0^{S'} e^{w(S-S')} p(t; 1) \exp \left[-\lambda t - \frac{a}{\alpha} (e^{\alpha t} - 1) \right] \lambda e^{-\lambda u} \\
 & \quad \times p(S' - u; 1) H(\nu(S-S'); S' - u, 1) du t (6.31)
 \end{aligned}$$

Finding the likelihood function, even a quadrature approximation to it, is more than difficult. Furthermore, current symbol manipulation programs (e.g., Mathematica, Maple) do not have the capability of doing the work. Accordingly, it must be done by hand. Approximately 1.5 person years were required to obtain a quadrature approximation to the likelihood. Before starting this activity, we had no idea of the rather practical difficulties involved. However, the activity was not without reward.

We found estimates for the parameter values using a data set consisting of 116 women who presented with primary breast cancer at the Curie-Sklodowska Cancer Institute in Warsaw (time units in months, volume units in cells): $a = .17 \times 10^{-9}$, $b = .23 \times 10^{-8}$, $\alpha = .31$, and $\lambda = .0030$. Using these parameter values, we found excellent agreement between the proportion free of metastasis versus time obtained from the data and that obtained from the model, using the parameter values given above. When we tried to fit the model to the data with the constraint that $\lambda = 0$ (that is, disregarding the systemic process as is generally done in oncology), the attempt failed.

One thing one always expects from a model-based approach is that, once the relevant parameters have been estimated, many things one had not planned to look for can be found. For example, tumor doubling time is 2.2 months. The median time from primary origination to detection is 59.2 months and at this time the tumor consists of 9.3×10^7 cells. The probability of metastasis prior to detection of the primary is .069, and so on. A model-based approach generally yields such serendipitous results, as a nonparametric approach generally does not. It is worth mentioning that, more frequently than one realizes, we need an analysis which is flexible, in the event that at some future time we need to answer questions different from those originally posed. The quadrature approximation of the likelihood is relatively inflexible compared to the simulation-based approach we shall develop shortly.

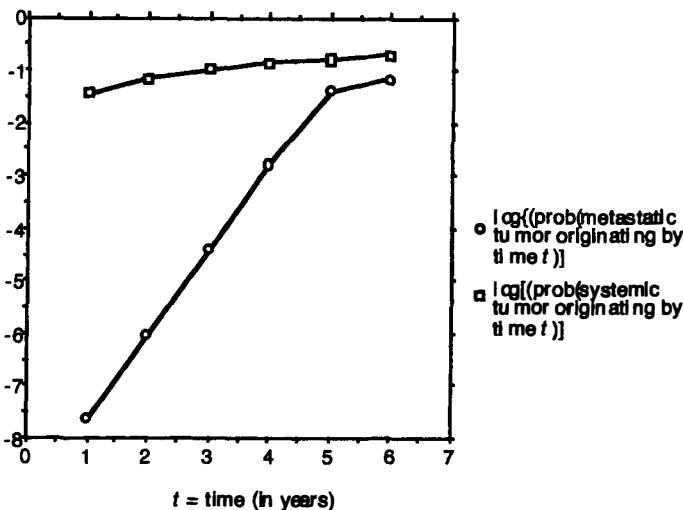


Figure 6.3. Metastatic and systemic effects.

Insofar as the relative importance of the systemic and metastatic mechanisms, in causing secondary tumors associated with breast cancer, it would appear from Figure 6.3 that the systemic is the more important. This result is surprising, but is consistent with what we have seen in our exploratory analysis of another tumor system (melanoma). Interestingly, it is by no means true that for all tumor systems the systemic term has such dominance. For primary lung cancer, for example, the metastatic term appears to be far more important.

It is not clear how to postulate, in any definitive fashion, a procedure for testing the null hypothesis of the existence of a systemic mechanism in the progression of cancer. We have already noted that when we suppress the systemic hypothesis, we cannot obtain even a poor maximum likelihood fit to the data. However, someone might argue that a different set of nonsystemic axioms should have been proposed. Obviously, we cannot state that it is simply impossible to manage a good fit without the systemic hypothesis. However, it is true that the nonsystemic axioms we have proposed are a fair statement of traditional suppositions as to the growth and spread of cancer.

As a practical matter, we had to use data that were oriented toward the life of the patient rather than toward the life of a tumor system. This is due to the fact that human *in vivo* cancer data is seldom collected with an idea toward modeling tumor systems. For a number of reasons, including the difficulty mentioned in obtaining the likelihood function, deep stochastic modeling has not traditionally been employed by many investigators in oncology. Modeling frequently precedes the collection of the kinds of data of greatest use in the estimation of the parameters of the model. Anyone who has gone through a modeling exercise such as that covered in this section

is very likely to treat such an exercise as a once in a lifetime experience. It simply is too frustrating to have to go through all the flailing around to come up with a quadrature approximation to the likelihood function. As soon as a supposed likelihood function has been found, and a corresponding parameter estimation algorithm constructed, the investigator begins a lengthy "debugging" experience. The algorithm's failure to work might be due to any number of reasons (e.g., a poor approximation to the likelihood function, a poor quadrature routine, a mistake in the code of the algorithm, inappropriateness of the model, etc.). Typically, the debugging process is time consuming and difficult. If one is to have any hope for coming up with a successful model-based investigation, an alternative to the likelihood procedure for aggregation must be found.

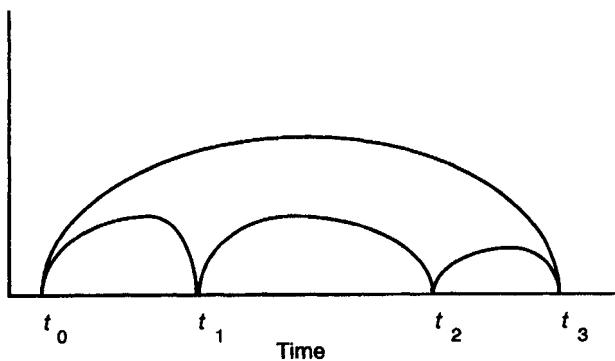


Figure 6.4. Two possible paths from primary to secondary.

To decide how best to construct an algorithm for parameter estimation which does not have the difficulties associated with the classical closed-form approach, we should try to see just what causes the difficulty with the classical method of aggregating from the microaxioms to the macro level, where the data lives. A glance at Figure 6.4 reveals the problem with the closed-form approach.

The axioms of tumor growth and spread are easy enough to implement in the forward direction. Indeed, they follow the natural forward formulation used since Poisson's work of 1837 [17]. Essentially, we are overlaying stochastic processes, one on top of the other, and interdependently to boot. But when we go through the task of finding the likelihood, we are essentially seeking all possible paths by which the observables could have been generated. The secondary tumor, originating at time t_3 , could have been thrown off from the primary at time t_3 , or it could have been thrown off from a tumor that itself was thrown off from another tumor at time t_2 which itself was thrown off from a tumor at time t_1 from the primary that originated at time t_0 . The number of possibilities is, of course, infinite.

In other words, the problem with the classical likelihood approach in the present context is that it is a backward look from a database generated in the forward direction. To scientists before the present generation of

fast, cheap computers, the backward approach was, essentially, unavoidable unless one avoided such problems (a popular way out of the dilemma). However, we need not be so restricted.

Once we realize the difficulty when one uses a backward approach with a concatenation of forwardly axiomitized mechanisms, the way out of our difficulty is clear [1, 23]. We need to analyze the data using a forward formulation. The most obvious way to carry this out is to pick a guess for the underlying vector of parameters, put this guess in the micro-axiomitized model and simulate many times of appearance of secondary tumors. Then, we can compare the set of simulated quasidata with that of the actual data.

The greater the concordance, the better we will believe we have done in our guess for the underlying parameters. If we can quantitize this measure of concordance, then we will have a means for guiding us in our next guess. One such way to carry this out would be to order the secondary occurrences in the data set from smallest to largest and divide them into k bins, each with the same proportion of the data. Then, we could note the proportions of quasidata points in each of the bins. If the proportions observed for the quasidata, corresponding to parameter value Θ , were denoted by $\{\pi_j(\Theta)\}_{j=1}^k$, then a Pearson goodness-of-fit statistic would be given by

$$\chi^2(\Theta) = \sum_{j=1}^k \frac{(\pi_j(\Theta) - 1/k)^2}{\pi_j(\Theta)}. \quad (6.32)$$

The minimization of $\chi^2(\Theta)$ provides us with a means of estimating Θ .

Typically, the sample size, n , of the data will be much less than N , the size of the simulated quasidata. With mild regularity conditions, assuming there is only one local maximum of the likelihood function, Θ_0 , as $n \rightarrow \infty$ (which function we of course do not know), then as $N \rightarrow \infty$, as n becomes large and k increases in such a way that $\lim_{n \rightarrow \infty} k = \infty$ and $\lim_{n \rightarrow \infty} k/n = 0$, the minimum χ^2 estimator for Θ_0 will have an expected mean square error which approaches the expected mean square error of the maximum likelihood estimator. This is, obviously, quite a bonus. Essentially, we will be able to forfeit the possibility of knowing the likelihood function and still obtain an estimator with asymptotic efficiency equal to that of the maximum likelihood estimator. The price to be paid is the acquisition of a computer swift enough and cheap enough to carry out a very great number, N , of simulations, say 10,000. This ability to use the computer to get us out of the “backward trap” is a potent but, as yet seldom used, bonus of the computer age. Currently, the author is using SIMEST on a 2 gigaherz personal computer, amply adequate for the task, which now costs around \$1000.

First, we observe how the forward approach enables us to eliminate those hypotheses which were, essentially a practical necessity if a likelihood function was to be obtained. Our new axioms are simply:

Hypothesis 1. For any patient, each tumor originates from a single cell and grows at exponential rate α .

Hypothesis 2. The probability that the primary tumor will be detected and removed in $[t, t + \Delta t]$ is given by $bY_0(t)\Delta t + o(\Delta t)$. The probability that a tumor of size $Y(t)$ will be detected in $[t, t + \Delta t]$ is given by $bY(t)\Delta t + o(\Delta t)$.

Hypothesis 3. The probability of a metastasis in $[t, t + \Delta]$ is $a\Delta t \times$ (total tumor mass present).

Hypothesis 4. The probability of a systemic occurrence of a tumor in $[t, t + \Delta t]$ equals $\lambda\Delta t + o(\Delta t)$, independent of the prior history of the patient.

In order to simulate, for a given value of (α, a, b, λ) , a quasidata set of secondary tumors, we must first define:

t_D = time of detection of primary tumor;

t_M = time of origin of first metastasis;

t_S = time of origin of first systemic tumor;

t_R = time of origin of first recurrent tumor;

t_d = time from t_R to detection of first recurrent tumor;

t_{DR} = time from t_D to detection of first recurrent tumor.

Now, generating a random number u from the uniform distribution on the unit interval, if $F(\cdot)$ is the appropriate cumulative distribution function for a time, t , we set $t = F^{-1}(u)$. Then, assuming that the tumor volume at time t is

$$v(t) = ce^{\alpha t}, \text{ where } c \text{ is the volume of one cell,} \quad (6.33)$$

we have

$$F_M(t) = 1 - \exp\left(-\frac{ac}{\alpha}e^{\alpha t_M}\right). \quad (6.34)$$

Similarly, we have

$$\begin{aligned} F_D(t_D) &= 1 - \exp\left(-\int_0^{t_D} bce^{\alpha\tau}d\tau\right) \\ &= 1 - \exp\left(-\frac{bc}{\alpha}e^{\alpha t_D}\right), \end{aligned} \quad (6.35)$$

$$F_S = 1 - e^{-\lambda t_S}, \quad (6.36)$$

and

$$F_d(t_d) = 1 - \exp\left(-\frac{bc}{\alpha}e^{\alpha t_d}\right). \quad (6.37)$$

Using the actual times of discovery of secondary tumors $t_1 \leq t_2 \leq \dots \leq t_n$, we generate k bins. In actual tumor situations, because of recording

protocols, we may not be able to put the same number of secondary tumors in each bin. Let us suppose that the observed proportions are given by (p_1, p_2, \dots, p_k) . We shall generate N recurrences $s_1 < s_2 < \dots < s_N$. The observed proportions in each of the bins will be denoted $\pi_1, \pi_2, \dots, \pi_k$. The goodness of fit corresponding to (α, λ, a, b) will be given by

$$\chi^2(\alpha, \lambda, a, b) = \sum_{j=1}^k \frac{(\pi_j(\alpha, \lambda, a, b) - p_j)^2}{\pi_j(\alpha, \lambda, a, b)}. \quad (6.38)$$

As a practical matter, we may replace $\pi_j(\alpha, \lambda, a, b)$ by p_j , since with (α, λ, a, b) far away from truth, $\pi_j(\alpha, \lambda, a, b)$ may well be zero. Then the following algorithm generates the times of detection of quasisecondary tumors for the particular parameter value (α, λ, a, b) .

Secondary Tumor Simulation (α, λ, a, b)

```

Generate  $t_D$ 
 $j = 0$ 
 $i = 0$ 
Repeat until  $t_M(j) > t_D$ 
   $j = j + 1$ 
  Generate  $t_M(j)$ 
  Generate  $t_{dM}(j)$ 
   $t_{dM}(j) \leftarrow t_{dM}(j) + t_M(j)$ 
  If  $t_{dM}(j) < t_D$ , then  $t_{dM}(j) \leftarrow \infty$ 
  Repeat until  $t_S > 10t_D$ 
     $i = i + 1$ 
    Generate  $t_{dS}(i)$ 
     $t_{dS}(i) \leftarrow t_{dS}(i) + t_S(i)$ 
     $s \leftarrow \min [t_{dM}(j), t_{dS}(i)]$ 
  Return  $s$ 
End Repeat

```

This algorithm does still have some simplifying assumptions. For example, we assume that metastases of metastases will probably not be detected before the metastases themselves. We assume that the primary will be detected before a metastasis, and so on. Note, however, that the algorithm uses much less restrictive simplifying assumptions than those that led to the terms of the closed-form likelihood such as (6.31). Even more importantly, the Secondary Tumor Simulation algorithm can be discerned in a few minutes, whereas a likelihood argument is frequently the work of months.

Another advantage of the forward simulation approach is its ease of modification. Those who are familiar with “backward” approaches based on the likelihood or the moment generating function are only too familiar with the experience of a slight modification causing the investigator to go back to the

start and begin anew. This is again a consequence of the tangles required to be examined if a backward approach is used. However, a modification of the axioms generally causes slight inconvenience to the forward simulator.

For example, we might add:

Hypothesis 5. A fraction γ of the patients ceases to be at systemic risk at the time of removal of the primary tumor if no secondary tumors exist at that time. A fraction $1 - \gamma$ of the patients remains at systemic risk throughout their lives.

Clearly, adding Hypothesis 5 will cause considerable work if we insist on using the classical aggregation approach of maximum likelihood. However, in the forward simulation method we simply add the following lines to the Secondary Tumor Simulation code:

Generate u from $U(0,1)$

If $u > \gamma$, then proceed as in the Secondary Tumor Simulation code

If $u < \gamma$, then proceed as in the Secondary Tumor Simulation code except replace the step “Repeat until $t_S > 10t_D$ ” with the step “Repeat until $t_S(i) > t_D$.”

In the discussion of metastasis and systemic occurrence of secondary tumors, we have used a model supported by data to try to gain some insight into a part of the complexities of the progression of cancer in a patient. Perhaps this sort of approach should be termed *speculative data analysis*. In the current example, we were guided by a nonparametric intensity function estimate [2], which was surprisingly nonincreasing, to conjecture a model, which enabled us to test systemic origin against metastatic origin on something like a level playing field. The fit without the systemic term was so bad that anything like a comparison of goodness-of-fit statistics was unnecessary.

It is interesting to note that the implementation of SIMEST is generally faster on the computer than working through the estimation with the closed-form likelihood. In the four-parameter oncological example we have considered here, the running time of SIMEST was 10% of the likelihood approach. As a very practical matter, then, the simulation-based approach would appear to majorize that of the closed-form likelihood method in virtually all particulars. The running time for SIMEST can begin to become a problem as the dimensionality of the response variable increases past one. Up to this point, we have been working with the situation where the data consists of failure times. In the systemic versus metastatic oncogenesis example, we managed to estimate four parameters based on this kind of one-dimensional data. As a practical matter, for tumor data, the estimation of five or six parameters for failure time data are the most one can hope for. Indeed, in the oncogenesis example, we begin to observe the beginnings

of singularity for four parameters, due to a near trade-off between the parameters a and b . Clearly, it is to our advantage to be able to increase the dimensionality of our observables. For example, with cancer data, it would be to our advantage to use not only the time from primary diagnosis and removal to secondary discovery and removal, but also the tumor volumes of the primary and the secondary. Such information enables one to postulate more individual growth rates for each patient. Thus, it is now appropriate to address the question of dealing with multivariate response data.

Gaussian Template Criterion. In many cases, it will be possible to employ a procedure using a criterion function. Such a procedure has proved quite successful in another context (see [22]). First, we transform the data $\{X_i\}_{i=1}^n$ by a linear transformation such that for the transformed data set $\{U_i\}_{i=1}^n$ the mean vector becomes zero and the covariance matrix becomes I :

$$U = AX + b. \quad (6.39)$$

Then, for the current best guess for Θ , we simulate a quasidata set of size N . Next, we apply the same transformation to the quasidata set $\{Y_j(\Theta)\}_{j=1}^N$, yielding $\{Z_j(\Theta)\}_{j=1}^N$. Assuming that both the actual data set and the simulated data set come from the same density, the likelihood ratio $\Lambda(\Theta)$ should increase as Θ gets closer to the value of Θ , say Θ_0 , which gave rise to the actual data, where

$$\Lambda(\Theta) = \frac{\prod_{i=1}^n \exp[-\frac{1}{2}(u_{1i}^2 + \dots + u_{pi}^2)]}{\prod_{i=1}^N \exp[-\frac{1}{2}(z_{1i}^2 + \dots + z_{pi}^2)]}. \quad (6.40)$$

As soon as we have a criterion function, we are able to develop an algorithm for estimating Θ_0 . The closer Θ is to Θ_0 , the smaller will $\Lambda(\Theta)$ tend to be.

This procedure above that uses a single Gaussian template will work well in many cases where the data has one distinguishable center and a falling off away from that center which is not too "taily." However, there will be cases where we cannot quite get away with such a simple approach. For example, it is possible that a data set may have several distinguishable modes and/or exhibit very heavy tails. In such a case, we may be well advised to try a more local approach. Suppose that we pick one of the n data points at random—say x_1 —and find the m nearest-neighbors amongst the data. We then treat this m nearest-neighbor cloud as if it came from a Gaussian distribution centered at the sample mean of the cloud and with covariance matrix estimated from the cloud. We transform these $m + 1$ points to zero mean and identity covariance matrix, via

$$U = A_1 X + b_1. \quad (6.41)$$

Now, from our simulated set of N points, we find the $N(m + 1)/n$ simulated points nearest to the mean of the $m + 1$ actual data points. This will

give us an expression like

$$\Lambda_1(\Theta) = \frac{\prod_{i=1}^{m+1} \exp[-\frac{1}{2}(u_{1i}^2 + \dots + u_{pi}^2)]}{\prod_{i=1}^{N(m+1)/n} \exp[-\frac{1}{2}(z_{1i}^2 + \dots + z_{pi}^2)]}. \quad (6.42)$$

If we repeat this operation for each of the n data points, we will have a set of local likelihood ratios $\{\Lambda_1, \Lambda_2, \dots, \Lambda_n\}$. Then one natural measure of concordance of the simulated data with the actual data would be

$$\Lambda(\Theta) = \sum_{i=1}^n \log(\Lambda_i(\Theta)). \quad (6.43)$$

We note that this procedure is not equivalent to one based on density estimation, because the nearest-neighbor ellipsoids are not disjoint. Nevertheless, we have a level playing field for each of the guesses for Θ and the resulting simulated data sets.

A Simple Counting Criterion. Fast computing notwithstanding, with n in the 1000 range and N around 10,000, the template procedure can become prohibitively time consuming. Accordingly, we may opt for a subset counting procedure:

For data size n , pick a smaller value, say nn .

Pick a random subset of the data points of size nn .

Pick a nearest neighbor outreach parameter m , typically $0.02n$.

For each of the nn data points, X_j , find the Euclidean distance to the m th nearest neighbor, say $d_{j,m}$.

For an assumed value of the vector parameter Θ , generate N simulated observations.

For each of the data points in the random subset of the data, find the number of simulated observations within $d_{j,m}$, say $N_{j,m}$.

Then the criterion function becomes

$$\chi^2(\Theta) = \sum_{j=1}^{nn} \frac{((m+1)/n - N_{j,m}/N)^2}{(m+1)/n}.$$

Experience indicates that whatever nn size subset of the data points is selected should be retained throughout the changes of Θ . Otherwise, practical instability may obscure the path to the minimum value of the criterion function.

A SIMDAT SIMEST Stopping Rule. In Section 4.2 we considered a situation where we compared the results from resampled data points with

those from model-based simulations. SIMDAT is not a simple resampling so much as it is a stochastic interpolator. We can take the original data and use SIMDAT to generate a SIMDAT pseudodata set of N values.

Then, for a particular guess of Θ , we can compute a SIMEST pseudodata set of N values. For any region of the space of the vector observable, the number of SIMEST-generated points should be approximately equal to the number of SIMDAT-generated points. For example, let us suppose that we pick nn of the n original data points and find the radius $d_{j,m}$ of the hypersphere which includes m of the data points for, say, point X_j . Let $N_{j,SD}$ be the number of SIMDAT-generated points falling inside the hypersphere and $N_{j,SE}$ be the number of SIMEST-generated points falling inside the hypersphere. Consider the empirical goodness-of-fit statistic for the SIMDAT cloud about point X_j :

$$\chi_{j,SD}^2(\Theta) = \frac{((m+1)/n - N_{j,SD}/N)^2}{(m+1)/n}.$$

For the SIMEST cloud, we have

$$\chi_{j,SE}^2(\Theta) = \frac{((m+1)/n - N_{j,SE}/N)^2}{(m+1)/n}.$$

If the model is correct and if our estimate for Θ is correct, then $\chi_{j,SE}^2(\Theta)$ should be, on the average, distributed similarly to $\chi_{j,SD}^2(\Theta)$. Accordingly, we can construct a sign test. To do so, let

$$\begin{aligned} W_j &= +1 \text{ if } \chi_{j,SD}^2(\Theta) \geq \chi_{j,SE}^2(\Theta) \\ &= -1 \text{ if } \chi_{j,SD}^2(\Theta) < \chi_{j,SE}^2(\Theta). \end{aligned}$$

So, if we let

$$Z = \frac{\sum_{j=1}^{nn} W_j}{\sqrt{nn}},$$

we might decide to terminate our search for estimating Θ when the absolute value of Z falls below 3 or 4.

Problems

6.1. Generate a sample of size 100 of firings at a bull's-eye of radius 5 centimeters where the distribution of the shots is circular normal with mean the center of the bullseye and deviation 1 meter.

- (a) Generate and display a bootstrapped sample of size 1000. Do you find any simulated points inside the bull's-eye?
- (b) Then using $m=10$, generate a SIMDAT pseudosample of size 1000. Do you find any simulated points inside the bull's-eye?

6.2. A multivariate distribution with heavy tales may be generated as follows. First, we generate a χ^2 variable v with 2 degrees of freedom. Then we generate p independent univariate normal variates from a normal distribution with mean 0 and variance 1. Then $\mathbf{X}' = (X_1, X_2, \dots, X_p)$ will have the multivariate normal distribution $\mathcal{N}_p(\mathbf{0}, \mathbf{I})$. Moreover,

$$t_{2,p}(\mu) = \frac{\mathbf{X}}{\sqrt{v/2}} + \mu$$

will have a shifted t distribution with 2 degrees of freedom. Generate a sample of size 500 from the mixture distribution

$$f = .9\mathcal{N}_3(\mathbf{0}, \mathbf{I}) + .1t_{2,3}(\mathbf{0}).$$

Can you tell the difference between the sample from the mixture distribution above and a sample of size 500 from $\mathcal{N}_3(\mathbf{0}, \mathbf{I})$?

6.3. Let us consider von Bortkiewicz's suicide data from eight German states over a period of 14 years, as shown in Table 6.1

Table 6.1. Actual and Expected Numbers of Suicides per Year												Sum
Suicides	0	1	2	3	4	5	6	7	8	9	≥ 10	Sum
Freq.	9	19	17	20	15	11	8	2	3	5	3	112
$E(\text{freq.})$	3.5	12.1	21	24.3	21	14.6	8.5	4.2	1.9	.7	.2	112

Using the Poisson model for k , the number of suicides generated in a year,

$$P(k|\theta) = e^{-\theta} \frac{\theta^k}{k!},$$

find a SIMEST estimator for θ , using as the criterion function

$$\sum_{i=0}^{10} \frac{(X_i - E_i)^2}{E_i} \approx \chi^2(k-1).$$

(We recall here that the category "10" here is ≥ 10 .)

6.4. Generally speaking, before using an algorithm for parameter estimation on a set of data, it is best to use it on a set of data simulated from the proposed model. Returning to the problem in Section 6.4.1, generate a set of times of discovery of secondary tumor (time measured in months past discovery and removal of primary) of 400 patients with $a = .17 \times 10^{-9}$, $b = .23 \times 10^{-8}$, $\alpha = .31$, and $\lambda = .0030$. Using SIMEST, see if you can recover the true parameter values from various starting values.

6.5. Using the parameter values given in Problem 6.4, generate a set of times of discovery of first observed secondary tumor and second observed secondary tumor. Using SIMEST, see whether you can recover the true parameter values from various starting values.

References

- [1] Atkinson, E.N., Bartoszyński, R., Brown, B. W., and Thompson, J. R. (1983). "Simulation techniques for parameter estimation in tumor related stochastic processes," *Proceedings of the 1983 Computer Simulation Conference*. New York: North-Holland, 754–757.
- [2] Bartoszyński, Robert, Brown, B. W., McBride, C. M., and Thompson, J. R. (1981). "Some nonparametric techniques for estimating the intensity function of a cancer related nonstationary Poisson process," *Ann. Statist.*, **9**, 1050–1060.
- [3] Bartoszyński, R., Brown, B. W. and Thompson, J. R. (1982). "Metastatic and systemic factors in neoplastic progression," in *Probability Models and Cancer*, L. LeCam and J. Neyman, eds. New York: North-Holland., 253–264.
- [4] Boswell, S. B. (1983). *Nonparametric Mode Estimators for Higher Dimensional Densities*. Doctoral Dissertation. Houston: Rice University.
- [5] Bridges, E., Ensor, K. B., and Thompson, J. R. (1992). "Marketplace competition in the personal computer industry," *Decis. Sci.*, **23**, 467–477.
- [6] Cleveland, W. S. (1979). "Robust locally-weighted regression and smoothing scatterplots," *J. Amer. Statist. Assoc.*, **74**, 829–836.
- [7] Cox, D. D. (1983). "Asymptotics for M-type smoothing splines," *Ann. Statist.*, **17**, 530–551.
- [8] Efron, B. (1979). "Bootstrap methods—another look at the jackknife," *Ann. Statist.*, **7**, 1–26.
- [9] Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. New York: Chapman & Hall.
- [10] Elliott, M. N. and Thompson, J. R. (1993). "An exploratory algorithm for the estimation of mode location and numerosity in multidimensional data," *Proceedings of the Thirty-Eighth Conference on the Design of Experiments in Army Research Development and Testing*, B. Bodt, ed. Research Triangle Park, N.C.: Army Research Office, 229–244.
- [11] Eubank, R. L. (1988). *Smoothing Splines and Nonparametric Regression*. New York: Marcel Dekker.
- [12] Guerra, V., Tapia, R. A., and Thompson, J. R. (1976). "A random number generator for continuous random variables based on interpolation procedure of Akima's," *Comput. Sci. Statist.*, 228–230.

- [13] Härdle, W. (1990). *Applied Nonparametric Regression*. New York: Oxford University Press.
- [14] Hart, J. D. and Wehrly, T. E. (1986) "Kernel regression estimation using repeated measurements data," *J. Amer. Statist. Assoc.*, **81**, 1080–1088.
- [15] Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized Additive Models*. New York: Chapman & Hall.
- [16] Mack, Y. P. and Rosenblatt, Murray (1979). "Multivariate k -nearest neighbor density estimates," *J. Multivariate Analysis*, **9**, 1–15.
- [17] Poisson, S. D. (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile, précédées des réglés générales du calcul des probabilités*. Paris.
- [18] Rosenblatt, M. (1956). "Remarks on some nonparametric estimates of a density function," *Ann. Math. Statist.*, **27**, 832–835.
- [19] Scott, D.W. and Thompson, J.R. (1983). "Probability density estimation in higher dimensions," in *Computer Science and Statistics*, J. Gentle, ed., Amsterdam: North Holland, 173–179.
- [20] Scott, D.W. (1992). *Multivariate Density Estimation*. New York: John Wiley & Sons.
- [21] Taylor, M. S. and Thompson, J. R. (1982). "A data based random number generator for a multivariate distribution," *Proceedings of the NASA Workshop on Density Estimation and Function Smoothing*. College Station, Texas: Texas A & M University Department of Statistics, 214–225.
- [22] Taylor, M. S. and Thompson, J. R. (1986). "A data based algorithm for the generation of random vectors," *Comp. Statist. Data Anal.*, **4**, 93–101.
- [23] Thompson, J. R., Atkinson, E. N., and Brown, B. W. (1987). "SIMEST: an algorithm for simulation based estimation of parameters characterizing a stochastic process," in *Cancer Modeling*. Thompson, J. and Brown, B., eds. New York: Marcel Dekker. 387–415.
- [24] Thompson, J. R. (1989). *Empirical Model Building*. New York: John Wiley & Sons, 35–43.
- [25] Thompson, J. R. and Tapia, R. A. (1990). *Nonparametric Function Estimation, Modeling, and Simulation*, Philadelphia: SIAM, 154–161, 214–233.

- [26] Thompson, J. R. and Koronacki, J. (1993). *Statistical Process Control for Quality Improvement*. New York: Chapman & Hall.
- [27] Tukey, J. W. (1977). *Exploratory Data Analysis*. Reading, Mass.: Addison-Wesley.

Chapter 7

Exploratory Data Analysis

7.1 Introduction

The concept of exploratory data analysis (EDA) is generally dismissive of models. The concept of EDA was proposed by the late John W. Tukey [1] of Princeton University and Bell Labs (best known in the popular press for having coined the term “software,” to scientists as possibly the most important statistician ever). Tukey wrote an entire book on the subject (*Exploratory Data Analysis*, 1977), and an entire semester can easily be spent studying that book. Nevertheless, we can find out a great deal of EDA and demonstrate much of its utility in rather a short amount of time.

It is clear to the reader by this time that I personally am of the opinion that we learn by developing a chain of ever improving models. By “improving,” I mean, in general, getting closer to the truth. EDAers, however, are rather postmodern in that they question whether the notion of “truth” is meaningful or relevant. As one of them, Professor William Eddy of Carnegie Mellon University, put it:

The data analytic method denies the existence of “truth”; the only knowledge is empirical.

.....

The only purpose of models is to make formal implications. For far too long statisticians have concentrated on fitting models to data. And, for reasons I don’t fully understand, they have “tested” the parameters of these models. The relevance of models comes only from their implications and the interpre-

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

tation thereof. If we can make predictions without models, I think we should.

To forecast purely on the basis of empiricism, without understanding the driving forces of the system under investigation is hardly desirable. Of course, we have to do this often enough. If we see that a stock is 100 on Monday and then goes to 110 by the next Monday and 120 the Monday after that and 130 the Monday after that, we may well be tempted to forecast that on the Monday after the 130 reading the stock will be around 140. But few of us would like to "bet the farm on the forecast." Suppose, however, that we found out that the company had developed a drug for curing AIDS but had not formally announced the discovery, and insiders were discreetly buying blocks of the stock, that might make us feel better about forecasting the 140 price.

Eddy [1] does not just take it as a practical necessity that sometimes we must predict without understanding the underlying process; he proudly announces that dispensing with modeling is a good thing. This kind of radical empiricism has resonance with the "Prince of Nominalists" William of Ockham. Ockham did not deny the existence of reality. He simply thought that we, as a practical matter, are unwise to fail to act until we know that reality rather precisely.

In order to understand the highly anti-modeling paradigm of Exploratory Data Analysis, we naturally need to model EDA itself. Here are some personal observations which capture much of the essence of EDA:

1. The eye expects adjacent pixels to be likely parts of a common whole.
2. The eye expects continuity (not quite the same as (1)).
3. As points move far apart, the human processor needs training to decide when points are no longer to be considered part of the common whole. Because of the ubiquity of situations where the Central Limit Theorem, in one form or another, applies, a natural benchmark is the normal distribution.
4. Symmetry reduces the complexity of data.
5. Symmetry essentially demands unimodality.
6. The only function that can be identified by the human eye is the straight line.
7. A point remains a point in any dimension.

7.2 Smoothing

During the early days of the U.S. space program, photographs were taken of the surface of the moon. These were not very useful without some processing because there was a great deal of noise imposed on the signal. However, let us use the first principle given above

1. The eye expects adjacent pixels to be likely parts of a common whole.

In Figure 7.1, we show five adjacent pixels. If we believe that each of these has more or less the same light intensity, then we might simply replace the observed light intensity of the center point by the average of the intensities of all five points. We know if we have five measurements of the same thing, each contaminated by noise, then the average gives us a better measure of the uncontaminated intensity than any single observation.

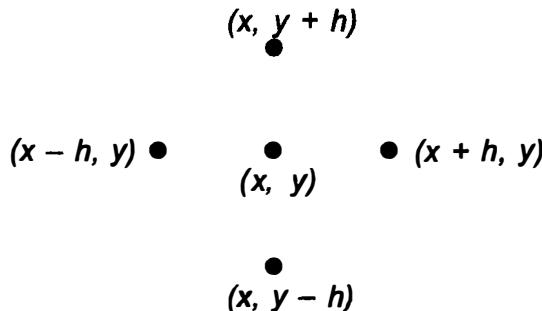


Figure 7.1. Five adjacent pixels.

So one way to improve our estimate of the reflected light intensity at the point x, y would be to use

$$I(x, y) \leftarrow \frac{I(x, y) + I(x+h, y) + I(x, y+h) + I(x-h, y) + I(x, y-h)}{5}. \quad (7.1)$$

But this means of proceeding might be a bit extreme, for we do not usually suppose that the adjacent points all have the same intensity. Accordingly, we might rather use

$$I(x, y) \leftarrow \frac{4I(x, y) + I(x+h, y) + I(x, y+h) + I(x-h, y) + I(x, y-h)}{8}. \quad (7.2)$$

Obviously, we could use this *digital filter* over and over until we felt we had about the right resolution. If we use the filter too many times, then every pixel will have roughly the same intensity. We will have a smooth picture, for we will have reduced the variability almost to zero. We will be close to the position of the stopped clock which is, of course, precisely correct twice in a 24-hour day.

However, if we do not smooth enough, we will have a photograph that is so jumpy we have difficulty making out any patterns on the moon's surface. The trick is knowing just when to stop. Human beings can achieve

this rather well by simply looking at the resulting picture after each smooth. Something of this sort was done by JPL (Jet Propulsion Laboratory) scientists with the Mariner probe pictures of the moon. The filter was used by the founder of exploratory data analysis, John Tukey. However, the name of the filter was given by Tukey to Julius von Hann, a nineteenth-century Austrian weather scientist, who used something like *Hanning* in obtaining land coordinate indexed temperature profiles. With Hanning, we put half the weight on the intensity of (x, y) and then distribute the rest of the weight equally on the surrounding points.

Table 7.1. Production.

1	133
2	155
3	199
4	193
5	201
6	235
7	185
8	160
9	161
10	182
11	149
12	110
13	72
14	101
15	60
16	42
17	15
18	44
19	60
20	86
21	50
22	40
23	34
24	40
25	33
26	67
27	73
28	57
29	85
30	90

$$I(h) \leftarrow \frac{2I(x) + I(x+h) + I(x-h)}{4}. \quad (7.3)$$

Consider a small company that makes van modifications transforming vans into campers. In Table 7.1 we show production figures on 30 consecu-

tive days. The Hanning smoother originally was developed for dealing with information indexed by two spatial dimensions. However, it can be used very nicely to smooth data indexed to a one-dimensional variable.

We graph this information in Figure 7.2. The points we see are real production figures. There are no errors in them. Nevertheless, most people would naturally smooth them. Perhaps the figures are real, but the human observer wants them to be smooth, not rough. One could look upon such a tendency to want smoothness as being some sort of natural Platonic notion hardwired into the human brain. We can rationalize this tendency by saying that the real world has smooth productions contaminated by shocks such as sudden cancellations of orders, sickness of workers, etc. But the fact is that most of us would hate to make plans based on such a jumpy plot.

The fact is that the world in which we live does tend to move rather smoothly in time. If there were so much turbulence that this was not the case, it would be hard to imagine how any sort of civilized society could ever have developed. So we ought not despise our apparently instinctive desire to see smooth curves rather than jumpy ones.

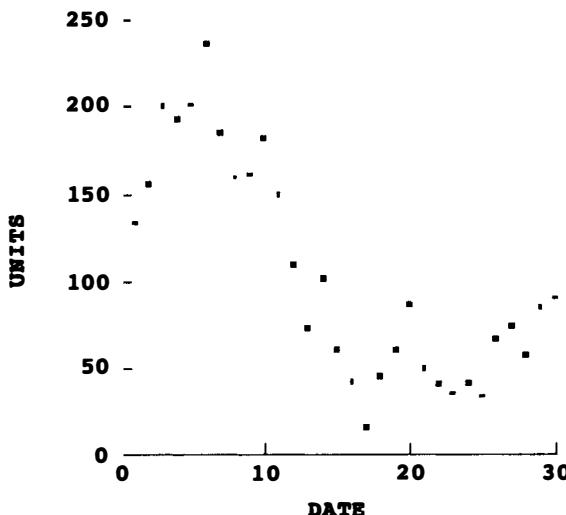


Figure 7.2. Daily production data.

So, how shall we smooth the data in Table 7.1 (and Figure 7.2)? There seemsto be no natural unique answer. There are obviously many contenders. One way would be simply to connect all the data points. We show this graph in Figure 7.3.

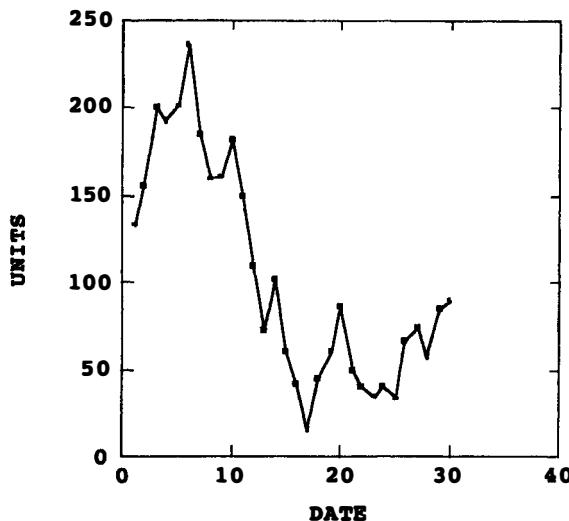


Figure 7.3. Piecewise linear graph.

We might try some sort of freehand smoothing such as that shown in Figure 7.4.

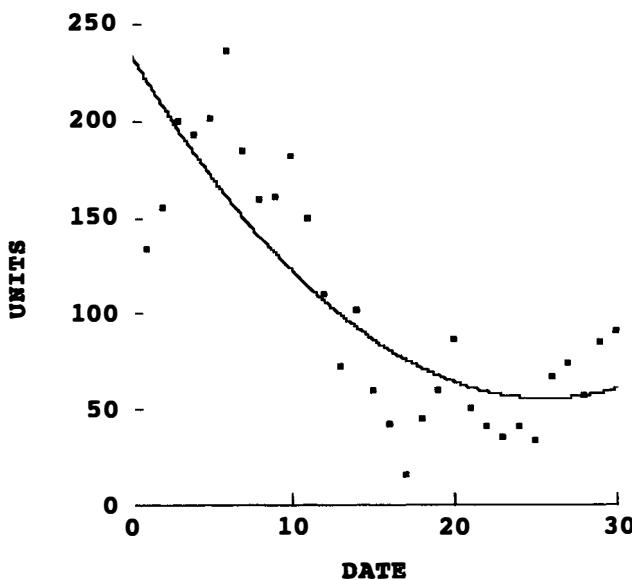


Figure 7.4. Freehand smooth.

In Figure 7.5 we observe the results of smoothing by Hanning one time. We note that the smooth appears to be rather local. We have not taken such

a global smoothing approach as we did in Figure 7.4.

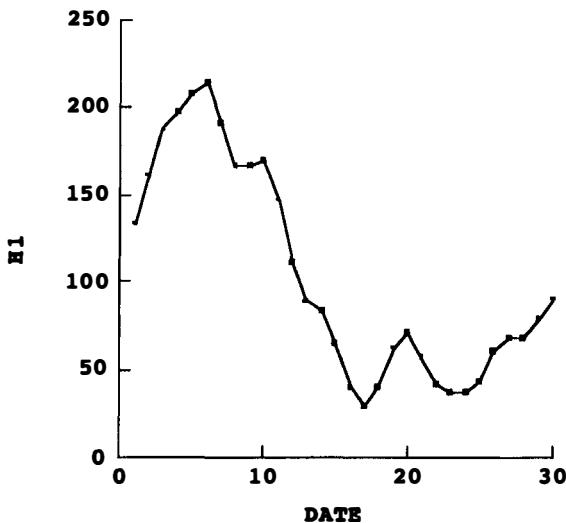


Figure 7.5. Hanning smooth.

Let us hann the hanning smooth to obtain the H2 graph in Figure 7.6.

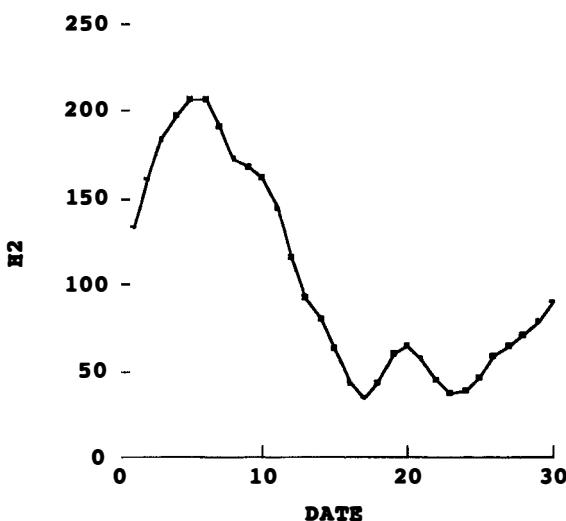


Figure 7.6. Double Hanning smooth.

Table 7.2. Various orders of Hannings.

Day	Production	H	HH	H5	H10
1	133	133	133	133	133
2	155	161	160	157	155
3	199	187	183	178	171
4	193	197	197	192	183
5	201	208	207	199	190
6	235	214	207	198	190
7	185	191	191	189	185
8	160	167	173	178	176
9	161	166	167	167	165
10	182	169	163	155	151
11	149	148	144	138	135
12	110	110	114	118	118
13	72	89	93	97	100
14	101	84	81	80	83
15	60	66	64	64	69
16	42	40	44	51	59
17	15	29	35	45	53
18	44	41	43	48	51
19	60	63	59	54	52
20	86	71	65	57	52
21	50	57	56	54	51
22	40	41	44	47	49
23	34	37	38	43	47
24	40	37	38	43	48
25	33	43	46	48	51
26	67	60	58	56	57
27	73	68	66	64	64
28	57	68	71	72	72
29	85	79	79	81	81
30	90	90	90	90	90

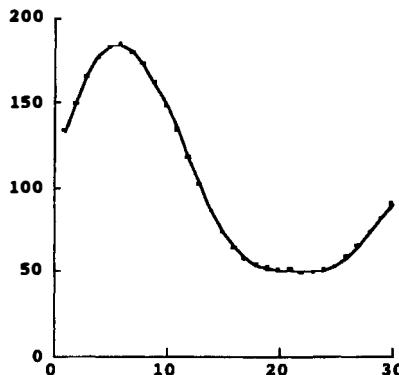


Figure 7.7. Fifteen-fold Hanning smooth.

For many purposes, the single Hanning or the double Hanning seems to give a reasonable smooth one that is more local than the “freehand smooth” one in Figure 7.4. One problem with Hanning is that “it does not know when to stop.” For any function other than a straight line, each successive application of Hanning will change the picture. Note that the 15-fold Hanning smooth has eliminated the local mode (peak) on day 20. There are other problems with hanning. One glitch will spread throughout the data set. In Table 7.3, we note that a string of ones to which has been added one 1,000 reading will, after a few hannings spread the glitch throughout the data.

Table 7.3. Contaminating Hanning Smooth.

Data	H	HH	HHH
1	1	1	1
1	1	1	16.61
1	1	63.44	94.66
1	250.75	250.75	235.14
1,000	500.50	375.62	313.18
1	250.75	250.75	235.14
1	1	63.44	94.66
1	1	1	16.61

A Hanning filter has advantages. It was used to provide smoothed pixels for flyby photographs of the moon in the 1960s. However, it would be nice if we had a filter that automatically functions without the necessity of a human observer to fine tune it. Consider Figure 7.8 which is one of many hundreds of shots taken of the moon by the Clementine Project in 1994.

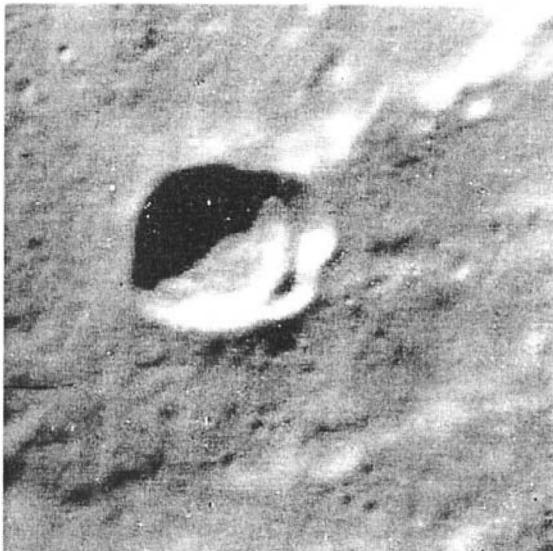


Figure 7.8. Clementine shot of moon crater.

For processing such material, we need something which will not destroy the edges of the crater by oversmoothing. The median smooth of Tukey [1] has such a desired effect. It replaces the averaging process in (7.2) by the median process

$$I(x, y) \leftarrow \text{Med} (I(x, y), I(x + h, y), I(x, y + h), I(x - h, y) + I(x, y - h)). \quad (7.4)$$

In other words, we replace the raw light intensity at (x, y) by the median intensity of the five points in the vicinity of (x, y) . This smoother “knows when to stop,” and it does not smooth out crater edges that need to be seen. We have here an example of a truly revolutionary result that requires practically no mathematics but is powerful nonetheless. It is also significant to note that it took decades from the time such a filter was needed until it was discovered.

Let us apply the median smooth (in one dimension, called the “3 smooth”) to the production data in Table 7.1. We note that there is no change in the table after we “smooth by 3R” the second time. Indeed, only one value changes from the first to the second smoothing by 3’s. And this is the big advantage of the median smooth. Unlike the Hanning smooth, it will not, if applied ad infinitum, simply smooth the data to destruction. It automatically stops, usually after no more than three iterations. We show these results in Table 7.4, where changed values are indicated in boldface. The expression “3R” means that the median smooth has been applied until no further changes were observed.

We note that there is a tendency for the 3R (median) smooth to give artificially broad peaks and valleys. This is generally a very minor problem. For many applications, the 3R smooth is all we need.

Day	Production	3	33=3R
1	133	133	133
2	155	155	155
3	199	193	193
4	193	199	199
5	201	201	201
6	235	201	201
7	185	185	185
8	160	161	161
9	161	161	161
10	182	161	161
11	149	149	149
12	110	110	110
13	72	101	101
14	101	72	72
15	60	60	60
16	42	42	42
17	15	42	42
18	44	44	44
19	60	60	60
20	86	60	60
21	50	50	50
22	40	40	40
23	34	40	40
24	40	34	40
25	33	40	40
26	67	67	67
27	73	67	67
28	57	73	73
29	85	85	85
30	90	90	90

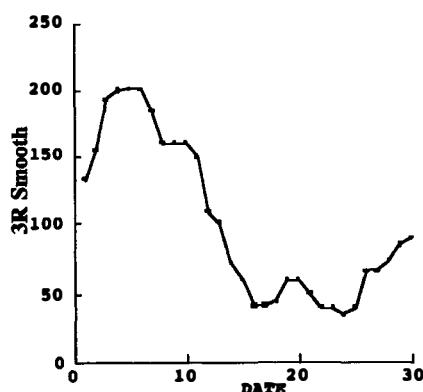


Figure 7.9. The 3R (median) smooth.

If we wish to make things look a little better aesthetically, we can take the results of the 3R smooth and hann them twice as we show in Table 7.5 and Figure 7.10. The advantage of the 3RHH smooth is that it is fully automated. We need not have an observer in the loop. This is a big advantage for most applications. The very big deal of this chapter on smoothing is the median smooth. It simply was done looked at until the 1960s. This is a really big discovery which requires very little in the way of mathematical sophistication. It is amazing how frequently such big discoveries are really simple mathematically. Frequently, genius consists in thinking of things to try and then seeing whether they work.

Table 7.5. 3RHH Smooth.

Day	Production	3	33=3R	3RH	3RHH
1	133	133	133	133	133
2	155	155	155	159	159
3	199	193	193	185	182
4	193	199	199	198	196
5	201	201	201	201	199
6	235	201	201	197	195
7	185	185	185	183	183
8	160	161	161	167	170
9	161	161	161	161	162
10	182	161	161	158	155
11	149	149	149	142	140
12	110	110	110	118	118
13	72	101	101	96	97
14	101	72	72	76	77
15	60	60	60	59	60
16	42	42	42	47	49
17	15	42	42	43	45
18	44	44	44	48	49
19	60	60	60	56	55
20	86	60	60	58	55
21	50	50	50	50	50
22	40	40	40	43	44
23	34	40	40	39	39
24	40	34	34	37	40
25	33	33	40	45	47
26	67	67	67	60	59
27	73	67	67	69	69
28	57	73	73	79	79
29	85	85	85	88	86
30	90	90	90	90	90

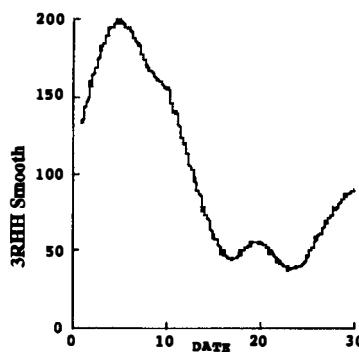


Figure 7.10. The 3RHH smooth.

7.3 The Stem and Leaf Plot

In 1965, a smalltown mayor takes a random sample of yearly incomes from 30 householders. The results are (in dollars) 5600, 8700, 9200, 9900, 10,100, 11,200, 13,100, 16,100, 19,300, 23,900, 25,100, 25,800, 28,100, 31,000, 31,300, 32,400, 35,800, 37,600, 40,100, 42,800, 47,600, 49,300, 53,600, 55,600, 58,700, 63,900, 72,500, 81,600, 86,400, 156,400. The mayor wants to get some notion of what the data might be saying about income distributions in his town. Of course simply looking at the data as a ordered text file above gives her some information. For example, the average income is 39,423. However, this one number contains only part of the important information in the 30 numbers. Next, consider the *stem and leaf plot*, demonstrated in Table 7.6.

Table 7.6. Stem and Leaf Plot.

0	5899
10	01369
20	3558
30	11257
40	0279
50	358
60	3
70	2
80	16
90	
100	
110	
120	
130	
140	
150	6

We note that the left hand column gives a marker graded in tens of thousands. The rows give the appropriate number of whole thousands of dollars of income. The table here is a kind of hybrid between a table and a graph. It looks very much like a histogram turned sideways. Such a table/graph is called a *stem and leaf plot*. It retains almost all the information of the original salary list but displays it in such a way that we immediately can make rough inferences from it. For example, the the highest income in the survey looks very different from the rest. The smallest income in the list does not look atypical of the rest of the incomes in the sense that there does not be a break in continuity of incomes between that lowest income and the rest.

7.4 The Five Figure Summary

Now, we know that there are 30 incomes in our list. The *sample median* income will be obtained by looking at the average of the fifteenth and sixteenth incomes.

$$M = \frac{1}{2}[31,300 + 32,400] = 31,850. \quad (7.5)$$

To obtain a notion of the spread of the income data, we look at the *lower quartile* and the *upper quartile*. One-fourth of 30 is 7.5. So, the lower quartile is found by taking the average of the seventh and eighth observations. We will call this value the *lower hinge*:

$$LH = \frac{1}{2}[13,100 + 16,100] = 14,600. \quad (7.6)$$

Similarly, for the *upper hinge*, we have

$$UH = \frac{1}{2}[53,600 + 55,600] = 54,600. \quad (7.7)$$

This gives us, then, for the *five figure summary* of the data:

Table 7.7. Five Figure Summary.

M15	31,850
H8	31,850
1	56,000

54,600

156,400

Here, we show the smallest and largest observations, the upper and lower quartiles (hinges) and the sample median.

7.5 Tukey's Box Plot

So far, things look rather intuitive. But we now wish to come up with a scheme for deciding whether some of the observations appear to be rather

untypical of the others. So we now come up with some spread measures which do not look, at first glance, to be so intuitive. The first is that of a *step*, which is defined to be 1.5 times the difference between the upper and lower quartiles (hinges). Here,

$$H = 1.5 \times (54,600 - 31,850) = 34,125. \quad (7.8)$$

The notion here is that any observation which is more than one *step* beyond the lower quartile or the upper quartile begins to be suspicious in consideration of whether it is typical of the rest of the observations. Any observation that is more than two *steps* beyond either of the quartiles is probably not typical of the rest of the observations. One step beyond the upper quartile is termed the *upper inner fence*. Two steps beyond the upper quartile (upper hinge) is termed the *upper outer fence*. Here the *upper outer fence* is seen to be

$$UOF = 2 \times 34,125 + 54,600 = 122,850. \quad (7.9)$$

The largest income of 156,400 is seen to be even greater than the *UOF*, so it is probably not representative of the other data points. Perhaps, this individual came into a windfall of profit due to some sale of property. Perhaps he/she comes from an unusually rich part of the town. The data alone cannot tell us what is going on. It does tell us that the income here is unusually high in comparison to the incomes of the group.

Now for the income data at hand, the *schematic plot* also called a *box plot* is given in Figure 7.11.

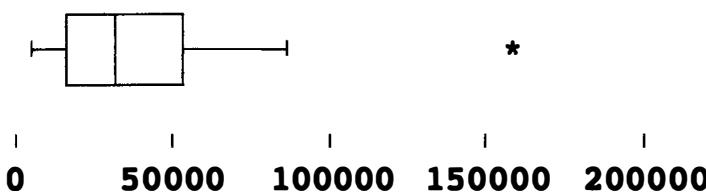


Figure 7.11. Box plot of income data.

We note that the median is clearly shown with a vertical line toward the middle of the "box." The upper and lower hinges (quartiles) are shown as vertical lines forming the limits of the "box." There is an asterisk showing any value outside the upper or lower outer fences. The first value inside the lower inner fence is indicated as the lower boundary of the "whisker" line. The first value inside the upper inner fence is indicated as the upper boundary of the "whisker" line. Other than the individual making 156,400, no person sampled has an income outside the inner fences. The incomes seem to be somewhat homogeneous except for that one very high income.

The box plot is a very handy device for getting a preliminary look at a data set. It does, in a fairly well structured way, things that have historically been done in a less orderly way. Let us consider an example of Tukey [1]. During the winter of 1893–94, W.J. Rayleigh was examining determinations of the densities of nitrogen obtained from a variety of sources. We show these densities in Table 7.8. Now, the “OTHER” refers to any source other than air. Rayleigh noted that regardless of the source of nitrogen (other than air), all the densities appeared to be roughly the same. He noted this in a somewhat nonspecific search pattern going through his data. But we can perceive Rayleigh’s discovery easily by making a box plot of his “OTHER” data as we show in Figure 7.12.

Table 7.8. Measurements of N Densities.

AIR	OTHER
2.31024	2.30143
2.31030	2.29816
2.31028	2.30182
2.31017	2.29890
2.30986	2.29940
2.31010	2.29849
2.31001	2.29889
	2.29889

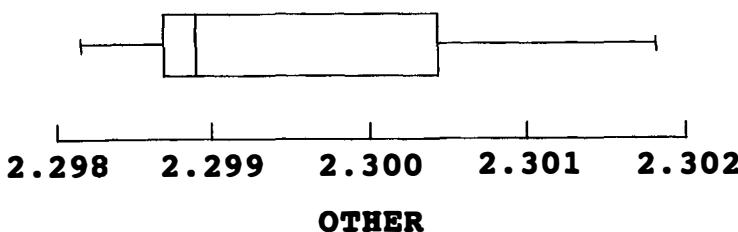


Figure 7.12. Nitrogen densities from sources other than AIR.

We note that the density of nitrogen as measured from the AIR sources are higher than those from the OTHER sources. To see how much this is so, let us take the smallest of AIR source observations, 2.30986, and add it to the OTHER data points, making a box plot of the pooled values (as shown in Figure 7.13).

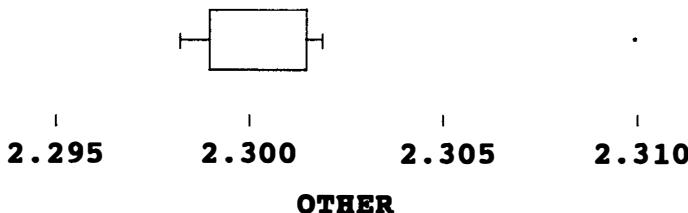


Figure 7.13. Nitrogen densities from OTHER sources plus one AIR source observation.

We note that the AIR source observation is well outside the upper outer fence. We should strongly question whether the AIR source nitrogen is really the same as the OTHER source nitrogen. Indeed, it turns out that the AIR sources contained traces of a then unknown element, namely argon. This empirically obtained observation of Lord Rayleigh contributed toward his receiving a Nobel Prize in Physics in 1906.

Problems

7.1. By means of a sketch supported by an argument, show a hypothetical data plot for which a Hanning smooth carried out until the smooth looks satisfactory will be disastrous , but a 3R smooth will work well.

7.2. By means of a sketch and an argument show a hypothetical data plot for which a 3R smooth is unsatisfactory but a Hanning smooth is satisfactory.

7.3. Construct a stem-and-leaf plot for county populations in the state of Texas or the state of California or the state of New York. Carry out appropriate transformations to symmetrize the plot (if this should be necessary). Construct five figure summaries. Find the inner and outer fences and construct a schematic plot.

7.4. Given the following table, infer the functional relationship between X and Y.

X	Y
.40	0.048
1.10	0.363
1.91	1.094
4.22	5.343

References

- [1] Eddy, W. (1979). "Discussion of a paper by Emanuel Parzen," *Journal of the American Statistical Association*, 1979, pp. 124-125.
- [2] Tukey, J. W. (1977). *Exploratory Data Analysis*. Reading, Mass.: Addison-Wesley Publishing Company.

Chapter 8

Noise Killing Chaos

Anything could be true. The so-called laws of nature were nonsense. The law of gravity was nonsense. “If I wished,” O’Brien had said, “I could float off this floor like a soap bubble.” Winston worked it out. “If he *thinks* he floats off the floor, and I simultaneously *think* I see him do it, then the thing happens.”

George Orwell, 1984

8.1 Introduction

Many regard noise and chaos to be the same thing. This impression is very wrong. Noise is a naturally occurring phenomenon. Willy-nilly, the television may go dark. The salmon harvest may explode. And, we can go backward in time to see what caused the disruption. The stock market may plummet. By looking backward in time, we can generally find the causes of the fall.

In chaos theory, it is not randomness that is the villain. It is some underlying part of the mechanism from which we cannot escape. Hence, weather forecasts are essentially impossible as are forecasting changes in the stock market.

In the ensuing argument, we will attempt to show that if chaos existed in the real world, it would be destroyed by naturally occurring noise.

Postmodernism is one of the latest intellectual schools to be inflicted on the West since the French Revolution. Postmodernism does more than cast doubt on objective reality; it flatly denies it. The perception of a particular reality is deemed to be highly subjective and hugely nonstationary, even for

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

the same observer. Those schooled in the implicit Aristotelian modality of reason and logic (and that still includes most scientists) can easily follow the postmodernist train of thought to the point where they note that the assumptions of the postmodernists really bring one to such a level of chaos that conversation itself becomes impractical and useless. 30 years ago one could find few people presenting papers on such subjects as *postmodern science*; this has become more frequent lately. Postmodernism, as the latest of the pre-Socratic assaults on reason based on facts, was gently smiled at by scientists simply as a bizarre attempt of their humanistic colleagues to appear to be doing something useful. In the last few years, however, the leakage of postmodern modalities of thinking into the sciences has increased significantly. When one looks in library offerings under such topics as chaos theory or fuzzy logic, one frequently finds "postmodern" as a correlative listing. For some of us Aristotelians, this tendency is alarming. But many mathematicians are pleased to find that their abstract ideas are appreciated by literati and the popular press.

Every generation or so, there crops up a notion in *mainstream science* which is every bit as antilogical and antirealistic as postmodernism itself. These notions either do not admit of scientific validation or, for some reason, are somehow exempted from it. Once these *new ideas* have been imprimaturized by the scientific establishment and developed into a systematology, large amounts of funding are provided by government agencies, and persons practicing such arts are handsomely rewarded by honors and promotion until lack of utility causes the systematology to be superseded or subsumed by some other *new idea*.

It is intriguing (and should be noncomforting to stochasticians) that such *new ideas* tend to have as a common tendency the promise that practitioners of the new art will be able to dispense with such primitive notions as probability. Stochastics, after all, was simply a *patch*, an empirical artifice, to get around certain bumps in the scientific road—which jumps have now been smoothed level by the new art.

Frequently, not only will the new art not give any hope for solving a scientific problem, but also it will actually give comfort to an economist, say, or a meteorologist, or the members of some area of science which has not lived up to hopes, over the *fact* that their area simply does not lend itself to the solution of many of its most fundamental problems. It does not merely say that these problems are hard, but it argues that they cannot be solved now or ever. Those who have labored for decades on these hard problems may now safely down shovels, knowing that they gave their best to do what could not be done.

We have seen scientists take problems that were supposed to be insoluble and solve them. The "hopeless" hairtrigger of nuclear war was obviated by one of the founders of Monte Carlo techniques, Herman Kahn. Kahn essentially created the escalation ladder, giving the great powers a new grammar of discourse that enabled staged response to crises.

High-speed computing broke out of "natural" bounds of feasibility by the invention first of the transistor and then the microchip. Weather forecasting is still admittedly primitive, but apparently we are now in a position to carry out such tasks as forecasting severe as opposed to mild hurricane seasons.

To introduce the subject of chaos, let us first consider the Mandelbrot model,

$$z_{n+1} = z_n^2 + c, \quad (8.1)$$

where z is a complex variable and c is a complex constant. Let us start with

$$z_1 = c = -0.339799 + 0.358623i.$$

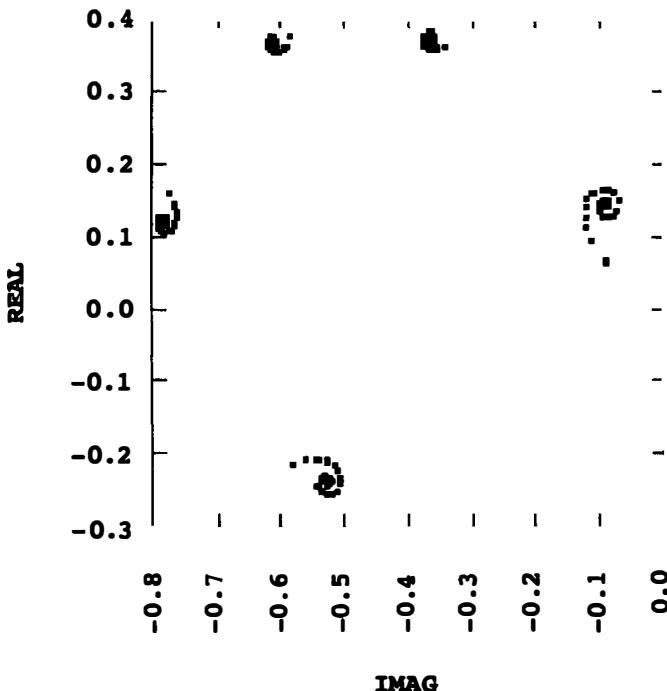


Figure 8.1. Mendelbrot set.

The iterative structure of this set is given in Figure 8.1. All the points exist in thin, curved manifolds. Most of the complex plane is empty. And, to make things more interesting, we do not trace smoothly within each manifold, but rather jump from manifold, to manifold. It is as though we started to drive to the opera hall in Houston, but suddenly found ourselves in downtown Vladivostok. Of course, it could be interesting for somebody who fantasizes about zipping from one part of the galaxy to another.

Now, the Mandelbrot model may itself be questioned as a bit bizarre when describing any natural phenomenon. Were it restricted to the reals,

then it would be, for a growth model, something that we would be unlikely to experience, because it is explosive. But let that pass. Since we are in the realm of *Gedankenspiel* anyway, let us see what happens when we introduce noise.

Consider the model

$$z_{n+1} = z_n^2 + c + 0.2u_{n+1} + 0.2v_{n+1}, \quad (8.2)$$

where the u and v are independent $\mathcal{N}(0,1)$ variables. We have been able to penetrate throughout the former empty space, and now we do not make quantum leaps from Houston to Vladivostok. Here, in the case of a formal mathematical structure, we simply note in passing how the introduction of noise can remove apparent pathology. Noise can and does act as a powerful smoother in many situations based on aggregation (Figure 8.2).

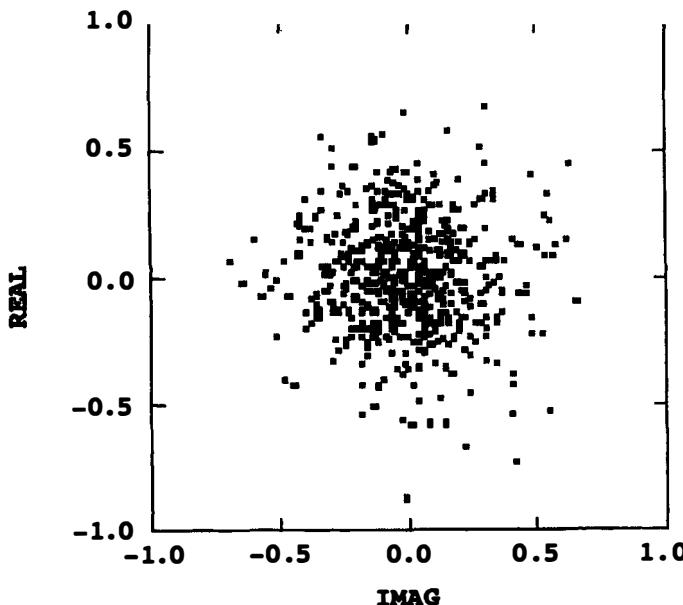


Figure 8.2. Mandelbrot plus noise.

8.2 The Discrete Logistic Model

Let us turn now to an example of a model of a real-world phenomenon, namely the growth of a population with finite food supply. One of the earliest was the 1844 logistic model of Verhulst:

$$\frac{dX}{dt} = X(\alpha - X), \quad (8.3)$$

where α essentially represents the limit, in population units, of the food supply. A solution to this model was obtained by Verhulst and is simply

$$X(t) = \frac{\alpha X(0) \exp(\alpha t)}{\alpha + X(0)[\exp(\alpha t) - 1]}. \quad (8.4)$$

Naturally, this model is only an approximation to the real growth of a population, but the mathematical solution is perfectly regular and without pathology.

Lorenz [2] has examined a discrete version of the logistic model:

$$X_n = X_{n-1}(a - X_{n-1}) = -X_{n-1}^2 + aX_{n-1}. \quad (8.5)$$

Using $X_0 = a/2$, he considers the time average of the modeled population size:

$$\bar{X} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N X_n. \quad (8.6)$$

For a values below $1 + \sqrt{6}$, the graph of \bar{X} behaves quite predictably. Above this value, great instability appears, as we show in Figure 8.3.

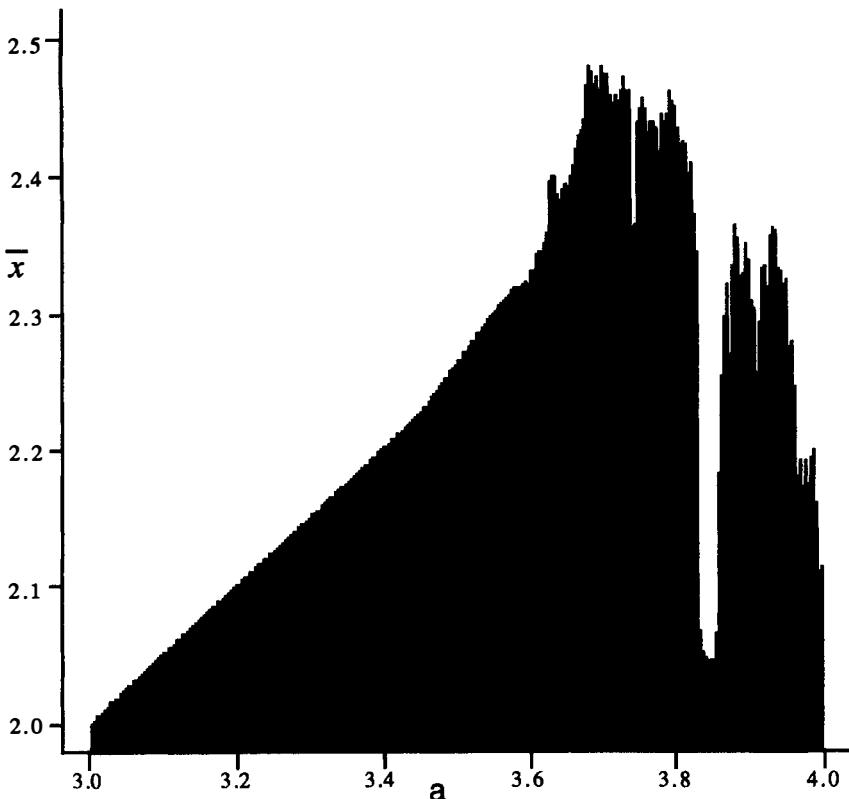


Figure 8.3. Discrete logistic model.

We note in Figures 8.4 and 8.5 how this *fractal* structure is maintained at any level of microscopic examination we might choose.

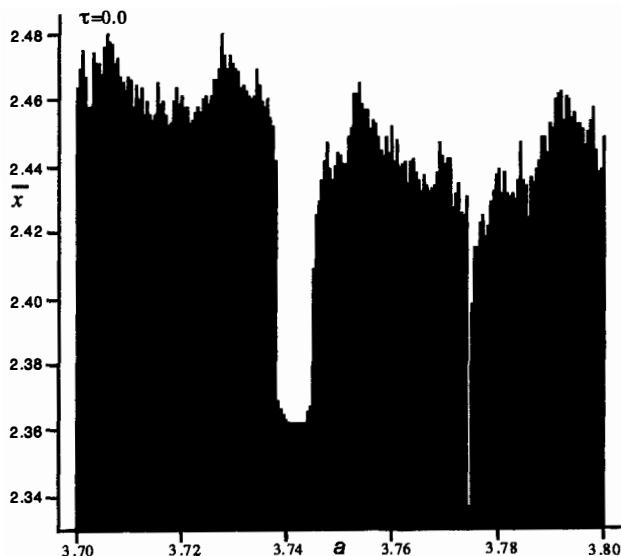


Figure 8.4. Discrete logistic model at small scale.

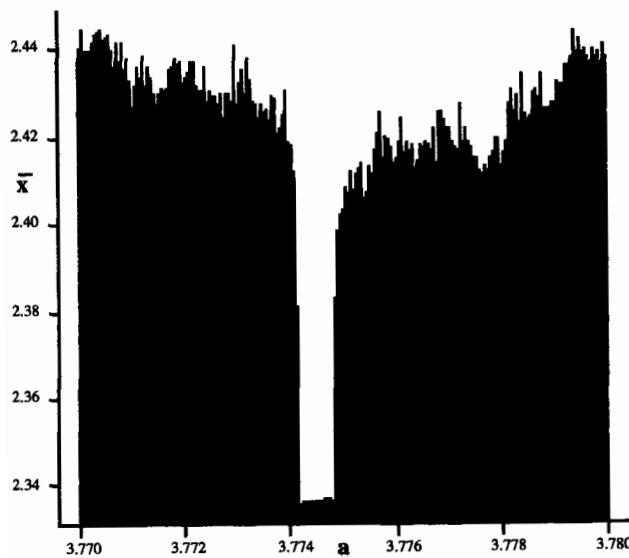


Figure 8.5. Discrete logistic model at very small scale.

Let us look at Figure 8.3 in light of real-world ecosystems. Do we know of systems where we increase the food supply slightly and the supported population crashes, and where we increase it again and it soars? What

should be our point of view concerning a model that produces results greatly at variance with the real world? And we recall that actually we have a perfectly good continuous 150 year old solution to the logistic equation. The use of the discrete logistic model is really natural only in the sense that we wish to come up with a model that can be put on a digital computer. In the case of chaos theory it is frequently the practice of enthusiasts to question not the model but the reality. So it is argued that, in fact, it is the discrete model which is the more natural. Let us walk for a time in that country.

For the kinds of systems the logistic model was supposed to describe, we could axiomatize by a birth-and-death process as follows:

$$\begin{aligned} P(\text{birth in } [t, t + \Delta t]) &= \beta(\gamma - X)X\Delta t \\ P(\text{death in } [t, t + \Delta t]) &= \eta X\Delta t. \end{aligned} \quad (8.7)$$

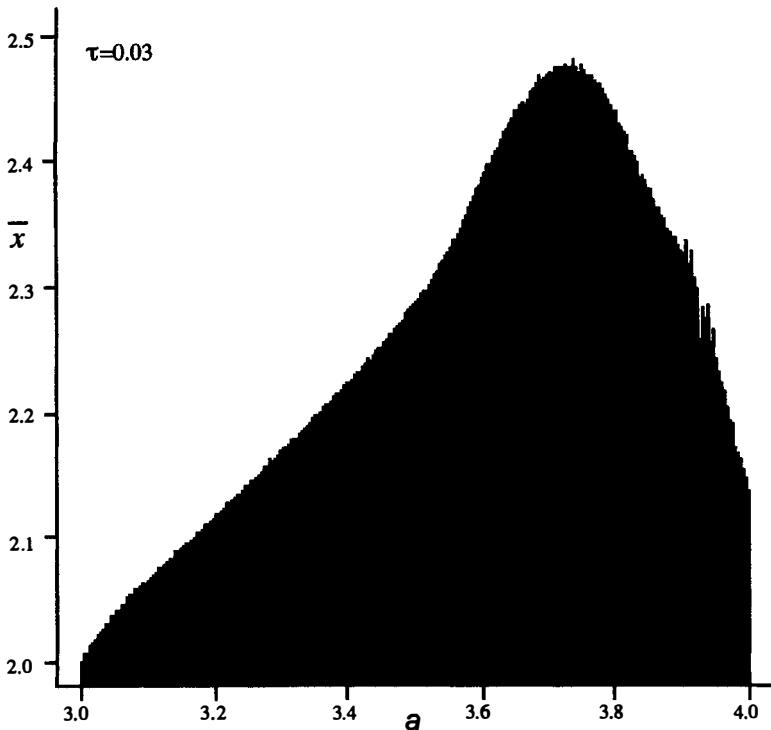


Figure 8.6. Discrete logistic model plus noise.

Perhaps Verhulst would have agreed that what he had in mind to do was to aggregate from such microaxioms but had not the computational ability to do so. Equation (8.5) was the best he could do. We have the computing power to deal directly with the birth-and-death model. However, we can make essentially the same point by adding a noise component to the logistic

model. We do so as follows:

$$\begin{aligned} X_n &= X_{n-1}(a - X_{n-1}) + \mu_{n-1}X_{m-1} \\ &= X_{n-1}([a + \mu_{n-1}] - X_{n-1}). \end{aligned} \quad (8.8)$$

where μ_{n-1} is a random variable from the uniform distribution on $(-\tau, \tau)$.

As a convenience, we add a bounceback effect for values of the population less than zero. Namely, if the model drops the population to $-\epsilon$, we record it as $+\epsilon$. In Figure 8.6 we note that the stochastic model produces no chaos (the slight fuzziness is due to the fact that we have averaged only 5000 successive X_n values). Nor is there fractal structure at the microscopic level, as we show in Figure 8.7 (using 70,000 successive X_n values).

Clearly, the noisy version of (8.5) is closer to the real world than the purely deterministic model. The food supply changes; the reproductive rate changes; the population is subject to constant change. However, the change itself induces stability into the system. The extreme sensitivity to a in Figure 8.3 is referred to as the butterfly effect. The notion is that if (8.5) described the weather of the United States, then one butterfly flying across a backyard could dramatically change the climate of the nation. Such an effect, patently absurd, is a failure of the model (8.5), not of the real world.

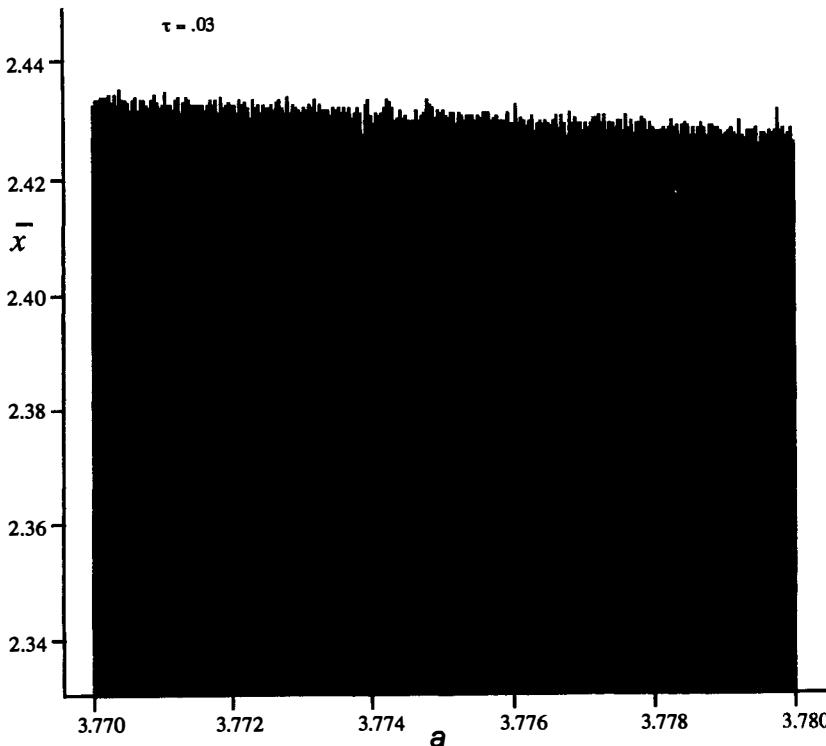


Figure 8.7. Discrete logistic model plus noise at small scale.

8.3 A Chaotic Convection Model

In 1963, after rounding off initial conditions, Lorenz [1] discovered that the following model was extremely sensitive to the initial value of (x, y, z) .

$$\begin{aligned}\frac{dx}{dt} &= 10(y - x) \\ \frac{dy}{dt} &= -xz + 28x - y \\ \frac{dz}{dt} &= xy - \frac{8}{3}z.\end{aligned}\tag{8.9}$$

In Figure 8.8 we show a plot of the system for 2000 steps using $\Delta t = .01$ and $\tau = 0.0$.

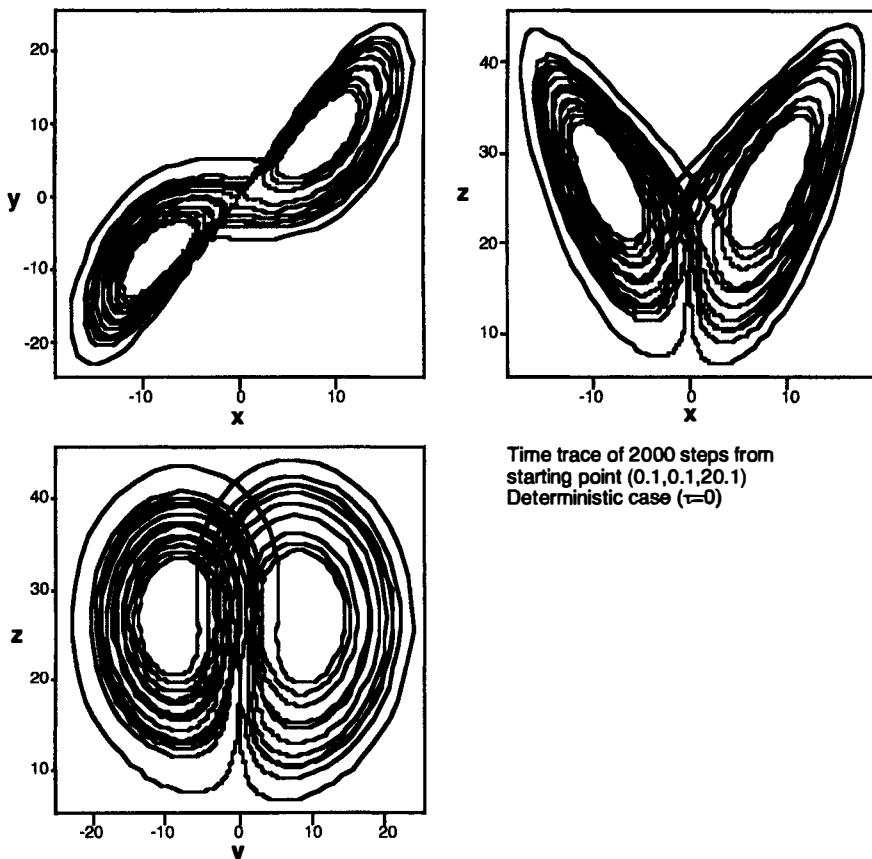


Figure 8.8. Lorenz weather model.

We observe the nonrepeating spiral which characterizes this chaotic model. The point to be made here is that, depending on where one starts the process, the position in the remote future will be very different. Lorenz uses

this model to explain the poor results one has in predicting the weather. Many have conjectured that such a model corresponds to a kind of uncertainty principle operating in fields ranging from meteorology to economics. Thus, it is opined that in such fields, although a deterministic model may accurately describe the system, there is no possibility of making meaningful long range forecasts, because the smallest change in the initial conditions dramatically changes the end result. The notion of such an uncertainty principle has brought comfort to such people as weather scientists and econometricians who are renowned for their inability to make useful forecasts. What better excuse for poor performance than a mathematical basis for its inevitability?

The philosophical implications of (8.9) are truly significant. Using (8.9), we know precisely what the value of (x, y, z) will be at any future time if we know precisely what the initial values of these variables are (an evident impossibility). We also know that the slightest change in these initial values will dramatically alter the result in the remote future. Furthermore, (8.9) essentially leads us to a dead end. If we believe this model, then it cannot be improved, for if at some time a better model for forecasting were available so that we really could know what (x, y, z) would be at a future time, then, since the chaos spirals are nonrepeating, we would be able to use our knowledge of the future to obtain a precise value of the present. Since infinite precision is not even a measurement possibility, we arrive at a practical absurdity.

If one accepts the ubiquity of chaos in the real world (experience notwithstanding), then one is essentially driven back to pre-Socratic modalities of thought, where experiments were not carried out, since it was thought that they would not give reproducible results. Experience teaches us that, with few exceptions, the models we use to describe reality are only rough approximations. Whenever a model is thought to describe a process completely, we tend to discover, in retrospect, that factors were missing, that perturbations and shocks entered the picture which had not been included in the model. A common means of trying to account for such phenomena is to assume that random shocks of varying amplitudes are constantly being applied to the system. Let us, accordingly, consider a discretized noisy version of (8.9):

$$\begin{aligned} x_n &= (1 + \mu_{x,n-1})x_{n-1} + \Delta t 10(y_{n-1} - x_{n-1}) \\ y_n &= (1 + \mu_{y,n-1})y_{n-1} + \Delta t(-x_{n-1}y_{n-1} + 28x_{n-1} - y_{n-1}) \\ z_n &= (1 + \mu_{z,n-1})z_{n-1} + \Delta t \left(x_{n-1}y_{n-1} - \frac{8}{3}z_{n-1} \right), \end{aligned} \quad (8.10)$$

where the μ 's are independently drawn from a uniform distribution on the interval from $(-\tau, \tau)$.

We consider in Figure 8.9 the final point at the 2500th step of each of 1000 random walks using the two initial points $(0.1, 0.1, 20.1)$ and $(-13.839,$

$-6.66, 40.289$), with $\tau = 0.0001$. These two starting points are selected since, in the deterministic case, the end results are dramatically different. In Figure 8.10, we show quantile-quantile plots for the two cases. We note that, just as we would expect in a model of the real world, the importance of the initial conditions diminishes with time, until, as we see from Figures 8.9 and 8.10, the distribution of end results is essentially independent of the initial conditions.

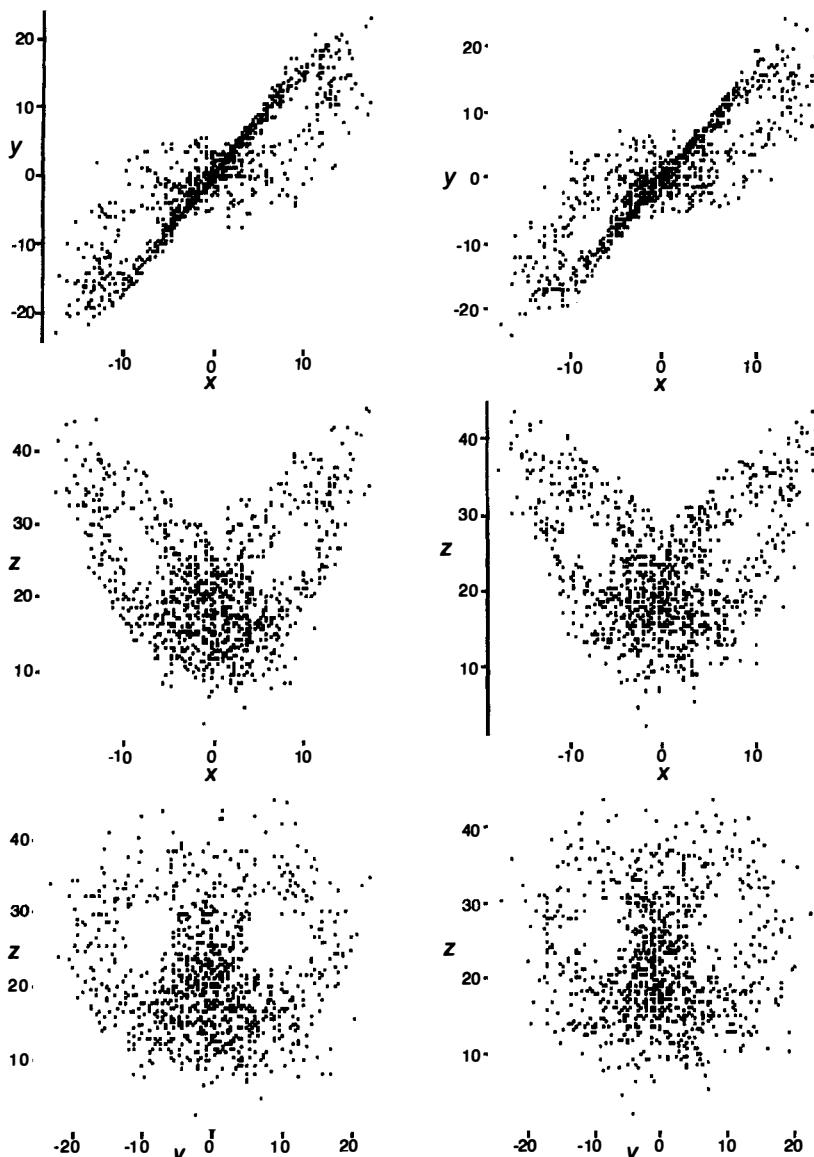


Figure 8.9. 10,000th step for each of 1000 time traces.

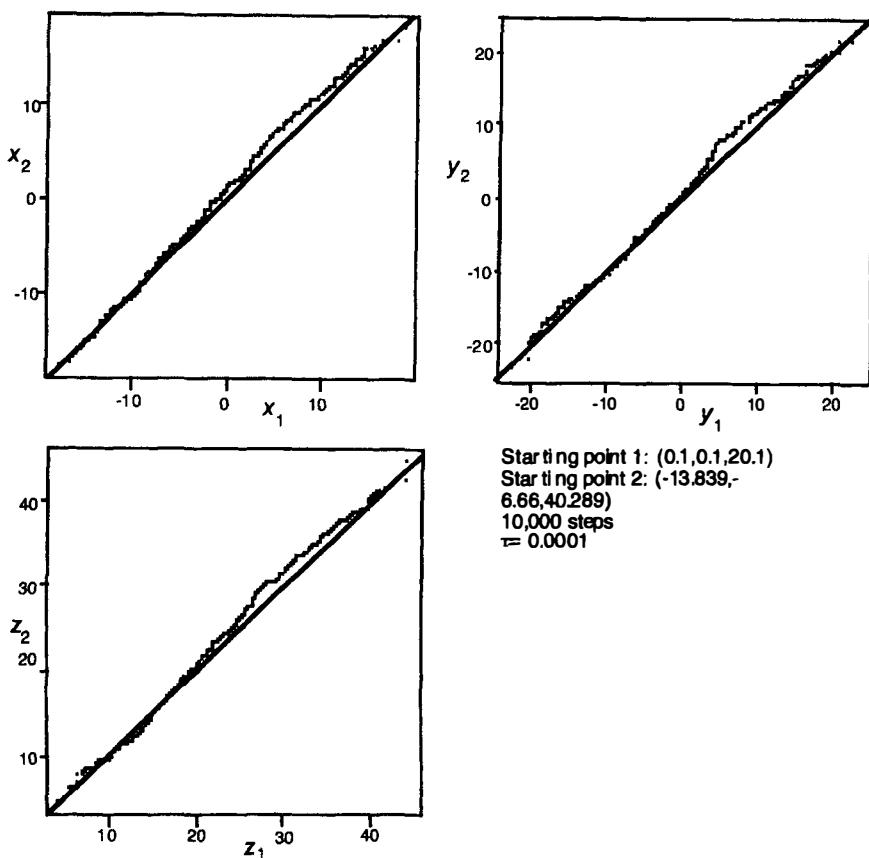


Figure 8.10. Quantile-quantile plots after 10,000 steps.

8.4 Conclusions

It is intriguing that so many scientists have been drawn to a theory which promises such a chaotic world that, if it were the real one, one would not even be able to have a conversation on the merits of the notion. Nevertheless, chaos theory should cause us to rethink aggregation to a supposed mean trace deterministic model. In many cases we will not get in trouble if we go down this route. For example, for most of the conditions one would likely put in a logistic model, one would lie in the smooth part of the curve (i.e., before $a = 1 + \sqrt{6}$). However, as we saw in looking at the logistic model for $a > 1 + \sqrt{6}$, if we consider a realistic model of a population, where the model must constantly be subject to noise, since that is the way the real world operates, then if we aggregate by simulating from the stochastic model as we did in Figure 8.6, we still obtain smooth, sensible results, even after the naive discrete logistic model has failed. But, as we have seen us-

ing two stochastic versions of examples from the work of Lorenz, there are times when the closed form is itself unreliable, even though the simulation-based aggregate, proceeding as it does from the microaxioms, is not. As rapid computing enables us to abandon more and more the closed form, we will undoubtedly find that simulation and stochastic modeling expand our ability to perceive the world as it truly is.

In 1993, Lorenz ([3], pp. 181–184) seriously addressed the question as to whether the flap of a butterfly’s wings in Brazil could cause a tornado in Texas. He reckoned that isolation of the Northern Hemisphere from the Southern in its wind currents might make it impossible, but within the same hemisphere, it was a possibility. In 1997, ([6], p. 360), remarking on his view that the weather is a chaotic phenomenon, Ian Stewart wrote, “Forecasts over a few days, maybe a week—that’s fine. A month? Not a hope!”

If one listens closely to these statements, it is possible to hear the throbbing of the shamans’ tom-toms celebrating the Festival of Unreason. Already, long-term models for the weather have been dramatically improved by noting the importance of driving currents in the jet stream. We now have the capability of forecasting whether winters will be warmer or colder than the norm in, say, New York. And high-frequency hurricane seasons are now being predicted rather well. It is large macroeffects—sunspots, El Niño, currents in the jet stream that drive the weather—not Brazilian butterflies. It is the aggregates which drive the weather. The tiny effects do not matter. In nature, smoothers abound, including noise itself.

Models for the forecasting of the economy have improved as well. Happily, rather than throwing up their hands at the futility of developing good models for real-world phenomena, a number of scientists are constantly drawing upon the scientific method for learning and making forecasts better. One of the bizarre traits of human beings is the tendency of some to believe models as stand-alones, unstressed by data. In Chapter 14 it is noted how this can cause disaster, as in the case of the LTCM investment fund. However, some people refuse to look at models as having consequences for decision making. We noted in Chapter 3 how this has characterized public health policy in the American AIDS epidemic. It is amazing how the same people can often take both positions simultaneously. For example, there are scientists who have worked hard to develop decision-free models for AIDS (this really takes some doing) and also work on trying to show how biological systems should be unstable (evidence to the contrary notwithstanding) as a result of the regions of instability in the discrete logistic model dealt with earlier in this chapter.

As time progresses, it is becoming ever more clear that naturally occurring realizations of mathematical chaos theory are difficult to find. The response of chaoticians is increasingly to include as “chaos” situations in which standard models for describing phenomena are being proved unsatisfactory or at least incomplete (see, e.g., Peters [4]). Thus, “chaos” has

been broadened to include “nonrobust” and “unstable.” That some models are simply wrong and that others claim a completeness that is mistaken I do not question. The point being made in this chapter is that just because one can write down a chaotic model, it need not appear in nature, and that when it does, we will probably view its effects only through the mediating and smoothing action of noise. Moreover, a philosophical orientation that we should give up on forecasting the weather or the economy because somebody has postulated, without validation, a chaotic model, is an unfortunate handshaking between the New Age and pre-Socratic times.

A more complete analysis of chaos and noise is given in Thompson et alia [7] and Thompson and Tapia [8].

Problems

8.1. Let us consider the following birth process amongst single cell organisms. X is the population size, t the time, F the limiting population.

$$\text{Prob(birth in } [t, t + \Delta t) \text{)} = .001X(F - X)\Delta t.$$

Starting with an initial population of 10 at time 0, simulate a number of population tracks using a time increment of .01 and going to a time limit of 20. Take the time averages of 500 such tracks. Do you see evidence of chaotic behavior as you change F ?

8.2. Again, consider the more general model

$$\text{Prob(birth in } [t, t + \Delta t) \text{)} = aX(F - X)\Delta t.$$

Here a is positive, the initial population is 10. Again, you should satisfy yourself that, although there is considerable variation over time tracks, the average time track does not exhibit chaotic behavior.

8.3. Many of the older differential equation models of real world phenomena suffered from the necessity of not easily being able to incorporate the reality of random shocks. For example, let us consider the well known predator-prey model of Volterra ([5, 9]):

$$\begin{aligned} \frac{dx}{dt} &= ax - bxy. \\ \frac{dy}{dt} &= -cy + cxy. \end{aligned}$$

A simple Simpson’s rule discretization works rather well:

$$\begin{aligned} x(t) &= x(t - \Delta) + [ax(t - \Delta) - bx(t - \Delta)y(t - \Delta)]\Delta \\ y(t) &= y(t - \Delta) + [-cy(t - \Delta) + dx(t)y(t - \Delta)]\Delta. \end{aligned}$$

Here, the x are the fish and the y are the sharks. For $x(0) = 10$, $y(0) = 1$, $a = 3$, $b = 2$, $c = 2$, $d = 1.5$, $\Delta = .01$, we show the resulting plot for the first 13 units of time in Figure 8.11.

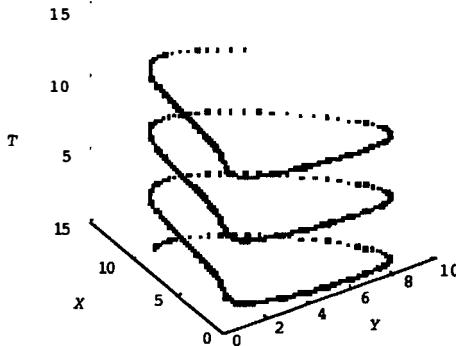


Figure 8.11. Volterra's predator-prey model.

Such regularity is, of course, not supported by the evidence. The system is certainly stable, but not so regular. Show that if you change the discretization slightly to:

$$\begin{aligned} x(t) &= x(t - \Delta)(1 + 10z\Delta) + [ax(t - \Delta) - bx(t - \Delta)y(t - \Delta)]\Delta \\ y(t) &= y(t - \Delta) + [-cy(t - \Delta) + dx(t)y(t - \Delta)]\Delta, \end{aligned}$$

where z is a $\mathcal{N}(0, 1)$ random variable, then the stability of the system is maintained, but the path is much more realistically random.

8.4. One realization of the Hénon attractor [5] is given, in discrete formulation, by

$$\begin{aligned} x_{n+1} &= y_n + 1 - 1.4x_n^2 \\ y_{n+1} &= 0.3x_n. \end{aligned}$$

For the Hénon attractor, as with that of Lorenz, slight differences in the initial conditions produce great differences in the trajectories. Examine this fact and then observe the effect of adding a slight amount of noise to the system.

8.5. Chaos is frequently the concern of those who are worried about a data set somehow “on the edge” of stability. Consider, for example, a sample of size n from the Gaussian $\mathcal{N}(0, 1)$. Then, we know that if we compute the sample mean \bar{X} and variance s^2 , then

$$t = \frac{\bar{X}}{s/\sqrt{n}}$$

has a t distribution with $\nu = n - 1$ degrees of freedom. For $n = 2$, this is the Cauchy distribution which has neither expectation nor variance but is symmetrical about zero. For $n = \infty$, this is $\mathcal{N}(0, 1)$. Let us consider what happens when we sample from $t(\nu)$ adding on observations one at a time and computing the sequential sample mean:

$$\bar{T}_\nu(N) = \frac{(N-1)\bar{T}_\nu(N-1) + t_{\nu,N}}{N}$$

- (a) Give plots of $\bar{T}_1(N)$ for N going from 1 to 5000.
- (b) What happens if you throw away the 10% smallest observations and the 10% largest before computing $\bar{T}_1(N)$ (show this for $N = 10, 50, 100$ and 5000). This “trimming” is generally an easy way for minimizing the effects of long-tailed distributions.
- (c) The Cauchy does not occur very often in nature (although it is easy to design a hardwired version: Just divide a $\mathcal{N}(0, 1)$ signal by another $\mathcal{N}(0, 1)$ signal; but that would not be a good design). Much more realistically, we carry out (a) but for $\nu = 3$.

References

- [1] Lorenz, E. N. (1963). “Deterministic nonperiodic flow,” *J. Atmospheric Sci.*, **20**, 130–141.
- [2] Lorenz, E. N. (1964). “The problem of deducing the climate from the governing equations,” *Tellus*, **16**, 1–11.
- [3] Lorenz, E. N. (1993). *The Essence of Chaos*. London: UCL Press.
- [4] Peters, E. E. (1994). *Fractal Market Analysis: Applying Chaos Theory to Investment and Economics*. New York: John Wiley & Sons.
- [5] Pritchard, J. (1996). *The Chaos Cookbook: A Practical Programming Guide*. Oxford: Butterworth-Heinemann, 94, 166.
- [6] Stewart, I. (1997). *Does God Play Dice?* London: Penguin Press.
- [7] Thompson, J. R., Stivers, D. N., and Ensor, K. B. (1991). “SIMEST: a technique for model aggregation with considerations of chaos,” in *Mathematical Population Dynamics*. O. Arino, D. Axelrod and M. Kimmel, eds. New York: Marcel Dekker, 483–510.
- [8] Thompson, J. R. and Tapia, R. A. (1990). *Nonparametric Function Estimation, Modeling, and Simulation*. Philadelphia: SIAM, 244–252.
- [9] Thompson, J. R. (1989). *Empirical Model Building*. New York: John Wiley & Sons, 71–76.

Chapter 9

A Primer in Bayesian Data Analysis

9.1 Introduction

Within the statistics community, there has for many years a group (we shall call the *Religious Bayesians*) who hold that all analysis should be carried out using *Bayes' Theorem*. That is not the position of this author. Nevertheless, there are many situations where Bayesian analysis is very useful. In this chapter, we give some examples.

Most of the work in this chapter based on one specific data set of Gehan and Freireich [4]. The convergence properties of the estimation procedures considered in this chapter, however, apply rather broadly. Those interested in convergence proofs are referred to Casella and George [1], Cowles and Carlin [2], or Tanner [6].

Let us consider a process in which failures occur according to an exponential distribution, that is,

$$F(t) = 1 - \exp(-\theta t). \quad (9.1)$$

Thus the probability no failure takes place on or before t is given by $\exp(-\theta t)$. Then, based on an independent sample of size n , the likelihood is given by

$$L(\theta) = \prod_{j=1}^n F'(t_j) = \theta^n \prod_{j=1}^n \exp(-\theta t_j) = \theta^n \exp\left(-\theta \sum_{j=1}^n t_j\right). \quad (9.2)$$

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

Then, the maximum likelihood estimator of θ can readily be obtained by taking the logarithm of the likelihood, differentiating with respect to θ , and setting the derivative equal to zero:

$$\log L(\theta) = n \log \theta - \theta \sum_{j=1}^n t_j \quad (9.3)$$

$$\frac{\partial \log L(\theta)}{\partial \theta} = \frac{n}{\theta} - \sum_{j=1}^n t_j = 0. \quad (9.4)$$

This yields

$$\hat{\theta} = \frac{n}{\sum_{j=1}^n t_j}, \quad (9.5)$$

where the $\{t_j\}$ represent the n failure times.

Next, let us consider the case where one of the subjects did not yield an observed failure because the study ended at censoring time T . Then the likelihood becomes

$$\begin{aligned} L(\theta) &= \theta^{n-1} \prod_{j=1}^{n-1} \exp(-\theta t_j) \text{ Prob (subject } n \text{ does not fail by } T) \\ &= \theta^{n-1} \prod_{j=1}^{n-1} \exp(-\theta t_j) \exp(-\theta T). \end{aligned} \quad (9.6)$$

Then the log likelihood becomes

$$\log L(\theta) = (n-1) \log(\theta) - \theta \sum_{j=1}^{n-1} t_j - \theta T, \quad (9.7)$$

yielding

$$\hat{\theta}_1 = \frac{n-1}{\sum_{j=1}^{n-1} t_j + T}. \quad (9.8)$$

9.2 The EM Algorithm

We note that our classical maximum likelihood estimator does not include any guess as to the when the n th subject might have failed had $T = \infty$. Yet, the times that $n-1$ individuals did fail does provide us with information relevant to guessing the n th failure time. Let us assume that our prior feelings as to the true value of θ are very vague: Essentially, we will take any value of θ from 0 to ∞ to be equally likely. Then, again, our log likelihood is given by

$$\log L(\theta) = n \log \theta - \theta \sum_{j=1}^{n-1} t_j - \theta t_n^*. \quad (9.9)$$

Now, the times of the $n - 1$ failures are a matter of fact. The modal value of θ and the time of the unobserved failure time t_n^* are matters of inference. The EM algorithm is an iterative procedure whereby the hypothesized failure time t_n^* can be conjectured and the log likelihood reformulated for another attempt to obtain a new estimate for the modal value $\hat{\theta}$. Using our naive maximum likelihood estimator for θ , we have that the expected value for t_n^* is given by

$$\begin{aligned} t_{n,1}^* &= \frac{\int_T^\infty t_n^* \exp(-\hat{\theta}_1 t_n^*) dt_n^*}{\exp(-\hat{\theta}_1 T)} \\ &= T + \frac{1}{\hat{\theta}_1}. \end{aligned} \quad (9.10)$$

Substituting this value in the log likelihood, we have

$$\log L(\theta) = n \log \theta - \theta \sum_{j=1}^{n-1} t_j - \theta \left(T + \frac{1}{\hat{\theta}_1} \right). \quad (9.11)$$

This gives us

$$\hat{\theta}_2 = \frac{n}{\sum_{j=1}^{n-1} t_j + \left(T + 1/\hat{\theta}_1 \right)}. \quad (9.12)$$

We then obtain a new expected value for the n 'th failure time via

$$t_{n,2}^* = T + \frac{1}{\hat{\theta}_2}; \quad (9.13)$$

and so on.

In fact, for the exponential failure case, we can handle readily the more complex situation where n of the subjects fail at times $\{t_j\}_{j=1}^n$ and k are censored at times $\{T_i\}_{i=1}^m$. The m expectation estimates for $\{t_i\}_{i=1}^m$ are given at the k th step by $\{T_i + 1/\hat{\theta}_{k-1}\}$, and the log likelihood to be maximized is given by

$$n \log \theta - \theta \left[\sum_{j=1}^{n-m} t_j + \sum_{i=1}^m \left(T_i + \frac{1}{\hat{\theta}_{k-1}} \right) \right]. \quad (9.14)$$

Next, let us apply the EM in the analysis of the times of remission of leukemia patients using a new drug and those using an older modality of treatment. The data we use are from a clinical trial designed by Gehan and Freireich [4]. The database has been used by Cox and Oakes [3] as an example of the EM algorithm. Here we use it to examine the EM algorithm, data augmentation, chained data augmentation, and the Gibbs sampler.

Table 9.1. Leukemia Remission Times

Ranked Survival	New Therapy	Old Therapy
1	6*	1
2	6	1
3	6	2
4	6	2
5	7	3
6	9*	4
7	10*	4
8	10	5
9	11*	5
10	13	8
11	16	8
12	17*	8
13	19*	8
14	20*	11
15	22	11
16	23	12
17	25*	12
18	32*	15
19	32*	17
20	34*	22
21	35*	23

In Table 9.1, an asterisk indicates that a patient's status was known to be remission until the (right censored time), at which time the status of the patient became unavailable. There is no pairing of patients in the two groups. We have simply ordered each of the two sets of 21 patients according to time of remission. Using (9.14) recursively to obtain θ for the new treatment, we have the results shown in Table 9.2.

The average survival time using the new therapy is 1/.025, or 40 months. For the old therapy, average survival was only 8.67 months. So, the new therapy seems relatively promising. We note that in our use of the EM algorithm, we have not allowed our experience with the old therapy to influence our analysis of survival times of the new therapy.

We have chosen to explicate the EM algorithm by the use of a relatively simple example. It should be noted that, like the other algorithms we shall explore in this chapter, it performs effectively under rather general conditions. Clearly, the EM algorithm is, in fact, a data augmentation approach. So is the Gibbs sampler. But the name *data augmentation algorithm* is generally reserved for the batch Bayesian augmentation approach covered in the next section.

Table 9.2. Iterations of EM Algorithm.

Iteration	θ
1	0.082569
2	0.041639
3	0.032448
4	0.028814
5	0.027080
6	0.026180
7	0.025693
8	0.025422
9	0.025270
10	0.025184
11	0.025135
12	0.025107
13	0.025091
14	0.025082
15	0.025077
16	0.025074
17	0.025072
18	0.025071
19	0.025070
20	0.025070

9.3 The Data Augmentation Algorithm

The EM procedure, although formally Bayesian, is, when one uses a diffuse prior, as we have done, an algorithm with which non-Bayesians generally feel comfortable. Many Bayesians, however, would be more comfortable with a procedure that gives the user, not simply the mode of a posterior distribution, but an estimate of the posterior distribution itself.

For example, let us suppose that the density function of a random variable X is given by $f(x; \theta)$, or, in Bayesian notation, $f(x|\theta)$. The joint density of a sample of x 's of size n is then given by

$$\prod_{i=1}^n f(x_i|\theta).$$

Generally speaking, we will be interested in making inferences about the parameter θ in the light of a random sample $\{x_j\}_{j=1}^n$. Before we take any observations, we may well have some feelings about the parameter θ . Seldom will these feelings be so strong as to be of the sort, "We know that θ is precisely equal to 150.3741." If we were really so certain, why bother to collect data concerning the random variable X ? It is much more likely that our feelings would be of the sort: "We are quite confident that θ is greater than 100 but less than 250." Expressing our prior feelings in terms

of a prior distribution on the parameter space is not easy, for most people. Perhaps the major difficulty with a Bayesian approach is not on the basis of logic but on that of practicality. We may well have ideas about θ , absent any data. But it is not so easy to express these as a probability density function.

One way out of the difficulty is to require that the prior density function be such that the functional form will be unchanged by the addition of data. That is, the posterior distribution will have the same functional form as that of the prior density.

Let us return to the problem of exponentially distributed failure times. Here, we recall that

$$f(t_1, t_2, \dots, t_n | \theta) = \theta^n \exp\left(-\theta \sum_{j=1}^n t_j\right). \quad (9.15)$$

Suppose that we decide to take as the prior density of θ , absent any data, a gamma density

$$p(\theta) = \frac{e^{-\lambda\theta} \lambda^\alpha \theta^{\alpha-1}}{\Gamma(\alpha)}. \quad (9.16)$$

Then

$$E(\theta) = \alpha/\lambda \quad (9.17)$$

$$\text{Var}(\theta) = \alpha/\lambda^2. \quad (9.18)$$

It is not unreasonable to suppose that we have some notion as to our prior feelings as to the mean and variance of θ . These feelings enable ready guesses as to appropriate values of λ and α . Furthermore, we can then write down the joint density of θ and the failure times as

$$\begin{aligned} p(\theta; t_1, t_2, \dots, t_n) &= p(\theta) f(t_1, t_2, \dots, t_n | \theta) \\ v &= \frac{\theta^{n+\alpha-1} \lambda^\alpha \exp[-\theta(\lambda + \sum t_j)]}{\Gamma(\alpha)}. \end{aligned} \quad (9.19)$$

If we then obtain the marginal density of t_1, t_2, \dots, t_n , $h(t_1, t_2, \dots, t_n)$, for example, we can obtain the posterior density of θ , via

$$p(\theta | t_1, t_2, \dots, t_n) = \frac{p(\theta; t_1, t_2, \dots, t_n)}{h(t_1, t_2, \dots, t_n)}. \quad (9.20)$$

Here we have

$$\begin{aligned} h(t_1, t_2, \dots, t_n) &= \frac{\lambda^\alpha}{\Gamma(\alpha)} \int_0^\infty \theta^{n+\alpha-1} \exp\left[-\theta(\lambda + \sum t_j)\right] d\theta \\ &= \frac{\Gamma(n+\alpha)}{\Gamma(\alpha)} \frac{\lambda^\alpha}{(\lambda + \sum t_j)^{n+\alpha}}. \end{aligned} \quad (9.21)$$

Then, readily, we have

$$p(\theta|t_1, t_2, \dots, t_n) = \frac{(\lambda + \sum t_j)^{n+\alpha}}{\Gamma(\alpha + n)} \theta^{n+\alpha-1} \exp\left[-\theta(\lambda + \sum t_j)\right]. \quad (9.22)$$

But this density is also of the gamma form with

$$\lambda^* = \lambda + \sum t_j$$

and

$$\alpha^* = \alpha + n,$$

that is,

$$p(\theta|t_1, t_2, \dots, t_n) = \frac{e^{-\theta\lambda^*} (\lambda^*)^{\alpha^*} \theta^{\alpha^*-1}}{\Gamma(\alpha^*)}, \quad (9.23)$$

where

$$E(\theta|t_1, t_2, \dots, t_n) = \frac{\alpha^*}{\lambda^*} = \frac{n + \alpha}{\lambda + \sum t_j} \quad (9.24)$$

and

$$\text{Var}(\theta|t_1, t_2, \dots, t_n) = \frac{\alpha^*}{(\lambda^*)^2} = \frac{n + \alpha}{(\lambda + \sum t_j)^2}. \quad (9.25)$$

In the case where some of the failure times were missing as a result of censoring, we could use the EM algorithm to find an improved estimate of the mode of $p(\theta|t_1, t_2, \dots, t_n)$. We shall, however, consider a strategy for estimating the posterior density itself, through obtaining either a knowledge of the function $p(\theta|t_1, t_2, \dots, t_n)$ or a pointwise (in θ) evaluator of $p(\theta|t_1, t_2, \dots, t_n)$.

Let us suppose we observe the failure times t_1, t_2, \dots, t_{n-m} but are missing $\{t_j\}_{j=n-m+1}^n$, since these individuals were lost from the study at times $\{T_j\}_{j=n-m+1}^n$. The *data augmentation algorithm* proceeds as described below:

Data Augmentation Algorithm

1. Sample θ_j from $p(\theta_j|t_1, t_2, \dots, t_n)$.
2. Generate t_{n-m+i}, \dots, t_n from $\theta_j \exp(-\theta_j t_{n-m+i})$ (with the restriction that $t_{n-m+i} \geq T_{n-m+i}$).
3. Repeat Step 2N times.
4. Compute $\bar{T} = 1/N \sum_{i=1}^N \sum_{j=1}^n t_{ji}$.
5. Let $\lambda^* = \lambda + \bar{T}$.
6. Let $\alpha^* = \alpha + n^*$.

7. Then the new iterate for the posterior distribution for θ is given by

$$p(\theta|t_1, t_2, \dots, t_n) = \frac{e^{-\theta\lambda^*}(\lambda^*)^{\alpha^*}\theta^{\alpha^*-1}}{\Gamma(\alpha^*)}.$$

8. Return to Step 1, repeating the cycle M times or until the estimates λ^* and α^* stabilize.

For the first pass through the cycle, we use for $p(\theta|t_1, t_2, \dots, t_n)$, simply the prior density for θ , $p(\theta)$. Typically, N is quite large, say on the order of 1000. Again, typically, we will go through the repeat cycle until the estimates of the posterior distribution for θ stabilize.

We note that, unlike the EM algorithm, there is no maximization step in data augmentation, rather a series of expectations. It is interesting to note that under rather general conditions, data augmentation does stabilize to a “fixed point” (in function space) under expectation iterations.

Let us consider using the data augmentation algorithm on the remission data for the new treatment in Table 9.1. First of all, we need to ask whether there is a reasonable way to obtain the parameters for the gamma prior distribution of θ . In clinical trials, there is generally the assumption that the newer treatment must be assumed to be no better than the old treatment. So, in this case, we might consider obtaining estimates of the two parameters by looking at the data from the older (control) therapy.

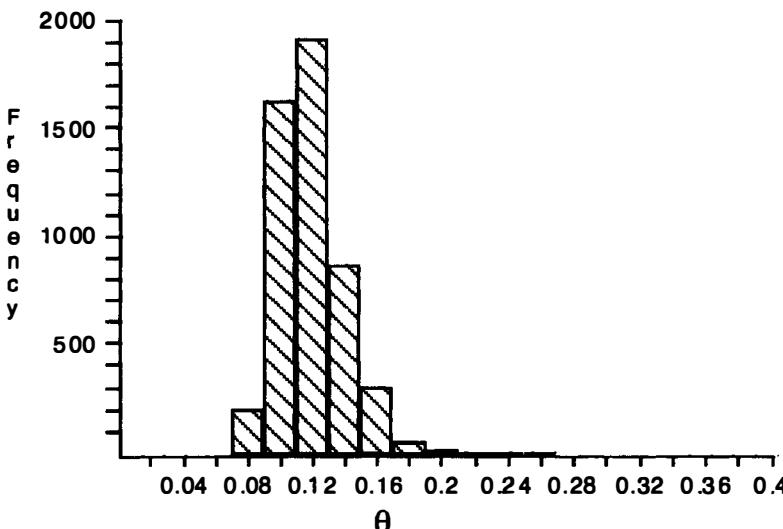


Figure 9.1. Resampled values of θ .

Let us use a simple bootstrap approach to achieve this objective. Sampling with replacement 5000 times from the 21 control remission times (we have used the *Resampling Stats* package [5], but the algorithm can easily be programmed in a few lines of Fortran or C code). Now, for each of the runs, we have computed an estimate for θ , which is simply the reciprocal of the average of the remission times. A histogram of these remission times is given in Figure 9.1.

For the control group, the bootstrap average of θ is .11877 and the variance is .00039018. Then, from (9.17) and (9.18), we have as estimates for λ and α , 304.40 and 36.15, respectively. Returning to the data from the new drug, let us use $N = 1000$.¹ After 400 simulations, we find that

$$\bar{T} = 1/400 \sum_{i=1}^{400} \sum_{j=1}^{12} t_{ji} = 539.75 \quad (9.26)$$

$$\begin{aligned} \alpha^* &= 36.15 + 21 = 57.15 \\ \lambda^* &= 304.4 + 539.75 = 844.15. \end{aligned} \quad (9.27)$$

Computing the posterior mean and variance of θ , we have

$$\begin{aligned} E(\theta|\bar{T}) &= \frac{\alpha^*}{\lambda^*} = .0677; \\ \text{Var}(\theta|\bar{T}) &= \frac{\alpha^*}{\lambda^{*2}} = .0000802. \end{aligned} \quad (9.28)$$

Essentially, we can approximate the posterior distribution of θ as being Gaussian with mean .0677 and variance .0000802.

It seems that waiting for such a large number (1000) of simulations to update our estimates for λ^* and α^* may be somewhat inefficient. Above we have used $N = 1000$ and $M = 400$.

Now, we need not use a large value for N in the generation of plausible t_{n-m+j} . In fact, by using a large value for N , we may be expending a large amount of computing time generating $t_{n-m+1}, t_{n-m+2}, \dots, t_n$ values for θ values generated from posterior distribution estimates which are far from the mark. Let us go through the data augmentation algorithm with $N = 50$ and $M = 10,000$.² This requires approximately the same number of computations as the first run ($N = 1000$ and $M = 400$) run. We now have the estimates for the posterior mean and variance of θ being .06783 and .0000848, respectively. Clearly, the use of a smaller value for N has not changed our estimate for the parameters for the posterior distribution of θ . In Figure 9.2, we show a histogram of the θ values generated for the 10,000

¹The author wishes to thank Patrick King and Mary Calizzi for the computations in the first data augmentation run.

²The author wishes to thank Otto Schwab for the computations and graphs in the balance of this chapter.

runs. On this, we superimpose two Gaussian distributions with variance .0000802, one centered at the mean of the generated θ 's and the other at the median.

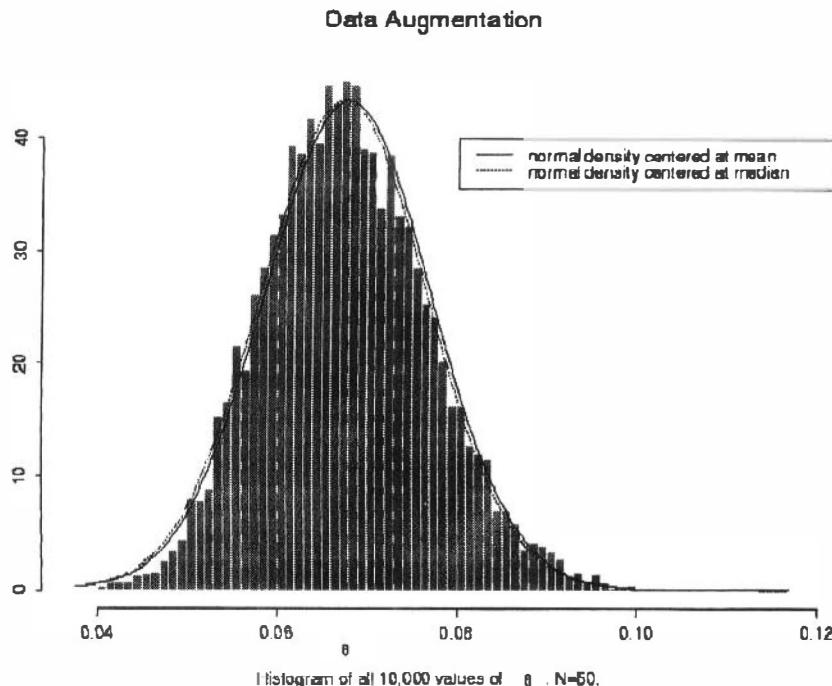


Figure 9.2. Histogram of θ .

It could be argued that we might, in fact, advocate recomputing our guess as to the true posterior distribution after each generation of a set of missing values. This is the *chained data augmentation algorithm*. We note that, now, when we go from Step 1 to Step 8, we “transit” from one value of θ to another. We also observe the one step memory of the process. Knowing $p(\theta_j|t_1, t_2, \dots, t_n)$ is sufficient to generate t_{n-m+1}, \dots, t_n , which, in turn, is sufficient to generate $p(\theta_{j+1}|t_1, t_2, \dots, t_{n_j}^*)$. Thus, the chained data augmentation algorithm is Markovian. And, clearly, the more general data augmentation algorithm is Markovian as well. In the present example, using $N = 1$ (i.e., the chained data augmentation algorithm) took less than 5% of the running time of the data augmentation algorithm with $N = 1000$ and $M = 400$ (with essentially the same results).

At the level of intuition, it would appear that the only reason to use an N greater than 1 would be to guarantee some sort of stability in the estimates. It turns out that this is not necessary, and this fact leads us immediately to

the dominant simulation-based paradigm by which orthodox Bayesians deal with missing values, namely, the *Gibbs Sampler*, which subject we address in the following section.

9.4 The Gibbs Sampler

Next, let us consider the situation where we decide to model failure times according to the normal distribution

$$f(t_1, t_2, \dots, t_n | \mu, h^2) = \left(\frac{h^2}{2\pi} \right)^{n/2} \exp \left[-\frac{1}{2} h^2 \sum_{j=1}^n (t_j - \mu)^2 \right], \quad (9.29)$$

where both μ and $h^2 = 1/\sigma^2$ are unknown. We start with a data set consisting of failure times t_1, t_2, \dots, t_{n-m} but are missing $t_{n-m+1}, t_{n-m+2}, \dots, t_n$. For the most recent estimates for μ and h^2 we shall generate surrogates for the m missing values from a Gaussian distribution, with these estimates for μ and h^2 imposing the restriction that we shall, in the generation of t_{n-m+j} , say, discard a value less than the censoring time T_{n-m+j} .

The natural conjugate prior here is

$$\begin{aligned} & p(\mu, h^2 | M', V', n', \nu') \\ &= K \exp \left[-\frac{1}{2} h^2 n' (\mu - M')^2 \right] \sqrt{h^2} \exp \left[-\frac{1}{2} h^2 V' \nu' \right] (h^2)^{\nu'/2-1}, \end{aligned} \quad (9.30)$$

where K is a constant of integration. (We shall, in the following, use K generically, i.e., one symbol, K will be used for all constants of integration.) Let

$$\begin{aligned} M &= \frac{1}{n} \sum_{i=1}^n t_i \\ \nu &= n - 1 \\ V &= \frac{1}{\nu} \sum_{i=1}^n (t_i - M)^2. \end{aligned}$$

Then the posterior distribution is given by

$$\begin{aligned} & p(\mu, h^2 | M'', V'', n'', \nu'') \\ &= K \exp \left[-\frac{1}{2} h^2 n'' (\mu - M'')^2 \right] \sqrt{h^2} \exp \left[-\frac{1}{2} (h^2) V'' \nu'' \right] (h^2)^{\nu''/2-1}, \end{aligned} \quad (9.31)$$

where

$$n'' = n' + n$$

$$\begin{aligned} M'' &= \frac{1}{n''}(n'M' + nM) \\ \nu'' &= \nu' + (n - 1) + 1 \\ V'' &= \frac{1}{\nu''}[(\nu'V' + n'M'^2) + (\nu V + nM^2 - n''M''^2)]. \end{aligned}$$

Immediately, then, we have the possibility of using the *chained data augmentation algorithm* via

1. Generate (μ, h^2) from $p(\mu, h^2 | M'', V'', n'', \nu'')$.
2. Generate $\{t_{n-m+j}\}_{j=1}^m$ from

$$f(t_{n-m+j} | \mu, h^2) = \left(\frac{h^2}{2\pi} \right)^{1/2} \exp \left[-\frac{h^2}{2} (t_{n-m+j} - \mu)^2 \right],$$

where $t_{n-m+j} > T$.

3. Return to Step 1.

Such a strategy is rather difficult to implement, since it requires the generation of two-dimensional random variables (μ, h^2) . And we can well imagine how bad things can get if the number of parameters is, say, four or more. We have seen in Chapters 1 and 2 that multivariate random number generation can be unwieldy. There is an easy way out of the trap, as it turns out, for according to the *Gibbs sampler* paradigm, we simply generate from the one-dimensional distributions of the parameters sequentially, conditional upon the last generation of the other parameters and the missing value data. Let us see how this is done in the case of the current situation.

The posterior density for h^2 (conditional on the data including the surrogates for the missing failure times) is given by

$$p_{h^2}(h^2 | V'', \nu'') = K \exp \left(-\frac{yV''\nu''}{2} \right) y^{\nu''/2-1}. \quad (9.32)$$

We may then obtain the conditional density for μ given h^2 and the data by

$$\begin{aligned} p_{\mu|h^2}(\mu | h^2, M'', V'', n'', \nu'') &= \frac{p(\mu, h^2 | M'', V'', n'', \nu'')}{p(h^2 | V'', \nu'')} \\ &= K \exp \left[-\frac{1}{2} h^2 n'' (\mu - M'')^2 \right]. \end{aligned} \quad (9.33)$$

Similarly, the posterior density for μ (conditional on the data including the surrogates for the missing failure times) is given by

$$p_{\mu}(\mu | M'', n'', V'', \nu'') = K \left[\nu'' + \frac{(\mu - M'')^2 n''}{V''} \right]^{-(\nu''+1)/2}. \quad (9.34)$$

We may then find the conditional distribution for h^2 given μ and the data by

$$\begin{aligned} p_{h^2|\mu}(h^2|\mu, M'', V'', n'', \nu'') \\ = \frac{p(\mu, h^2 | M'', V'', n'', \nu'')}{p_\mu(\mu | M'', n'', V'', \nu'')} \\ = K(h^2) \frac{\nu''+1}{2} e^{-\frac{1}{2}(n''(\mu - M'')^2 + V''\nu'')h^2}. \end{aligned} \quad (9.35)$$

In summary, the missing failure times are generated from a normal distribution with the current estimates of mean = μ and variance = $1/h^2$. μ , conditional on the generation both of data including missing values for failures and h^2 , is generated from a normal distribution with mean M'' and variance $1/(h^2 n'')$. h^2 , conditional on the data and pseudodata and μ , is a χ^2 variable with $\nu''+1$ degrees of freedom divided by $n''(\mu - M'')^2 + V''\nu''$. Clearly, such one-dimension-at-a-time samplings are extremely easy.

Let us suppose we observe the failure times t_1, t_2, \dots, t_{n-m} but are missing $\{t_j\}_{j=n-m+1}^n$, since these individuals were lost from the study at times $\{T_j\}_{j=n-m+1}^n$. The *Gibbs sampler algorithm* proceeds thusly: At the start, we shall set $M'' = M'$, $V'' = V'$, $n'' = n'$, $V'' = V'$.

1. Generate μ from $p_\mu(\mu | M'', n'', V'', \nu'')$.
2. Generate h^2 from $p_{h^2}(h^2 | V'', \nu'')$.
3. Generate $\{t_{n-m+j}\}_{j=1}^m$ from

$$f(t_{n-m+j} | \mu, h^2) = \left(\frac{h^2}{2\pi} \right)^{1/2} \exp \left[-\frac{h^2}{2} (t_{n-m+j} - \mu)^2 \right].$$

where $t_{n-m+j} > T$.

4. Return to Step 1.

By simply recording all the μ and h^2 (hence σ^2), recorded in sequence of generation, over, say, 10,000 iterations, we can obtain a histogram picture of the joint posterior density for (μ, h^2) .

We can do more. Suppose that

$$\begin{aligned} \bar{M}'' &= \text{Ave}\{M''\} \\ \bar{V}'' &= \text{Ave}\{V''\}. \end{aligned}$$

Then we have

$$p(\mu, h^2 | t_1, t_2, \dots, t_n) \approx p(\mu, h^2 | \bar{M}''', \bar{V}''', n'', \nu''). \quad (9.36)$$

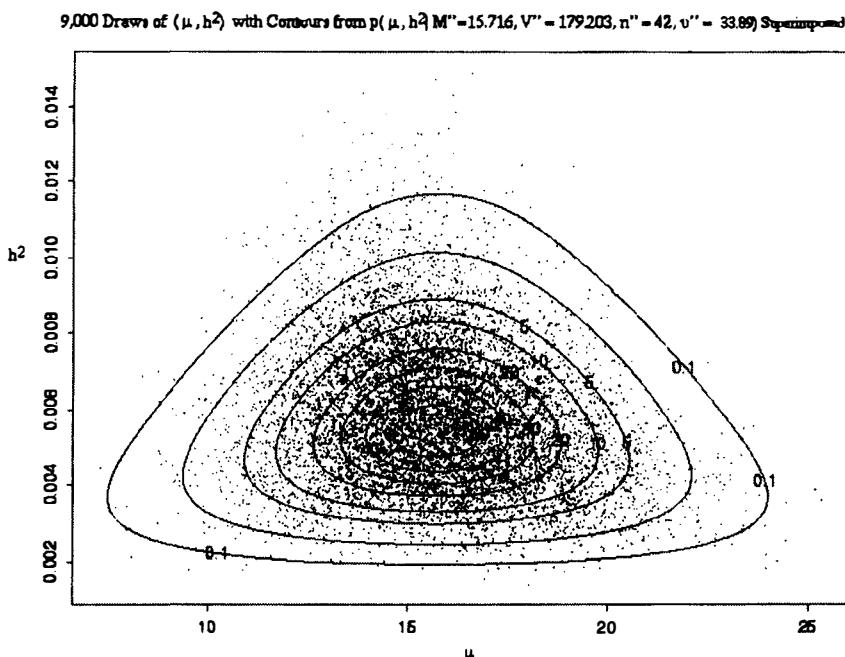


Figure 9.3. Posterior draws of (μ, h^2)

For the Gehan–Freireich data, we have $n = 21$. Using the control group to obtain the initial parameters of the prior, we have $M' = 8.667$, via a bootstrap sample of size 5000. Looking at the reciprocals of the variances of 5000 bootstrap samples of size $n' = 21$ from the control group, we have $E(h) = .028085$ and $\text{Var}(h) = .00012236$. This gives

$$\begin{aligned} V' &= \frac{1}{E(h)} = 35.606 \\ \nu' &= 2 \frac{E(h)^2}{\text{Var}(h)} = 2 \frac{(.028085)^2}{.00012236} = 12.89. \end{aligned}$$

Performing 10,000 samples of μ and h generated one after the other (with the first 1000 discarded to minimize startup effects), we arrive at

$$\begin{aligned} M'' &= 15.716 \\ V'' &= 179.203 \\ n'' &= n' + n = 42 \\ \nu'' &= 12.89 + 21n = 33.89. \end{aligned}$$

The results of these samplings of (μ, h^2) are shown in Figure 9.3. The marginal histograms of μ and h^2 are given in Figures 9.4 and 9.5 respectively.

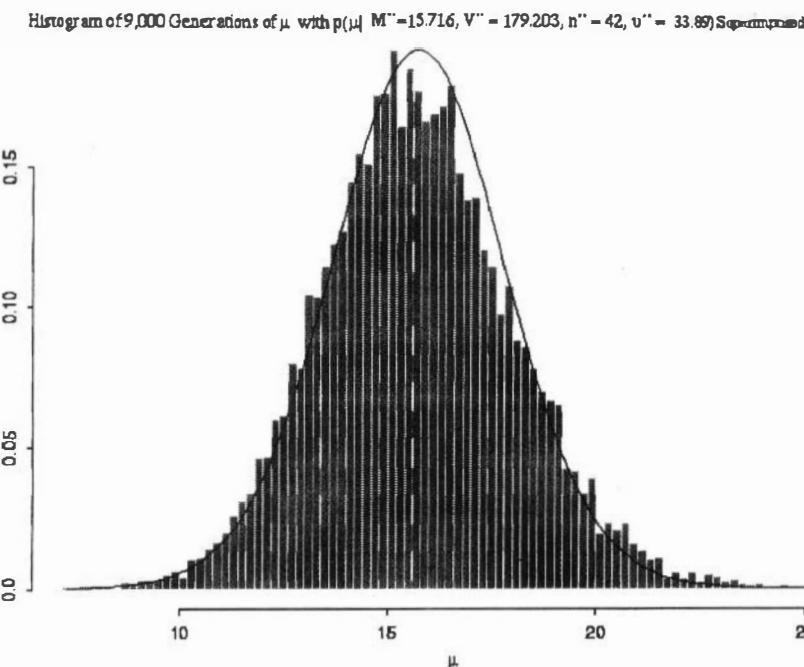


Figure 9.4. Posterior draws of μ .

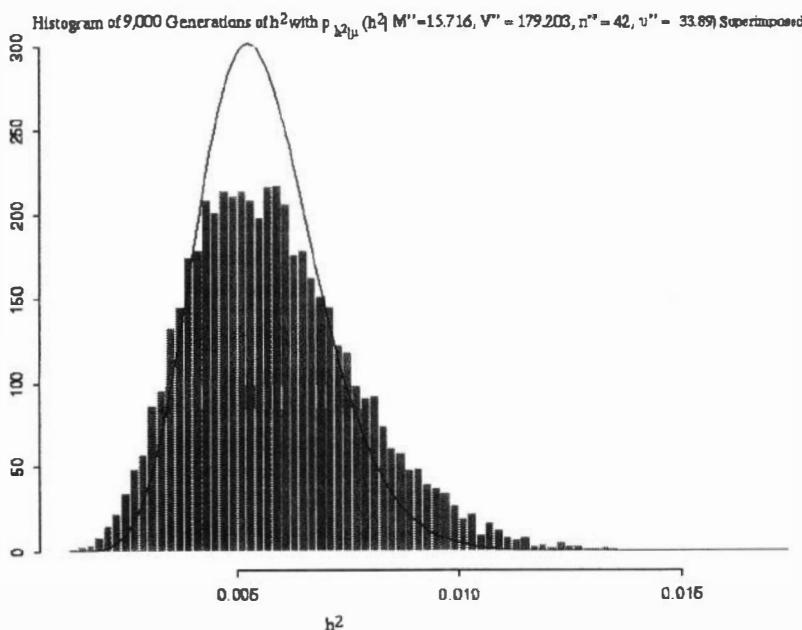


Figure 9.5. Posterior draws of h^2 .

By way of comparison with the earlier data augmentation approach taken in Figure 9.4, in Figure 9.6, we show a histogram taken from the values of 450 but coded as $\theta = 1/\mu$. We note that Figure 9.4 and Figure 9.6 are essentially the same. The average survival posterior mean for the survival time, in both cases is roughly 15.7 months.

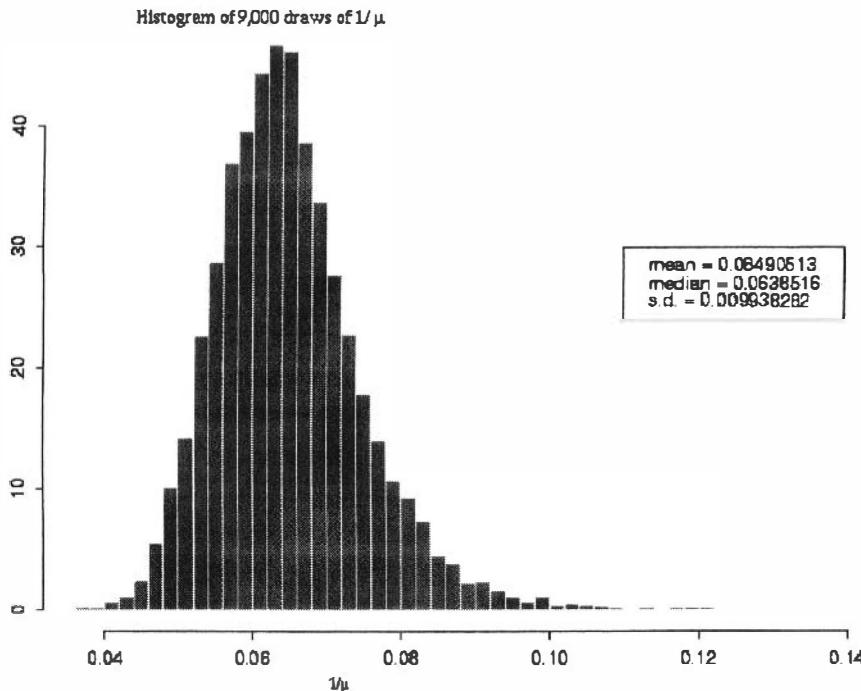


Figure 9.6. Posterior draws of $1/\mu$

9.5 Conclusions

We have noted that the data augmentation and the Gibbs sampler procedures employed both give, for the new treatment, approximately the same posterior mean (15.7 months) for the average survival time. This value is roughly twice as long as the average survival time for the old treatment (8.7 months). But in the case of the EM algorithm, our estimate for the average survival time was a much more optimistic 40 months. What is the reason for the discrepancy?

The discrepancy occurs because in the data augmentation and Gibbs sampler procedures, we elected to use a prior distribution for the parameters of interest which was obtained from data using the old treatment. In the EM algorithm, on the other hand, we used a “noninformative” prior distribution (i.e., one with a very large variance).

Philosophically, some Bayesians would object to our utilization of the results from the old treatment to estimate the parameters of the prior distribution for the new treatment. It could be argued that such a step was unduly pessimistic and smacked of an underlying frequentist mindset whereby the prior distribution was formed with a de facto null hypothesis in mind, namely that the new therapy produced no better average survival times than the old. A Bayesian statistician would probably prefer to consult a panel of advocates of the new treatment and using their insights, attempt to obtain the parameters of the appropriate prior distribution.

In the long run, after the new treatment has been employed on many individuals, it will not make much difference what (nondegenerate) prior distribution we elected to use. But a standard Bayesian claim is that Bayesian techniques enable a more rapid change from a less effective to a more effective procedure. In the case considered here, we have used what Bayesian would consider a rather pessimistic prior. This is unusual in Bayesian analysis. Most Bayesians would tend to use a prior no more pessimistic than the noninformative prior which we used in the case of the EM algorithm. And such a prior is generally roughly equivalent to a standard nonBayesian, frequentist approach. Thus, there is kind of practical bias in favor of the new therapy in most Bayesian analyses. If there is physician opinion indicating the new treatment is much better, then that is incorporated into the prior. But, absent good news, the statistician is supposed to default to a noninformative prior. That may indeed give a running start to those wishing to change the protocol.

In the problems, we give examples of a variety of possible prior assumptions which might be used.

Problems

9.1. Consider the situation where a random variable X has the normal distribution $\mathcal{N}(\theta, 1)$ and θ has the normal distribution $\mathcal{N}(0, 1)$. Create a sample of size 1000 by first sampling a θ and then an X 1000 times.

- (a) Create a two-dimensional histogram of (θ, X) .
- (b) Create a one-dimensional histogram of the marginal density of X .
- (c) Find explicitly the marginal density of X .

9.2. We recall that a Poisson random variable has the probability function

$$P(X) = e^{-\theta} \frac{\theta^X}{X!}.$$

Suppose that θ has the exponential density

$$\begin{aligned} f(\theta) &= e^{-\theta} \text{ for } \theta \geq 0 \\ &= 0 \text{ for } \theta < 0. \end{aligned}$$

Generate first a θ , then an X . Do this 1000 times.

- (a) Show a two-dimensional histogram of (θ, X) .
- (b) Create a one-dimensional histogram of the probability function of X .
- (c) Find the one-dimensional probability function of X .

The time that it takes for a member of a group of small inner city businesses in a particular city to develop credit problems seems to be short. The lending agency is considering making an attempt to increase the time to failure by giving companies free counseling at monthly intervals. Two subgroups of size 25 each have their times until first credit difficulty recorded. The members of the first group (A) do not receive the counseling. The members of the second group (B) do. Below, we show the results of the first two years. The times (in months) till time of first credit problem are given in Table 9.3.

Table 9.3.
Time Until First Credit Problem

Rank	Group A	Group B
1	0	1
2	1	1
3	1	1
4	1	4
5	1	4
6	1	4
7	2	5
8	2	5
9	2	6
10	3	7
11	5	8
12	5	8
13	7	10
14	7	11
15	8	12
16	9	12
17	10	14
18	11	16
19	12	24*
20	14	24*
21	15	24*
22	20	24*
23	24*	24*
24	24*	24*
25	24*	24*

9.3. Assuming that the time to first problem is exponentially distributed, use the EM algorithm to estimate θ_A and θ_B .

9.4. One useful Bayesian approach in deciding whether to make the counseling a standard protocol would be to compute the posterior distributions of θ_A and θ_B . Do this, utilizing the data augmentation algorithm developed in Section 9.3.

9.5. Next, let us consider the situation where we decide to model failure times according to the normal distribution

$$f(t_1, t_2, \dots, t_n | \mu, h^2) = \left(\frac{h^2}{2\pi} \right)^{n/2} \exp \left[-\frac{1}{2} h^2 \sum_{j=1}^n (t_j - \mu)^2 \right],$$

where both μ and $h^2 = 1/\sigma^2$ are unknown. Using the approach in Section 9.4, obtain Gibbs sampler estimates for the posterior distributions of (μ_A, h_A^2) and (μ_B, h_B^2) .

9.6. As has been mentioned earlier in this chapter, the results for finding an estimate for θ in the case of the Gehan–Freireich leukemia data are similar for the data augmentation and the Gibbs sampler procedures. But these results are quite different from those obtained by the use of the EM algorithm. This would appear to be due to the fact that in the case of our data augmentation and Gibbs sampler analyses we have used a prior distribution based on the survival results of the old therapy. Moreover, the sample size of the old therapy was the same as that of the new therapy. We noted that the bootstrap estimate for θ based on the old therapy was .11877. The EM estimate for θ was .02507. The average of these two estimates is .07192, a figure which is roughly similar to the estimate we obtained with data augmentation, .0677. It could be argued, therefore, that by using the old procedure to obtain a prior density for θ , we have blended the good results of the new therapy with the poorer results of the old therapy.

We recall from equation (9.22) that the posterior distribution of θ is given by

$$p(\theta | t_1, t_2, \dots, t_n) = \frac{e^{-\theta\lambda^*} (\lambda^*)^{\alpha^*} \theta^{\alpha^*-1}}{\Gamma(\alpha^*)},$$

where

$$E(\theta | t_1, t_2, \dots, t_n) = \frac{\alpha^*}{\lambda^*} = \frac{n + \alpha}{\lambda + \sum t_j}$$

and

$$\text{Var}(\theta | t_1, t_2, \dots, t_n) = \frac{\alpha^*}{(\lambda^*)^2} = \frac{n + \alpha}{(\lambda + \sum t_j)^2}.$$

Suppose that we decide to use our resampled value of the prior distribution's mean using the data from the old therapy (.11877), but decide that we wish to increase dramatically the estimate for the variance of the prior beyond the old therapy resampled value (.00390). Go through the data augmentation algorithm using higher and higher values for the variance of the prior distribution for θ and see what the data augmentation posterior means for θ are. One might expect that as the variance of the prior increases without limit, the data augmentation and Gibbs sampler results

would look much more like those obtained by the EM algorithm. Is this the case?

References

- [1] Casella, G. and George, E. I. (1992). "Explaining the Gibbs Sampler," *Amer. Statist.*, 167–174.
- [2] Cowles, M. K. and Carlin, B. P. (1996). "Markov chain Monte Carlo convergence diagnostics: a comparative review," *J. Amer. Statist. Assoc.*, 883–904.
- [3] Cox, D. R., and Oakes, D. (1984). *Analysis of Survival Data*, 8, London: Chapman & Hall, 166–168.
- [4] Gehan, E. A. and Freireich, J. (1965). "A generalized Wilcoxon test for comparing arbitrarily single-censored samples," *Biometrika*, 203–23.
- [5] Simon, J. L. (1990). *Resampling Stats*. Arlington, Va.: Resampling Stats, Inc.
- [6] Tanner, M. A. (1993). *Tools for Statistical Inference*. New York: Springer-Verlag.

Chapter 10

Multivariate and Robust Procedures in Statistical Process Control

10.1 Introduction

In terms of its economic impact, statistical process control (SPC) is among the most important topics in modern statistics. Although some statisticians (see, e.g., Banks [2]) have considered SPC to be trivial and of scant importance, the market seems to have reacted quite differently. For example, the American Society for Quality is vastly larger in its membership than the American Statistical Association. It is clear is that SPC is not going away, even should many professional statisticians continue in their disdain for it.

At the end of the World War II, Japan was renowned for shoddy goods produced by automatons living in standards of wretchedness and resignation. W. Edwards Deming began preaching the paradigm of statistical process control (originally advocated by Walter Shewhart) in Japan in the early 1950s. By the mid 1960s, Japan was a serious player in electronics and automobiles. By the 1980s, Japan had taken a dominant position in consumer electronics and, absent tariffs, automobiles. Even in the most sophisticated areas of production, such as computing, the Japanese had achieved a leadership role. The current situation of the Japanese workers is among the best in the world. A miracle, to be sure, and one far beyond that of, say postwar Germany, which was a serious contender in all levels of production before World War II.

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

There is little doubt that the SPC paradigm facilitated these significant changes in Japanese production. Nevertheless, SPC is based on some very basic notions:

- The key to optimizing the output of a system is the optimization of the system itself.
- Although the problem of modifying the output of a system is frequently one of linear feedback (easy), the problem of optimizing the system itself is one of nonlinear feedback (hard).
- The suboptimalities of a system are frequently caused by a small number of assignable causes. These manifest themselves by intermittent departures of the output from the overall output averages.
- Hence, it is appropriate to dispense with complex methods of system optimization and replace these by human intervention whenever one of these departures is noted.
- Once an assignable cause of suboptimality has been removed, it seldom recurs.
- Thus, we have the indication of an apparently unsophisticated but, in fact, incredibly effective, paradigm of system optimization.

Perhaps there is a valid comparison between Shewhart and Adam Smith, who had perceived the power of the free market. But there appears to be no single implementer of the free market who was as important in validating *The Wealth of Nations* as Deming has been in validating the paradigm of statistical process control. There has never been, in world history, so large scale an experiment to validate a scientific hypothesis as Deming's Japanese validation and extension of the statistical process control paradigm of Deming and Shewhart.

It is not our intention to dwell on the philosophy of SPC. That topic has been extensively dealt with elsewhere (see, e.g., Thompson and Koronacki [9]). We will develop here a modeling framework for SPC and then indicate natural areas for exploration. Both Shewhart and Deming held doctorates in mathematical physics, so it is reasonable to assume that there was some reason they did not resort to exotic mathematical control theory type strategies. In Figure 10.1 we indicate a standard feedback diagram for achieving the desired output of a system.

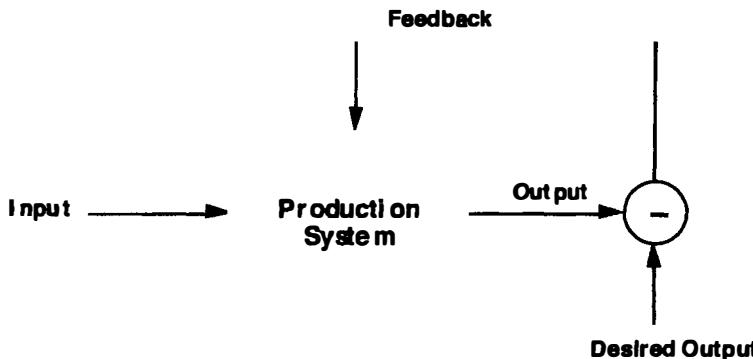


Figure 10.1. Control of output.

One might pose this as an optimization problem where we desire to minimize, say

$$\int_0^T [\text{output}(t) - \text{target}(t)]^2 dt. \quad (10.1)$$

Such a problem is generally linear and tractable.

However, the task of SPC is not to optimize the output directly, but rather, to achieve optimization of the system itself. Generally, such an optimization is an ill-posed problem. At some future time, artificial intelligence and expert systems may bring us to a point where such problems can be handled, to a large extent, automatically. But Shewhart and Deming lacked such software/hardware (as we all still lack it). So they resorted to a piecewise (in time) control strategy based on human intervention (cf. [3] and [5]).

10.2 A Contamination Model for SPC

It is paradoxical that W. Edwards Deming, one of the most important statistical figures of all time, never really published a model of his paradigm of statistical process control (SPC). Deming rightly argued that the key to quality control of an industrial product was to understand the system that produced it. But, in the case of the SPC paradigm itself, he was rather didactic, like the Zen masters of Japan, the country whose economy and standard of living he did so much to improve. Careful analysis of Deming's paradigm led Thompson and Koronacki to their model-based analysis of the SPC system [9].

To understand one of the key aspects of SPC, let us first of all assume that there is a "best-of all-possible -worlds" mechanism at the heart of the process. For example, if we are turning out bolts of 10-cm diameter, we can assume that there will be, in any lot of measurements of diameters, a variable, say X_0 , with mean 10 and a variance equal to an acceptably small number. When we actually observe a diameter, however, we may not

be seeing only X_0 but a sum of X_0 plus some other variables which are a consequence of flaws in the production process. These are not simply measurement errors but actual parts of the total diameter measurements which depart from the “best of all possible worlds” distribution of diameter as a consequence of imperfections in the production process. One of these imperfections might be excessive lubricant temperature, another bearing vibration, another nonstandard raw materials, and so on. These add-on variables will generally be intermittent in time. This intermittency enables us to find measurements which appear to show “contamination” of the basic production process. We note how different the situation would be without the intermittency, if, say, an output variable were the sum of the “best of all possible worlds” variable X_0 and an “out of control” variable X_1 . Then, assuming both variables were Gaussian, the output variable would simply have the distribution $\mathcal{N}(\mu_0 + \mu_1, \sigma_0^2 + \sigma_1^2)$, and the SPC control charts *would not work*. Perhaps the greatest statistical contribution of Shewhart was noting the general presence of intermittent contamination in out-of-control systems.

It is important to remember that the Deming–Shewhart paradigm of SPC is *not* oriented toward the detection of faulty lots. Rather, SPC seeks for atypical lots to be used as an indication of epochs when the system exhibits possibly correctable suboptimalities. If we miss a bad lot or even many bad lots, that is not a serious matter from the standpoint of SPC. If we are dealing with a system that does not produce a sufficiently low proportion of defectives, we should use 100% end inspection, as we note from the following argument frequently referred to as Deming’s Theorem:

- Let a be the cost of passing a bad item.
- Let b be the cost of inspecting an item.
- Let x be the proportion of items inspected.
- Let y be the proportion of bad items.
- Let N be the number of items produced.

Then the cost of inspecting some items and not inspecting others is given by

$$\begin{aligned} C &= bxN + ay(1 - x)N \\ &= (b - ay)xN + ayN. \end{aligned}$$

Clearly, then, if $ay > b$, we should inspect all the lots. (If $ay < b$, we should inspect none.) Thus, from a sampling end product cost model, we should sample all or none.

End-product inspection is really not SPC but *quality assurance*. Most companies use some sort of quality assurance. SPC, however, is different from quality assurance. In fact, the experience of Thompson and Koronacki

when implementing SPC in a number of factories in Poland was that it was better to leave the quality assurance (a.k.a. quality control) groups in place and simply build up new SPC departments.

Now, for SPC, we simply cannot have the alarms constantly ringing, or we shall be wasting our time with false alarms. Accordingly, in SPC, we should be interested in keeping the probability of a Type I error small. Thus, the testing rules in control charting are typically of the type

$$P[\text{declaring "out of control" | in control}] = 0.002. \quad (10.2)$$

With such conservatism we may well find an out-of-control situation later rather than sooner. However, we shall tend to avoid situations where the alarms are always ringing, frequently to no good purpose. And by the theory of SPC, suboptimalities, if missed, will occur in the same mode again.

Proceeding with the contamination model, let us assume that the random variables are added. In any lot, indexed by the time t of sampling, we will assume that the measured variable can be written as

$$Y(t) = X_0 + \sum_{i=1}^k I_i(t) X_i, \quad (10.3)$$

where X_i comes from distribution F_i having mean μ_i and variance σ_i^2 , $i = 1, 2, \dots, k$, and indicator

$$\begin{aligned} I_i(t) &= 1 \text{ with probability } p_i \\ &= 0 \text{ with probability } 1 - p_i. \end{aligned} \quad (10.4)$$

When such a model is appropriate, then, with k assignable causes, there may be in any lot, 2^k possible combinations of random variables contributing to Y . Not only do we assume that the observations within a lot are independent and identically distributed, we assume that there is sufficient temporal separation from lot to lot that the parameters driving the Y process are independent from lot to lot. Also, we assume that an indicator variable I_i maintains its value (0 or 1) throughout a lot. Let \mathcal{I} be a collection from $i \in 1, 2, \dots, k$. Then

$$Y(t) = X_0 + \sum_{i \in \mathcal{I}} X_i \text{ with probability } \left(\prod_{i \in \mathcal{I}} p_i \right) \left(\prod_{i \in \mathcal{I}^c} (1 - p_i) \right). \quad (10.5)$$

Restricting ourselves to the case where each distribution is Gaussian (normal), the observed variable $Y(t)$ is given by

$$\begin{aligned} Y(t) &= \mathcal{N} \left(\mu_0 + \sum_{i \in \mathcal{I}} \mu_i, \sigma_0^2 + \sum_{i \in \mathcal{I}} \sigma_i^2 \right), \\ &\text{with probability } \left(\prod_{i \in \mathcal{I}} p_i \right) \left(\prod_{i \in \mathcal{I}^c} (1 - p_i) \right). \end{aligned} \quad (10.6)$$

Moreover, it is a straightforward matter to show that

$$E(Y) = \mu_0 + \sum_{i=1}^k p_i \mu_i \quad (10.7)$$

$$\text{Var}(Y) = \sigma_0^2 + \sum_{i=1}^k p_i \sigma_i^2 + \sum_{i=1}^k p_i (1 - p_i) \mu_i^2.$$

A major function of the Shewhart control chart is to find epochs of time which give lots showing characteristics different from those of the "in control" distribution. We note that no assumption is made that this dominant distribution necessarily conforms to any predetermined standards or tolerances. Deming proposes that we find estimates of the dominant μ_0 and σ_0^2 and then find times where lots significantly depart from the dominant. Personal examination of what was unusual about the system when the unusual lot was observed enables us to search for the "assignable cause" of the trouble and fix it. Not a particularly elegant way to proceed perhaps, but plausible *prima facie* and proved amazingly effective by experience.

Consider the flowchart of a production process in Figure 10.2. (For reasons of simplicity, we shall neglect effects of time delays in the flow).

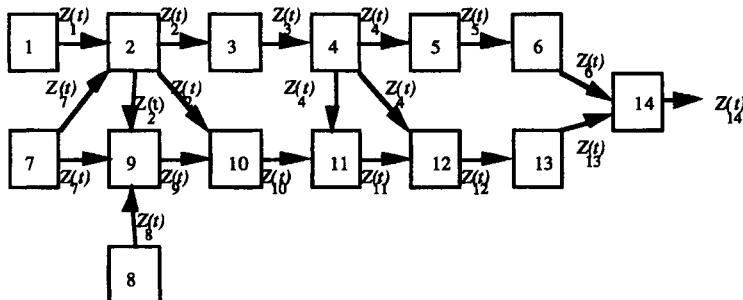


Figure 10.2. Simple flowchart of production.

When Deming writes of statistical process control imparting "profound knowledge," he is not resorting to hype or boosterism. On the contrary, this "profound knowledge" is hardnosed and technical. At a very early stage in the optimization process of SPC, we are urged to draw a flowchart of the production process. In many cases, that very basic act (i.e., the composition of the flowchart) is the single most valuable part of the SPC paradigm. Those who are not familiar with real industrial situations might naively assume that a flowchart is composed long before the factory is built and the production begins. Unhappily, such is not the case. To a large extent, the much maligned ISO 9000 protocol for selling goods in the European Economic Community is simply the requirement that a manufacturer write down a flowchart of his or her production process.

The SPC flowchart continually monitors the output of each module and seeks to find atypical outputs at points in time that can be tracked to a

particular module. This module is then considered as a candidate for immediate examination for possible suboptimalities, which can then be corrected. The SPC flowchart approach will respond rather quickly to suboptimalities. Let us consider an example.

First of all, let us suppose that in Figure 10.2, the output of module 2 is an aqueous solution where the key measure of Z_2 is the exiting concentration of compound *A* from module 2. We note that Z_9 is the measured strength of compound *B* from module 9. In module 10, a compound *AB* is produced, and the output variable Z_{10} is the strength of that by-product compound. Let us suppose that the end product Z_{14} is the measured reflectivity of a strip of metal. Clearly, the system described above may be one of great complexity. A primitive quality assurance paradigm would simply examine lots of the end product by looking at lot averages of Z_{14} and discarding or reworking lots which do not meet specification standards. If the output is our only measured variable, then any notion of correcting problems upstream of Z_{14} is likely to be attempted by a uniform harangue of the personnel concerned with the management of each of the modules "to do better" without any clue as to how this might be done.

The philosophy of Deming's SPC suggests that we would do much better to find the source of the defect in the system so that it can be rectified. This will be achieved by monitoring each of the intermediate output values Z_1, Z_2, \dots, Z_{14} . Simply looking at Z_{14} will not give us an indication that, say, there is a problem with the control of Z_{10} . This is an example of the truism that "you cannot examine quality into a system."

Let us recall the *Maxim of Pareto*: ¹ *The failures in a system are usually the consequence of a few assignable causes rather than the consequence of a general malaise across the system.* Suppose that we make the following further extension of this Maxim: *the failures in a module are usually bunched together in relatively short time epochs, where contamination intervenes, rather than being uniformly distributed across the time axis.* Thus, we are postulating that there will be periods where malfunctioning in a flawed module will be particularly prominent. Statistical process control gives us a simple means for searching for atypical epochs in the record of observations. Whenever we find such atypicality, we will attempt to examine the functioning of the module closely in the hope that we can find the problem and fix it.

Let us return to Figure 10.2. Suppose that we find an atypical epoch of Z_{10} . Since the effect of Z_{10} flows downstream to modules 11 through 14, it may well be that a glitch in Z_{10} will cause glitches in some or all of these as

¹ Interestingly, we may look on Pareto's maxim in the light of Bayes' axiom (postulate 1). If we have a discrete number of causes of failure, Bayes' axiom suggests that we put equal prior probability on each cause. Pareto's maxim (which, although not explicitly stated in his works is fairly deemed to be consistent with them) tells us that it is most likely the probabilities will actually be skewed rather than uniform. The differences represented by the two postulates are consistent with the different philosophies of Bayes and Pareto, the first optimistic and democratic, the second pessimistic and oligarchic.

well. However, our best course will be to find the earliest of the modules in a series where the glitch occurs, since that is the one where the assignable cause is most likely to be found (and fixed).

In the example above, let us suppose that we find no atypicality from lot to lot until we get to module 10. Then we also find atypicality in modules 11 through 14. It seems rather clear that we need to examine module 10 for an assignable cause of the system behaving suboptimally. Once the assignable cause is found, it can generally be fixed. Once fixed, it will not soon recur.

10.3 A Compound Test for SPC Data in Higher Dimensions

The basic control chart procedure of Deming is not oriented toward seeing whether a particular lot of items is within predetermined limits, but rather whether the lot is typical of the dominant distribution of items produced. In the one-dimensional case, the interval in which we take a lot sample mean to be “typical,” and hence the production process to be “in control” is given by

$$\bar{x} - 3 \frac{\hat{\sigma}}{\sqrt{n}} \leq \bar{x} \leq \bar{x} + 3 \frac{\hat{\sigma}}{\sqrt{n}}. \quad (10.8)$$

where n is the lot size, \bar{x} is the mean of a lot, $\hat{\sigma}$ is an estimator for the standard deviation of the dominant population of items, and \bar{x} is an estimator for the mean of the dominant population. Assuming that \bar{x} is normally distributed, then the probability that a lot of items coming from the dominant (i.e., “typical”) population will fall outside the interval is roughly 0.002. Generally speaking, because we will usually have plenty of lots, taking the sample variance for each lot, and taking the average of these will give us, essentially, σ^2 .

Now let us go from the one-dimensional to the multivariate situation. Following Thompson and Koronacki [9], let us assume that the dominant distribution of output \mathbf{x} data is p -variate normal, that is,

$$f(\mathbf{x}) = |2\pi\Sigma|^{-1/2} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right], \quad (10.9)$$

where $\boldsymbol{\mu}$ is a constant vector and Σ is a constant positive definite matrix. By analogy with the use of control charts to find a change in the distribution of the output and/or the input of a module, we can describe a likely scenario of a process going out of control as the mean suddenly changes from, say, $\boldsymbol{\mu}_0$ to some other value. Let us suppose that for j th of N lots, the sample mean is given by $\bar{\mathbf{x}}_j$ and the sample covariance matrix by \mathbf{S}_j . Then the

natural estimates for μ_0 and Σ are

$$\bar{\bar{\mathbf{x}}} = \frac{1}{N} \sum_{j=1}^N \bar{\mathbf{x}}_j$$

and

$$\bar{\mathbf{S}} = \frac{1}{N} \sum_{j=1}^N \mathbf{S}_j,$$

respectively. Now the Hotelling T^2 -like statistic for the j th lot assumes the form

$$T_j^2 = n(\bar{\mathbf{x}}_j - \bar{\bar{\mathbf{x}}})' \bar{\mathbf{S}}^{-1} (\bar{\mathbf{x}}_j - \bar{\bar{\mathbf{x}}}), \quad (10.10)$$

where $j = 1, 2, \dots, N$. Alt has shown [1] that

$$\frac{nN - N - p + 1}{p(p-1)(N-1)} T_j^2$$

has the F distribution with p and $nN - N - p + 1$ degrees of freedom. Thus, we consider the j th lot to be out of control if

$$T_j^2 > \frac{p(p-1)(N-1)}{nN - N - p + 1} F_{p, nN - N - p + 1}(\alpha), \quad (10.11)$$

where $F_{p, nN - N - p + 1}(\alpha)$ is the upper (100α) th percentile of the F distribution with p and $nN - N - p + 1$ degrees of freedom. In SPC, it is generally a fair assumption that N is large, so that we can declare the j th lot to be out of control if

$$T_j^2 > \chi_p^2(\alpha), \quad (10.12)$$

where $\chi_p^2(\alpha)$ is the upper (100α) th percentile of the χ^2 distribution with p degrees of freedom.

The dispersion matrix (i.e., the covariance matrix of the set of estimates $\hat{\mu}$) is given by

$$\mathbf{V}_{(p \times p)} = \begin{bmatrix} \text{Var}(\hat{\mu}_1) & \text{Cov}(\hat{\mu}_1, \hat{\mu}_2) & \dots & \text{Cov}(\hat{\mu}_1, \hat{\mu}_p) \\ \text{Cov}(\hat{\mu}_1, \hat{\mu}_2) & \text{Var}(\hat{\mu}_2) & \dots & \text{Cov}(\hat{\mu}_2, \hat{\mu}_p) \\ \vdots & \vdots & & \vdots \\ \text{Cov}(\hat{\mu}_1, \hat{\mu}_p) & \text{Cov}(\hat{\mu}_2, \hat{\mu}_p) & \dots & \text{Var}(\hat{\mu}_p) \end{bmatrix}, \quad (10.13)$$

where

$$\hat{\mu}_j = \bar{x}_j = \frac{1}{n} \sum_{i=1}^n x_{ij} \quad (10.14)$$

for each j .

Let us investigate the power (probability of rejection of the null hypothesis) of the T_j^2 test as a function of the noncentrality:

$$\lambda = (\mu - \mu_0)' \mathbf{V}_{(p \times p)}^{-1} (\mu - \mu_0). \quad (10.15)$$

One can use the approximation [8] for the power of the T_j^2 test:

$$P(\lambda) = \int_{[(p+\lambda)/(p+2\lambda)]\chi_{\alpha}^2(p)}^{\infty} d\chi^2 \left(p + \frac{\lambda^2}{p+2\lambda} \right), \quad (10.16)$$

where $d\chi^2(p)$ is the differential of the cumulative distribution function of the central χ^2 distribution with p degrees of freedom and $\chi_{\alpha}^2(p)$ its $100(1-\alpha)\%$ point. Now, in the current application, any attempt at a numerical approximation technique is unwieldy, due to the fact that we shall be advocating a multivariate strategy based on a battery of nonindependent test. Here is just one of the myriad of instances in real-world applications where simulation can be used to provide quickly and painlessly to the user an excellent estimate of what we need to know at the modest cost of a few minutes crunching on a modern desktop computer.

Current practice is for virtually all testing of p -dimensional data to be carried out by a battery of p one-dimensional tests. Practitioners rightly feel that if glitches are predominant in one or another of the p dimensions, then the information in the p -dimensional multivariate Hotelling statistic will tend to be obscured by the inclusion of channels that are in control.

As an example, let us consider the case where $p = 5$. Thompson and Koronacki [9] have proposed the following compound test:

1. Perform the five one-dimensional tests at nominal Type I error of $\alpha = 0.002$ each.
2. Next, perform the ten two-dimensional tests at nominal $\alpha = 0.002$ for each.
3. Then perform the ten three-dimensional tests.
4. Then perform the five four-dimensional tests
5. Finally, perform the one five-dimensional test.

If all the tests were independent, we would expect a pooled Type I error of

$$\alpha = 1 - (1 - 0.002)^{31} = 0.06. \quad (10.17)$$

Table 10.11. Type I Errors of Compound Test.

p	1-d Tests	$n = 5$	$n = 10$	$n = 15$	$n = 20$	$n = 50$	$n = 100$	$n = 200$
2	.004	.00434	.00466	.00494	.00508	.00512	.00528	.00546
3	.006	.00756	.00722	.00720	.00720	.00704	.00826	.00802
4	.008	.01126	.01072	.01098	.01098	.01108	.01136	.01050
5	.010	.01536	.01582	.01552	.01544	.01706	.01670	.01728

However, the ten tests are not really independent [so we cannot use (10.16)]. For dimensions two through five, using uncorrelated vector components, we show in Table 10.1, the Type I errors based on 50,000 simulations per table

entry. In the second column, we show the Type I errors if only the one-dimensional tests are carried out (i.e., $p \times .002$). The subsequent columns give the Type I errors for various lot sizes (n) assuming we use all the possible $2^p - 1$ tests. The resulting Type I errors for the pooled tests for all dimensions are not much higher than those obtained simply by using only the one-dimensional tests. We note that in the five-dimensional case, if we use only the five one-dimensional tests, we have a Type I error of .01. Adding in all the two-dimensional, three-dimensional, four-dimensional and five-dimensional tests does not even double the Type I error. As a practical matter, a user can, if desired, simply continue to use the one-dimensional tests for action, reserving the compound higher-dimensional tests for exploratory purposes.

We note here how the use of simulation, essentially as a “desk accessory,” enabled us quickly to determine the downside risks of using a new testing strategy. Slogging through analytical evaluations of the compound test would have been a formidable task indeed. Using a simulation-based evaluation, we were able quickly to see that the price for the compound test was small enough that it should be seriously considered.

10.4 Rank Testing with Higher-Dimensional SPC Data

In statistical process control, we are looking for a difference in the distribution of a new lot, anything out of the ordinary. That might seem to indicate a nonparametric density estimation based procedure. But the general ability to look at averages in statistical process control indicates that for many situations, the central limit theorem enables us to use procedures that point to distributions somewhat close to the normal distribution as the standards. In the case where data are truly normal, the functional form of the underlying density can be based exclusively on the mean vector and the covariance matrix. However, as we show below, it is a rather easy matter to create multivariate tests that perform well in the normal case and in heavy tailed departures from normality.

Consider the case where we have a base sample of N lots, each of size n , with the dimensionality of the data being given by p . For each of these lots, compute the sample mean $\bar{\mathbf{X}}_i$ and sample covariance matrix \mathbf{S}_i . Moving on, compute the average of these N sample means, $\bar{\bar{\mathbf{X}}}$, and the average of the sample covariance matrices $\bar{\bar{\mathbf{S}}}$. Then, use the transformation

$$\mathbf{Z} = \bar{\bar{\mathbf{S}}}^{-1/2}(\mathbf{X} - \bar{\bar{\mathbf{X}}}), \quad (10.18)$$

which transforms \mathbf{X} into a variate with approximate mean $\mathbf{0}$ and approximate covariance matrix \mathbf{I} . Next, apply this transformation to each of the N lot means in the base sample. For each of the transformed lots, compute

the transformed mean and covariance matrix, $\bar{\mathbf{Z}}_i$ and $\mathbf{S}_{\bar{\mathbf{Z}}_i}$, respectively. For each of these, apply, respectively, the location norm

$$\|\bar{\mathbf{Z}}_i\|^2 = \sum_{j=1}^p \bar{Z}_{i,j}^2, \quad (10.19)$$

and the scale norm

$$\|\mathbf{S}_i\|^2 = \sum_{j=1}^p \sum_{l=1}^p S_{i,j,l}^2. \quad (10.20)$$

Now, if a new lot has location norm higher than any of those in the base sample, we flag it as atypical. If its scale norm is greater than those of any lot in the base sample, we flag it as atypical. The Type I error of either test is given, approximately, by $1/(N + 1)$, that of the combined test is given very closely by

$$1 - \left(1 - \frac{1}{N + 1}\right)^2 = \frac{2N + 1}{(N + 1)^2}. \quad (10.21)$$

Let us now compare the performance of the location rank test with that of the parametric likelihood ratio test when we have as the generator of the “in control” data a p -variate normal distribution with mean $\mathbf{0}$ and covariance matrix \mathbf{I} , the identity. We consider as alternatives “slipped” normal distributions, each with covariance matrix \mathbf{I} but with a translated mean each of whose components is equal to the “slippage” μ . In Figure 10.3, using 20,000 simulations of 50 lots of size 5 per slippage value to obtain the base information, we compute the efficiency of the rank test to detect a shifted 51st lot relative to that of the likelihood ratio test [i.e., the ratio of the power of the rank test to that of the $\chi^2(p)$ test (where p is the dimensionality of the data set)]. In other words, here, we assume that both the base data and the lots to be tested have identity covariance matrix and that this matrix is known. We note that the efficiency of the rank test here, in a situation favorable to the likelihood ratio test, is close to 1, with generally improving performance as the dimensionality increases. Here, we have used the critical values from tables of the χ^2 distribution. For such a situation, the χ^2 is the likelihood ratio test, so in a sense this is a very favorable case for the parametric test. In Figure 10.3 we apply the location test only for the data simulation delineated above.

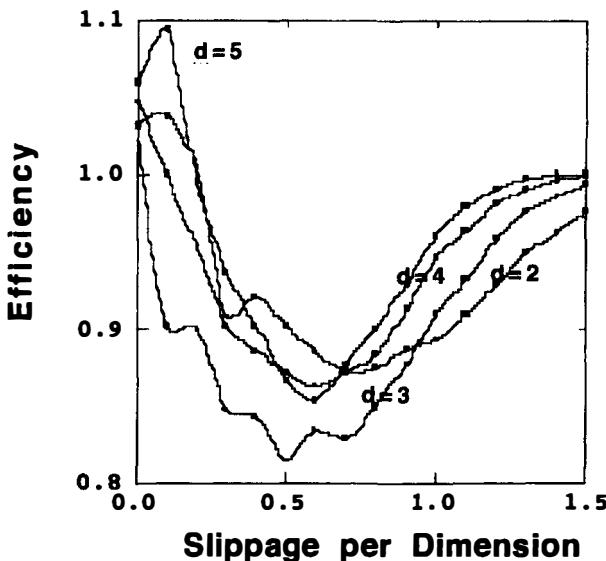


Figure 10.3. Monte Carlo estimates of efficiencies (normal data).

Next, we consider applying the rank test for location to $t(3)$ data generated in the obvious manner as shown in Figure 10.4. First, we generate a χ^2 variable v with 3 degrees of freedom. Then we generate p independent univariate normal variates $\mathbf{X}' = (X_1, X_2, \dots, X_p)$ from a normal distribution with mean 0 and variance 1. If we wish to have a mean vector μ and covariance matrix \mathbf{I} ,

$$t = \frac{\mathbf{X}}{\sqrt{v/3}} + \mu \quad (10.22)$$

will have a shifted t distribution with 3 degrees of freedom.

Once again the rank test performs well when its power is compared with that of the parametric test *even though we have computed the critical value for the parametric test assuming the data are known to be $t(3)$* . We should remember, however, that if we had assumed (incorrectly) that the data were multivariate normal, the likelihood ratio test would have been quite different and its results very bad. (Naturally, as the lot size becomes large, the central limit theorem will render the normal theory-based test satisfactory.) The rank test performs well whether the data are normal or much more diffuse, and it requires no prior information as to whether the data is normal or otherwise.

So far, we have been assuming that both the base lots and the new lots were known to have identity covariance matrices. In such a case, the appropriate parametric test is χ^2 if the data are normal, and if they are not, we have employed simulation techniques to find appropriate critical values for the distribution in question. Now, however, we shift to the situation where we believe that the covariance matrices of the new lots to be sampled

may not be diagonal. We have been assuming that the base lots (each of size 5) are drawn from $\mathcal{N}(\mathbf{0}, \mathbf{I})$. The sampled (bad) lot is drawn from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where

$$\boldsymbol{\mu} = \begin{pmatrix} \mu \\ \mu \\ \vdots \\ \mu \end{pmatrix} \quad (10.23)$$

and

$$\boldsymbol{\Sigma} = \begin{pmatrix} 1 & .8 & .8 & \dots & .8 \\ .8 & 1 & .8 & \dots & .8 \\ .8 & .8 & 1 & \dots & .8 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ .8 & .8 & .8 & \dots & \vdots \end{pmatrix}. \quad (10.24)$$

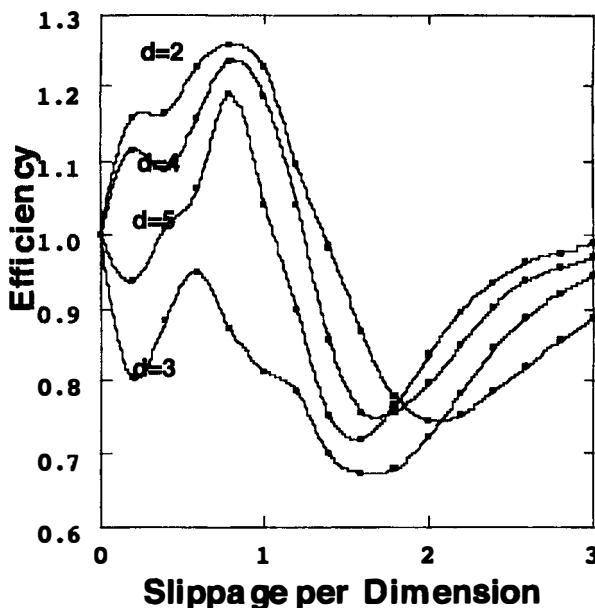


Figure 10.4. Monte Carlo estimates of efficiencies (t(3) data).

Thus, we are considering the case where the lot comes from a multivariate normal distribution with equal slippage in each dimension and a covariance matrix that has unity marginal variances and covariances .8. In Figure 10.5, we note the relative power of the “location” rank test when compared with that of the Hotelling T^2 procedure. The very favorable performance of the rank test is largely due to the effect that it picks up not only changes in location but also departures in the covariance matrix of the new lot from that of the base lots. The basic setting of statistical process control

lends itself very naturally to the utilization of normal distribution theory, since computation of lot averages is so customary. But as we have seen, for modest lot sizes it is possible to run into difficulty if the underlying distributions have heavy tails.

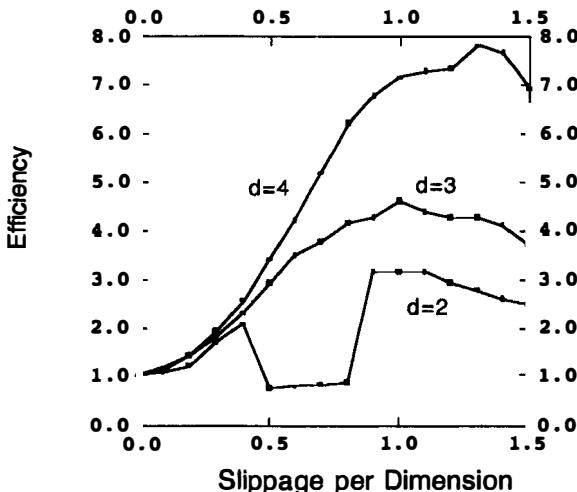


Figure 10.5. Monte Carlo estimates of efficiencies for correlated data.

In the construction of these rank tests by Thompson, Lawera and Karonacki, [7, 9], a substantial number of unsuccessful tests was examined before noting the testing procedure explicated here. Again, the utility of simulation is demonstrated. It is all very well to try an analytical approach for such tests, examining, for example, asymptotic properties. But few of us would willingly expend months of effort on tests that might well "come a cropper." Simulation gives us a means of quickly stressing potentially useful tests quickly and efficiently.

10.5 A Robust Estimation Procedure for Location in Higher Dimensions

Let us recall that, in the philosophy of Deming, one should not waste much time in determining whether a lot conforms to some predetermined standards. Many have thought themselves inspired because they came up with very strict standards for, say, manufacturing automobile transmissions. The very statement of strenuous goals is thought by many contemporary American managers (not to mention directors of Soviet *five-year plans* in a bygone age, or presidents of American universities in this age) to be a constructive act. SPC does not work that way. In SPC we seek to see epochs when lots

appear to have been produced by some variant of the dominant in-control process which produces most of the lots.

But, how shall we attempt to look into the wilderness of the past record of lots and determine what actually is, say, the location of the dominant output? It is not an easy task. Obviously, to the extent that we include out-of-control lots in our estimate of the location of the in-control lots we will have contaminated the estimate.

The following King of the Mountain algorithm of Lawera and Thompson [7] (see also Thompson and Koronacki, [9]) appears to be promising:

“King of the Mountain” Trimmed Mean Algorithm

Set the counter M equal to the historical proportion of bad lots
times number of lots.

For N lots compute the vector sample means of each lot $\{\bar{\mathbf{X}}_i\}_{i=1}^N$.

1. Compute the pooled mean of the means $\bar{\bar{\mathbf{X}}}$.

Find the two sample means farthest apart in the cloud of lot means.

From these two sample means, discard the farthest from $\bar{\bar{\mathbf{X}}}$.

Let $M = M - 1$ and $N = N - 1$.

If the counter is still positive, go to 1; otherwise exit and print out
 $\bar{\bar{\mathbf{X}}}$ as $\bar{\bar{\mathbf{X}}}_T$.

To examine the algorithm, we examine a mixture distribution of lot means

$$\gamma \mathcal{N}(\mathbf{0}, \mathbf{I}) + (1 - \gamma) \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}). \quad (10.25)$$

Here we assume equal slippage in each dimension, that is,

$$(\boldsymbol{\mu}) = (\mu, \mu, \dots, \mu). \quad (10.26)$$

Let us compare the trimmed mean procedure $\bar{\bar{\mathbf{X}}}_T$ with the customary procedure of using the untrimmed mean $\bar{\bar{\mathbf{X}}}$. In Tables 10.2 and 10.3 we show for dimensions two, three, four, and five, the average MSEs of the two estimators when $\gamma = 0.70$ for simulations of size 1000.

Table 10.2. MSEs for 50 Lots: $\gamma = 0.7$

μ	d=2 $\bar{\bar{\mathbf{X}}}_T$	d=2 $\bar{\bar{\mathbf{X}}}$	d=3 $\bar{\bar{\mathbf{X}}}_T$	d=3 $\bar{\bar{\mathbf{X}}}$	d=4 $\bar{\bar{\mathbf{X}}}_T$	d=4 $\bar{\bar{\mathbf{X}}}$	d=5 $\bar{\bar{\mathbf{X}}}_T$	d=5 $\bar{\bar{\mathbf{X}}}$
1	0.40	0.94	0.43	1.17	0.46	1.42	0.54	1.72
2	0.21	1.24	0.17	1.62	0.17	2.00	0.18	2.37
3	0.07	1.85	0.09	2.67	0.12	3.49	0.15	4.33
4	0.06	3.01	0.09	4.48	0.12	5.93	0.14	7.41
5	0.05	4.61	0.09	6.88	0.11	9.16	0.15	11.52
6	0.06	6.53	0.08	9.85	0.12	13.19	0.14	16.50
7	0.06	8.94	0.09	13.41	0.11	17.93	0.14	22.33
8	0.06	11.58	0.08	17.46	0.11	23.27	0.14	28.99
9	0.06	14.71	0.08	22.03	0.11	29.40	0.15	36.67
10	0.06	18.13	0.08	27.24	0.11	36.07	0.15	45.26

Table 10.3. MSEs for 100 Lots: $\gamma = 0.7$

μ	d=2	d=2	d=3	d=3	d=4	d=4	d=5	d=5
	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}
1	0.28	0.71	0.30	0.94	0.32	1.16	0.34	1.34
2	0.05	0.88	0.06	1.30	0.07	1.74	0.08	2.10
3	0.03	1.72	0.05	2.58	0.06	3.37	0.07	4.25
4	0.03	2.95	0.04	4.41	0.06	5.86	0.07	7.34
5	0.03	4.56	0.04	8.25	0.06	9.13	0.08	11.39
6	0.03	6.56	0.04	8.31	0.06	13.08	0.07	16.35
7	0.03	8.87	0.04	13.28	0.06	17.67	0.07	22.16
8	0.03	11.61	0.04	17.32	0.06	23.17	0.07	28.86
9	0.03	14.67	0.04	21.98	0.06	29.28	0.07	36.61
10	0.03	18.04	0.04	27.12	0.06	36.06	0.07	45.13

In Tables 10.4 and 10.5 we show the MSEs of the trimmed mean and the customary pooled sample mean for the case where $\gamma = 0.95$.

Table 10.4. MSEs for 50 Lots: $\gamma = 0.95$

μ	d=2	d=2	d=3	d=3	d=4	d=4	d=5	d=5
	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}
1	0.05	0.10	0.07	0.15	0.09	0.18	0.12	0.24
2	0.05	0.11	0.07	0.16	0.09	0.20	0.11	0.26
3	0.04	0.11	0.06	0.17	0.08	0.22	0.11	0.28
4	0.04	0.13	0.06	0.19	0.08	0.26	0.11	0.33
5	0.04	0.16	0.06	0.24	0.08	0.32	0.11	0.41
6	0.04	0.19	0.06	0.29	0.09	0.40	0.10	0.49
7	0.04	0.24	0.06	0.35	0.08	0.48	0.11	0.61
8	0.04	0.29	0.06	0.42	0.08	0.58	0.10	0.71
9	0.04	0.33	0.06	0.51	0.08	0.68	0.11	0.86
10	0.042	0.39	0.06	0.59	0.08	0.79	0.11	1.01

Table 10.5. MSEs for 100 Lots: $\gamma = 0.95$

μ	d=2	d=2	d=3	d=3	d=4	d=4	d=5	d=5
	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}	\bar{X}_T	\bar{X}
1	0.02	0.05	0.04	0.09	0.05	0.12	0.06	0.14
2	0.02	0.07	0.03	0.10	0.05	0.14	0.05	0.16
3	0.02	0.09	0.03	0.14	0.04	0.18	0.05	0.23
4	0.02	0.13	0.03	0.20	0.05	0.27	0.06	0.34
5	0.02	0.18	0.03	0.27	0.04	0.37	0.06	0.46
6	0.02	0.24	0.03	0.38	0.04	0.49	0.05	0.60
7	0.02	0.31	0.03	0.47	0.04	0.62	0.06	0.80
8	0.02	0.40	0.03	0.60	0.04	0.80	0.05	1.00
9	0.02	0.50	0.03	0.74	0.04	1.00	0.05	1.21
10	0.02	0.60	0.03	0.90	0.05	1.20	0.05	1.50

If the level of contamination is substantial (e.g., if $1 - \gamma = 0.3$), the use of a trimming procedure to find a base estimate of the center of the in-control distribution contaminated by observations from other distributions may be

strongly indicated. For more modest but still significant levels of contamination (e.g., if $1 - \gamma = 0.05$), simply using $\bar{\bar{X}}$ may be satisfactory. We note that the trimming procedure considered here is computer intensive and is not realistic to be performed on the usual hand-held calculator. However, it is easily computed on a personal computer or workstation. Since the standards for rejecting the null hypothesis that a lot is in control are generally done by off-line analysis regularly, we do not feel that the increase in computational complexity should pose much of a logistical problem.

Problems

10.1. It is desired to simulate a contamination model for training purposes. We wish to produce sheets of aluminum with thickness 1 mm. Suppose the dominant distribution is given by $\mathcal{N}(1, 1)$. Time is divided up into epochs of 10 minutes. Contamination will occur in each epoch with probability γ . The contaminating distribution will be $\mathcal{N}(\mu, \sigma^2)$. Simulate data for a variety of γ 's, μ 's, and σ 's. Using Shewhart control charting (10.8), see how effective you are in finding lots that are from contamination periods.

10.2. Let us consider the case where there are four measurables in a production process. After some transformation, the in control situation would be represented by a normal distribution with mean

$$\mu = \begin{pmatrix} \mu_1 & = & 0 \\ \mu_2 & = & 0 \\ \mu_3 & = & 0 \\ \mu_4 & = & 0 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \mathbf{I} = \begin{pmatrix} \sigma_{11} = 1 & \sigma_{12} = 0 & \sigma_{13} = 0 & \sigma_{14} = 0 \\ \sigma_{12} = 0 & \sigma_{22} = 1 & \sigma_{23} = 0 & \sigma_{24} = 0 \\ \sigma_{13} = 0 & \sigma_{23} = 0 & \sigma_{33} = 1 & \sigma_{34} = 0 \\ \sigma_{14} = 0 & \sigma_{24} = 0 & \sigma_{34} = 0 & \sigma_{44} = 1 \end{pmatrix}.$$

We wish to examine the effectiveness of the battery of tests procedure in Section 10.3, for the following situations, each with lot sizes of 10:

(a) There is 5% contamination from the Gaussian distribution with

$$\mu = \begin{pmatrix} \mu_1 & = & 1 \\ \mu_2 & = & 1 \\ \mu_3 & = & 1 \\ \mu_4 & = & 1 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \mathbf{I}.$$

(b) There is 5% contamination from a Gaussian distribution with

$$\mu = \begin{pmatrix} \mu_1 & = & 0 \\ \mu_2 & = & 2 \\ \mu_3 & = & -2 \\ \mu_4 & = & 0 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \mathbf{I}.$$

(c) There is 10% contamination from a Gaussian distribution with

$$\mu = \begin{pmatrix} \mu_1 & = & 1 \\ \mu_2 & = & 1 \\ \mu_3 & = & 1 \\ \mu_4 & = & 1 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_{11} = 1 & \sigma_{12} = .6 & \sigma_{13} = .6 & \sigma_{14} = .6 \\ \sigma_{12} = .6 & \sigma_{22} = 1 & \sigma_{23} = .6 & \sigma_{24} = .6 \\ \sigma_{13} = .6 & \sigma_{23} = .6 & \sigma_{33} = 1 & \sigma_{34} = .6 \\ \sigma_{14} = 0 & \sigma_{24} = 0 & \sigma_{34} = 0 & \sigma_{44} = 1 \end{pmatrix}.$$

10.3. Next, let us consider the case where there are three measurables in a production process. After some transformation, the in control situation would be represented by a normal distribution with mean

$$\mu = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$$

and covariance matrix

$$\Sigma = \mathbf{I}.$$

There is 10% contamination of lots (i.e., 10% of the lots are from the contaminating distribution). The contaminating distribution is multivariate t with 3 degrees of freedom (see (10.22)) and translated to

$$\mu = \begin{pmatrix} \mu_1 & = & 1 \\ \mu_2 & = & 1 \\ \mu_3 & = & 1 \end{pmatrix}.$$

If lot sizes are 8, compare the effectiveness of the battery of tests above with the rank test given in Section 10.4.

10.4. You have a six-dimensional data set that turns out to be

$$\gamma \mathcal{N}(\mathbf{0}, \mathbf{I}) + (1 - \gamma) t_3(1, \mathbf{I}).$$

The problem is to recover the mean of the uncontaminated Gaussian distribution for the case where $\gamma = .2$. Lot sizes are 10. See how well the King of the Mountain Algorithm works for this case compared to using pooled lot means.

References

- [1] Alt, F. B. (1985). "Multivariate quality control," in *Encyclopedia of Statistical Sciences*, 6, S. Kotz, and N. Johnson eds. New York: John Wiley & Sons, 110–122,
- [2] Banks, D. (1993). "Is industrial statistics out of control?" *Statist. Sc.*, 8, 356–409.
- [3] Deming, W. E. (1982). *Quality, Productivity and Competitive Position*. Cambridge, Mass.: Center for Advanced Engineering Study, Massachusetts Institute of Technology.
- [4] Deming, W. E. (1986). *Out of the Crisis*. Cambridge, Mass.: Center for Advanced Engineering Study, Massachusetts Institute of Technology.
- [5] Juran, J. M. (1995) "The Non-Pareto Principle; Mea Culpa," *Quality Prog.*, 149–150.
- [6] Koronacki, J. and Thompson, J.R. (1997). "Statystyczne sterowanie procesem wielowymiarowym" ("Statistical process control in high dimensions"), *Problemy Jakości*, xxix, 8–15.
- [7] Lawera, M. and Thompson, J. R. (1993). "Some problems of estimation and testing in multivariate statistical process control," *Proceedings of the Thirty-Fourth Conference on the Design of Experiments in Army Research Development and Testing*. Research Triangle Park, N.C.: Army Research Office, 99–126.
- [8] Stuart, A. and Ord, J. K. (1991). *Kendall's Advanced Theory of Statistics*, 2. New York: Oxford University Press, 870.
- [9] Thompson, J. R. and Koronacki, J. (1992). *Statistical Process Control for Quality Improvement*. New York: Chapman & Hall.
- [10] Thompson, J. R. (1989). *Empirical Model Building*. New York: John Wiley & Sons.

Chapter 11

Considerations for Optimization and Estimation in the Real (Noisy) World

11.1 Introduction

In 1949, Abraham Wald [19] attempted to clean up the work of Fisher by proving that the maximum likelihood estimator $\hat{\theta}_n$ of the parameter characterizing a probability density function converged to the true value of the parameter θ_0 . He was indeed able to show under very general conditions that if $\hat{\theta}_n$ (globally) maximized the likelihood, it did converge almost surely to θ_0 . From a practical standpoint, there is less to the Wald result than one might have hoped. The problem is that, in most cases, we do not have good algorithms for global optimization.

Let us suppose that we seek a minimum to a function $f(x)$. In the minds of many, we should use some variant of Newton's method to find the minimum. Now, Newton's method does not seek to find the minimum of a function, but rather the (hopefully unique) point where the first derivative of the function is equal to zero. We recall, then, that the simplest of the Newton formulations is an iterative procedure, where

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)}. \quad (11.1)$$

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

Returning to Wald's result, let us suppose that we consider data from a Cauchy density

$$f(x) = \frac{1}{\pi[1 + (x - \theta_0)^2]}. \quad (11.2)$$

Then we can pose the maximum likelihood estimation by minimizing the negative of the log likelihood:

$$L(\theta|x_1, x_2, \dots, x_n) = n \log(\pi) + \sum_{i=1}^n \log[1 + (x_i - \theta)^2]. \quad (11.3)$$

If we use Newton's method here, if we have a starting guess to the left of the smallest point in the data, we will tend to declare that the smallest of the data is the maximum likelihood estimator for θ_0 , for that point is a local minimum of the negative of the log likelihood. If we start to the right of the largest data value, we will declare the largest data point to be a maximum likelihood estimator. Now as the data set becomes large, if we do not start on the fringes, we have much better fortune. We do not have multiple "bumps" of the log likelihood near the middle of a large data set. As the data set becomes larger and larger, a starting point that gave a misleading bump at a data far away from θ_0 will, in fact, lead us to an acceptable estimator for θ_0 . In one sense, this is true broadly for the problem of estimation by maximum likelihood or minimum χ^2 , for example, by Newton's method will become less and less as the sample size increases. We show this stabilization with increasing sample size from a Cauchy distribution with $\theta = 0$ in Figure 11.1.¹ We see, for example, that for the data set explored, for a sample size of 10, we have no false local modes if we start with a positive θ value less than approximately 13, but if we start with a greater value, we might, using Newton's method, wind up with a false maximum (i.e., one that is not equal to the global maximum). Picking any interval of starting values for θ and any values ϵ and δ , however, there will be a sample size such that the probability will be less than ϵ that Newton's method will converge to a value more than δ removed from the true global maximum, namely, θ_0 .

This is a phenomenon occurring much more generally than in the case of Cauchy data. In the case of the use of SIMEST in Chapter 5 in the estimation of parameters in a cancer model, for example, the use of the algorithm with sample sizes of 150 demonstrated problems with local maxima, unless one started very near the global maximum. As the number of patients increased past 700, the problem of local maxima of the likelihood (minima of the χ^2) essentially disappeared. That was due to the fact that a larger sample, for maximum likelihood estimation, brings a starting value, unacceptable for smaller samples, into the domain of attraction of the global maximum, and the bumps which existed for the smaller samples, tend to

¹This figure was created by Otto Schwalb.

"tail off" (i.e., appear remotely from reasonable starting values). As the sample sizes become large, problems of finding the global maximum of the sample likelihood tend to become less. So, as a practical matter, Wald's result is actually useful if the sample size be sufficiently large.

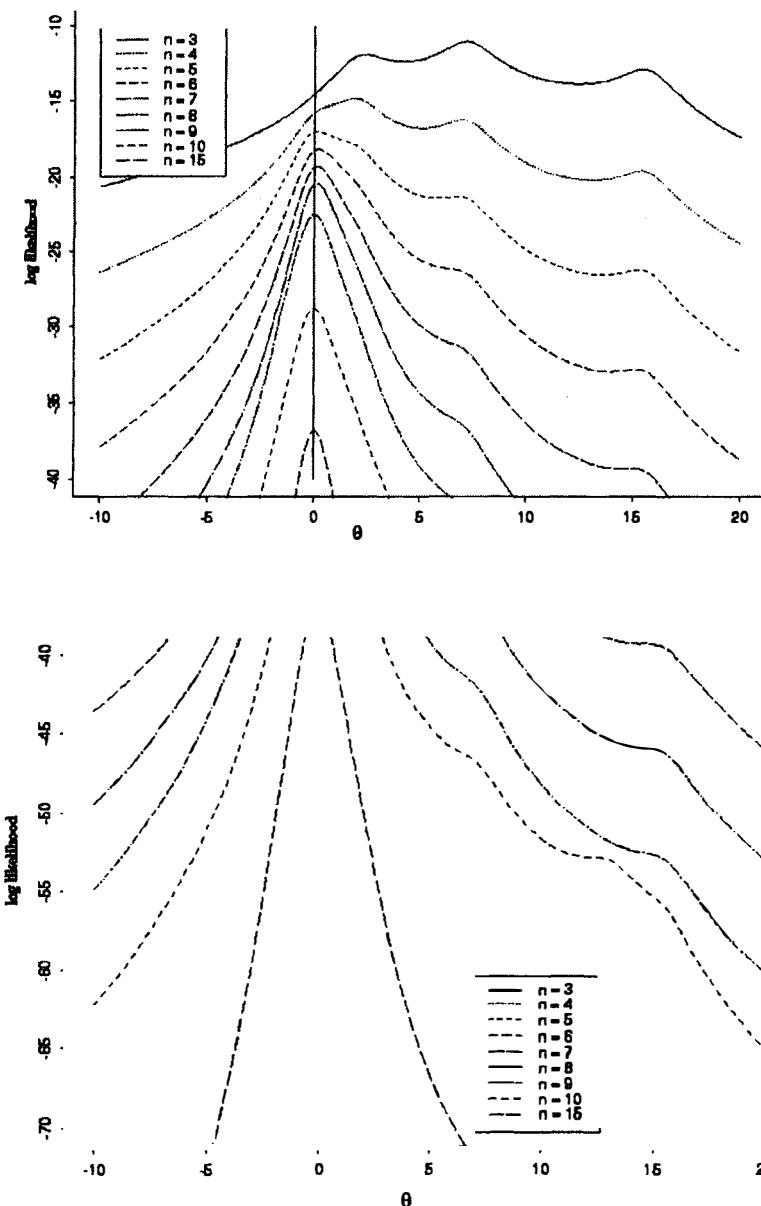


Figure 11.1. Cauchy log likelihoods for various n .

Practical Version of Wald's Result. Let $\{x_1, x_2, \dots, x_n\}$ be

a random sample with a density function $f(x|\theta_0)$ with positive mass throughout its support $a < x < b$. a and/or b may or may not be ∞ . For any fixed (i.e., not changing with n) starting point between a and b , as n goes to ∞ , if f is well behaved, a Newton's algorithm maximizer of the log likelihood function will converge almost surely to θ_0 .

In other words, for the statistician, the natural piling up of data points around regions of high density will cause a practical convergence of the naive maximum likelihood estimator to the truth. Of course, the number of points required to make this conjecture useful may be enormous, particularly if we use it for a multivariate random variable and multivariate characterizing parameter. Statisticians have an advantage over others who deal in optimization, for, generally speaking, the function to be maximized in most problems is not a density function, so that the possibility of never converging to a global maximum is a real one. This, of course, suggests the attractive possibility of trying to reformulate an objective function as a probability density function when feasible to do so.

We go further with a conjecture on simulation-based estimation (SIMEST discussed in Chapter 5). Since the time of Poisson [11], it has been taken as natural to model time-based processes in the forward direction. For example, “with a probability proportional to its size, a tumor will, in any time interval, produce a metastasis.” Easy to state, easy to simulate—not so easy to find the likelihood, particularly when the metastatic process is superimposed upon other simultaneously occurring processes. For example, the tumor is also growing in proportion to its size; the tumor may be discovered and removed with a probability proportional to its size; and so on. But the simulations superimpose quite readily. Hence, given fast computing, we are tempted to assume the unknown parameters characterizing the pooled processes, generate relevant events by simulation, and then use the difference between the simulated process and the actual, say, discovery of tumors, as a measure of the quality of assumed parameters.

SIMEST Conjecture. Let $\{x_1, x_2, \dots, x_n\}$ be a random sample with a (possibly not known in closed form) density function $f(x|\theta_0)$. Suppose we can generate, for a given θ , N pseudovalues of $\{y_j\}$ from the density function. Create bins in an intuitive way in the data space, for example putting n/k data points and N_i data points into each of k bins. If the maximum likelihood estimator for θ_0 , in the case where we know the closed form of the log likelihood converges to θ_0 , then so does an estimator $\hat{\theta}_n$, which maximizes the histogram log likelihood function

$$L_H(\theta|x_1, x_2, \dots, x_n) = \sum_{i=1}^k N_i \log(n/k)$$

based on N pseudodata, as n goes to ∞ and N goes to ∞ if we let k go to ∞ in such a way that $\lim_{n \rightarrow \infty} k/n = 0$.

The conjecture concerning simulation-based estimation is rather powerful stuff, because it raises the possibility of parameter estimation in incredibly complex modeling situations. Naturally, for the situation where X is vector valued, some care must be taken in finding appropriate binning strategies. A somewhat differently styled version of the SIMEST conjecture has been proved by Schwab [13].

As a matter of fact, the statistician is generally confronted with finding the maximum of an objective function which is contaminated by noise. In the case of simulation-based parameter estimation, the noise is introduced by the modeler himself. The use of Newton-like procedures will generally be inappropriate, since derivatives and their surrogates will be even more unstable than pointwise function evaluation. We shall discuss two ways of dealing with this problem. Interestingly, both the Nelder-Mead algorithm and the Box-Hunter algorithm were built, not by numerical analysts, but by statisticians, working in the context of industrial product optimization. The algorithm of Nelder and Mead essentially gives up on equivalents of "setting the first derivative equal to zero." Rather, it follows an *ad hoc* and frequently very effective zig-zag path to the maximum using pointwise function evaluations without any derivative-like evaluation. The essential idea is to approach the maximum indirectly and, therefore, hopefully, with some robustness. The mighty quadratic leaps toward the maximum promised by Newton's method are not available to the N-M user. On the other hand, neither are the real-world leaps to nowhere-in-particular that frequently characterize Newton's method.

The algorithm of Box and Hunter, however, takes noisy pointwise function evaluations over a relatively small hyperspherical region and uses them to estimate the parameters of a locally approximating second degree polynomial. Essentially, with Box-Hunter we do take the derivative of the fitting polynomial and proceed to the maximum by setting it equal to zero. But remembering that the fitting validity of the polynomial is generally credible only in a rather small region, we cannot take the giant leaps to glory (or perdition) associated with Newton's method.

11.2 The Nelder-Mead Algorithm

The problem of parameter estimation is only one of many. For most situations, we will not have samples large enough to enable Newton's method to do us much good. Newton method's is generally not very effective for most optimization problems, particularly those associated with data analysis.

A more robust algorithm, one pointing more clearly to the direct search for the minimum (or maximum) of a function, is needed. (And there is still the problem of trying to find the global minimum, not simple some

local minimum. We deal with this problem later). More than 30 years ago, two statisticians, Nelder and Mead [10], designed an algorithm which searches directly for a minimum rather than for zeros of the derivative of the function. It does not assume knowledge of derivatives, and it generally works rather well when there is some noise in the pointwise evaluation of the function itself. It is not fast, particularly in higher dimensions. This is due, in part, to the fact that the Nelder–Mead algorithm employs a kind of envelopment procedure rather than one which, as does Newton’s method, tries to move directly to the minimum. The Nelder–Mead algorithm is rather intuitive, and, once learned, is easy to construct. We give the algorithm below with accompanying graphs (Figures 11.2 and 11.3) showing the strategy of moving toward the minimum. Our task is to find the minimum of the function $f(x)$. Here, we consider a two-dimensional x .

Nelder–Mead Algorithm

Expansion

- $P = C + \gamma_R(C - W)$ (where typically $\gamma_R = \gamma_E = 1$)
- If $f(P) < f(B)$, then
- $PP = C + \gamma_E(C - W)$ [a]
- If $f(PP) < f(P)$, then
- Replace W with PP as new vertex [c]
- Else
- Accept P as new vertex [b]
- End If

Else

If $f(P) < f(2W)$, then

- Accept P as new vertex [b]
- Else

Contraction

If $f(W) < f(P)$, then

- $PP = C + \gamma_C(W - B)$ (typically, $\gamma_C = 1/2$) [a*]
- If $F(PP) < F(W)$, Then replace W with PP as new vertex [b*]
- Else replace W with $(W+B)/2$ and $2W$ with $(2W+B)/2$ (total contraction) [c*]
- End If

Else

Contraction

If $f(2W) < f(P)$, then

- $PP = C + \gamma_C(P - B)$ [aa]
- If $f(PP) < f(P)$, then Replace W with PP as new vertex [bb]
- Else replace W with $(W+B)/2$ and $2W$ with $(2W+B)/2$ (total contraction) [cc]

Else

- Replace W with P
- End If

End If

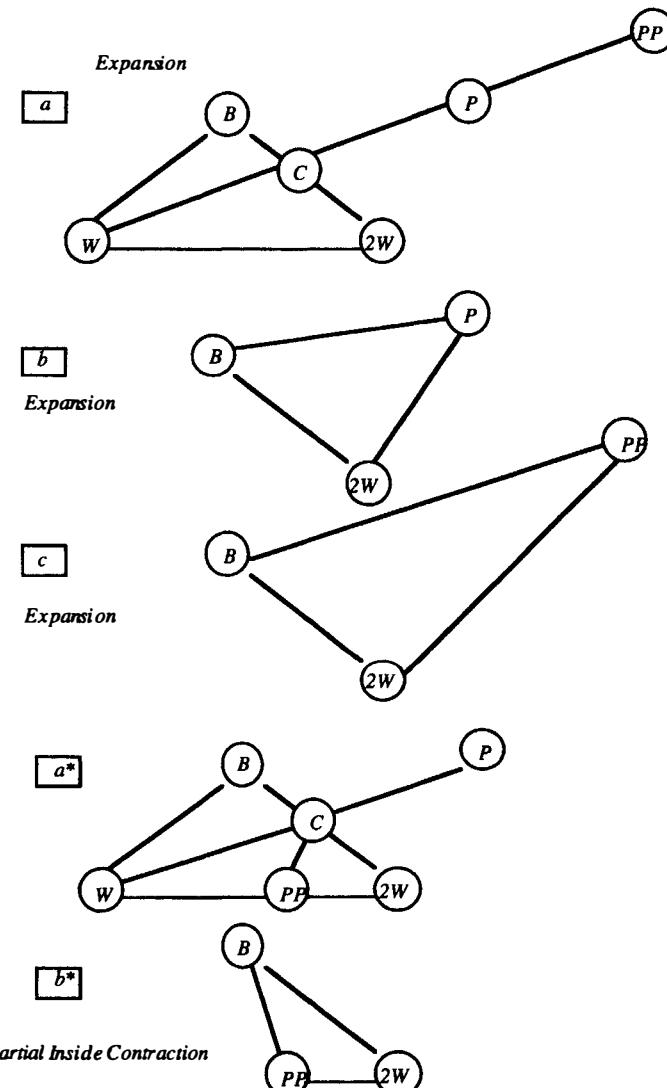


Figure 11.2. Nelder-Mead polytope expansions.

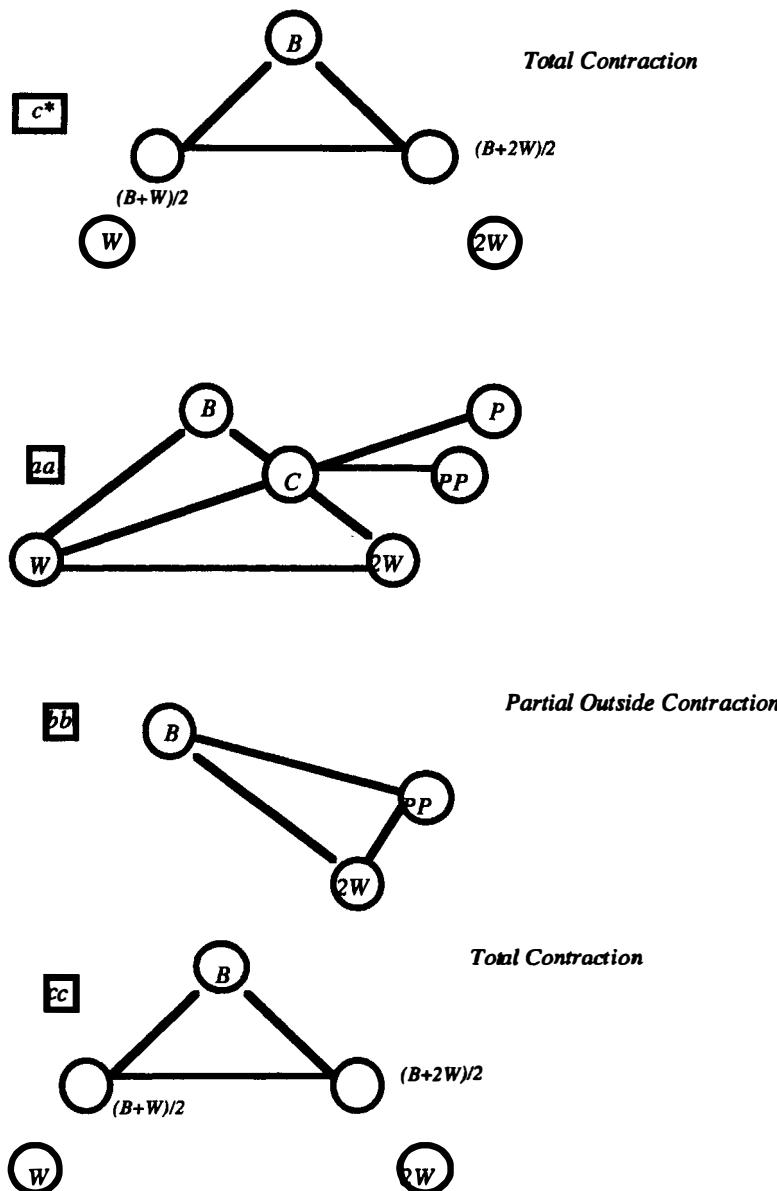


Figure 11.3. Nelder–Mead polytope contractions.

11.3 The Box–Hunter Algorithm

Both the Nelder–Mead [10] and Box–Hunter [5] algorithms were designed with an eye for use in the design of industrial experiments. In fact, they

have been used both as experimental design techniques and as computer optimization routines. However, the Box–Hunter designs clearly are the more used in industry, the Nelder–Mead approach the more used in computer optimization. The reasons are not hard to understand. When it comes to industrial experiments, where great costs are incurred, the rather free-wheeling nature of Nelder–Mead appears profligate. On the computer, where an “experiment” generally simply involves a function evaluation, the Box–Hunter approach appears overly structured, with a “batch” rather than a continuous-flow flavor. However, the natural parallelization possibilities for Box–Hunter should cause us to rethink whether it might be the basis for a new theory of computer optimization and estimation.

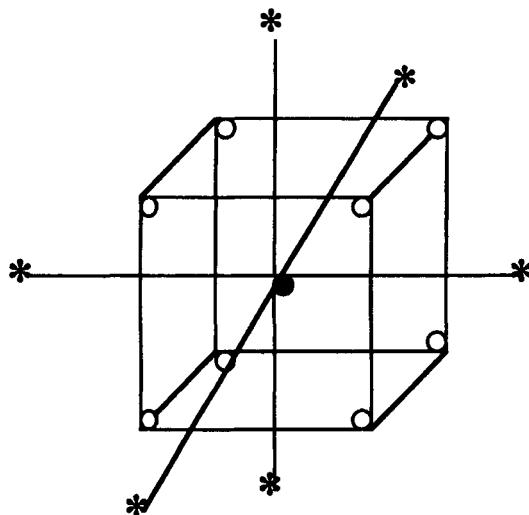


Figure 11.4. Box–Hunter three-dimensional design.

Essentially, the Box–Hunter rotatable design [17] approach centers at the current best guess for the optimum. Points are then placed at the degenerate sphere at that center, the coordinates rescaled so that a movement of one unit in each of the variables produces approximately the same change in the objective function. Then a design is created with points on a hypersphere close to the origin and then on another hypersphere farther out. The experiment is carried out, and the coefficients of an approximating quadratic are estimated. Then we move to the new apparent optimum and repeat the process.

In the ensuing discussion, we follow the argument of Lawera and Thompson [9]. The variation of Box–Hunter was created, in large measure, to deal with the application of SIMEST (see Chapter 5) to parameter estimation in stochastic processes. In Figure 11.4, we show a three-dimensional Box–Hunter design. In standardized scaling, the points on the inner hypersphere are corners of the hypercube of length two on a side. The second

sphere has “star points” at $2^{p/4}$, where p is the dimensionality of the independent variables over which optimization is taking place. For three dimensions, we have points at $(0,0,0)$, $(1,1,1)$, $(1,1,-1)$, $(1,-1,1)$, $(1,1,-1)$, $(-1,-1,1)$, $(-1,-1,-1)$, $(-1,1,1)$, $(-1,1,-1)$, $(2^{.75},0,0)$, $(-2^{.75},0,0)$, $(0,2^{.75},0)$, $(0,-2^{.75},0)$, $(0,0,2^{.75})$, and $(0,0,-2^{.75})$.

For dimensionality p we start with an orthogonal factorial design having 2^p points at the vertices of the (hyper)cube $(\pm 1, \pm 1, \dots, \pm 1)$. Then we add $2p$ star points at $(\pm \alpha, 0, \dots, 0)$, $(0, \pm \alpha, 0, \dots, 0)$, \dots , $(0, 0, \dots, 0, \pm \alpha)$. Then we generally add two points (for, say, $p \leq 5$, more for larger dimensionality) at the origin. A sufficient condition for rotatability of the design, i.e., that, as above, $\text{Var}(\hat{y})$ is a function only of

$$\rho^2 = X_1^2 + X_2^2 + \dots + X_p^2, \quad (11.4)$$

can be shown to be [4] that

$$\alpha = (2^p)^{.25}. \quad (11.5)$$

In Table 11.1 we show rotatable designs for dimensions 2, 3, 4, 5, 6 and 7.

Table 11.1. Some Box–Hunter Rotatable Designs.

Dimension	Num. Cube Points	Num. Center Points	Num. Star Points	α
2	4	2	4	$2^{.5}$
3	8	2	6	$2^{.75}$
4	16	4	8	2
5	32	4	10	$2^{1.25}$
6	64	6	12	$2^{1.5}$
7	128	8	14	$2^{1.75}$

When the response variable has been evaluated at the design points, we then use least squares to fit a quadratic polynomial to the results.

$$J_1(\Theta) = \beta_0 + \sum_{i=1}^p \beta_i \Theta_i + \sum_{i=1}^p \sum_{j=1}^p \beta_{ij} \Theta_i \Theta_j. \quad (11.6)$$

We then transform the polynomial to canonical form \mathbf{A} :

$$J_2(\Theta) = \beta_0 + \sum_{i=1}^p \beta_i \Theta_i + \sum_{i=1}^p \beta_{ii} \Theta_i^2. \quad (11.7)$$

Let us now flowchart the Lawera–Thompson version of the Box–Hunter algorithm. First, we define some notation:

- Θ_0 coordinates of current minimum
- D $n \times 2^n + 2n + n_0$ Box–Hunter design matrix
- R $n \times n$ diagonal matrix used to transform into “absolute” coordinate system
- T $n \times n$ matrix which rotates the axes of the design to coincide with the “absolute” axes

(Note that the design points as given by matrix D have “absolute” coordinates given by $T \times R \times D + \Theta_0$.)

$S(.)$	the objective function
EC	prespecified by the user upper limit on noise level
$CONV$	user-specified constant in the convergence criterion at level 2
X_R	$(2^n + 2n + n_0) \times [1 + n + n(n + 1)/2]$ matrix of regression points

X can be written as $[1, (R \times D)^T, rd_{11}, \dots, rd_{nn}]$, where 1 is a column vector obtained by elementwise multiplication of the i th and the j th columns of $(R \times D)^T$.

Lawera–Thompson Algorithm

Level 1

Input (initial guess): Θ_0, R_0, T_0

1. Perform the level 2 optimization starting from the initial guess. Output: Θ_1, R_1, T_1 .
2. Perform the level 2 optimization 10 times, starting each time from the results obtained in (1).
3. Find $S_{min}(\Theta_{min})$: the best of results obtained in (2).

Output: $\Theta_{min}, S_{min}(\Theta_{min})$

Level 2

Input: Θ_0, R_0, T_0, EC_0

1. $EC \leftarrow EC_0$
2. Perform the level 3 optimization using the input values. Output: $\Theta_1, R_1, T_1, S(\Theta_1)$
3. Calculate the distance between Θ_0 and Θ_1 , i.e., $\Delta\Theta = \sqrt{\|\Theta_1 - \Theta_0\|^2}$.
4. Calculate the gain from (2): $\Delta S = S(\Theta_0) - S(\Theta_1)$.
5. If $\Delta\Theta > 0$ and $\Delta S > 1.5 \times \sqrt{EC}$, then $\Theta_0 \leftarrow \Theta_1, R_0 \leftarrow R_1, T_0 \leftarrow T_1$ and go to (1).
6. Else if $EC > CONV$, then $EC \leftarrow EC/4$, and go to (2).
7. Else exit to level 1.

Output: $\Theta_1, R_1, T_1, S(\Theta_1)$.

Level 3

Input: $\Theta_0, R_0, T_0, EC_0, S(\Theta_0)$

1. Perform level 4. Output: R_{min} .
2. Set $Y \leftarrow NULL$, $X \leftarrow NULL$, $i \leftarrow 0$, $\Theta_1 \leftarrow \Theta_0$.
3. Set $R_{Cur} \leftarrow R_{min}$.
4. Increment i by one.
5. Evaluate $S(\Theta)^* = S(T_0 \times R_{Cur} \times D + \Theta_1)$.
6. Calculate $X_{R_{Cur}}$.
7. Set $X \leftarrow \begin{pmatrix} X \\ X_{R_{Cur}} \end{pmatrix}$.
8. Set $Y \leftarrow \begin{pmatrix} Y \\ S(\Theta)^* \end{pmatrix}$.
9. Regress Y on X . Obtain: vector of regression coefficients $\{\hat{\beta}\}$ and the r^2 statistic.
10. Perform the level 5 optimization. Output: Θ^* , R^* , T^* , $S(\Theta)^*$.
11. Calculate the gain from (5): $\Delta S = S(\Theta_1) - S(\Theta)^*$.
12. If $r^2 > 0.9$, $\Delta S > 1.5 \times \sqrt{EC}$, $i < 20$, then
 $\Theta_1 \leftarrow \Theta^*$, $R_1 \leftarrow R^*$, $T_1 \leftarrow T^*$
 $R_{Cur} \leftarrow 2 \times R_{Cur}$, and go to (4).
13. Else exit to level 2.

Output: Θ_1 , R_1 , T_1 , $S(\Theta_1)$

Level 4

Input: Θ_0 , R_0 , T_0 , EC_0 , $S(\Theta_0)$

1. Evaluate $S(\Theta_0)^* = S(T_0 \times R_0 \times D + \Theta_0)$.
2. Calculate X_{R_0} .
3. Regress $S(\Theta_0)^*$ on X_{R_0} . Obtain the r^2 statistic and the error sum of squares (ESS).
4. If $r^2 < 0.9$ and [$ESS < 2 \times EC_0$, or $\text{Max}(S(\Theta_0)) - \text{Min}(S(\Theta_0)) < 1.5 \times \sqrt{EC}$],
then
(a) Set $R_0 \leftarrow 2 \times R_0$.

(b) Repeat (1)–(4) until $r^2 > 0.9$, or [$ESS > 2 \times EC_0$, and $\text{Max}(S(\Theta_0)) - \text{Min}(S(\Theta_0)) < 1.5 \times \sqrt{EC}$].

(c) Exit to level 3.

5. Else

(a) Set $R_0 \leftarrow 0.5 \times R_0$.

(b) Repeat (1)–(4) until $r^2 < 0.9$, and $ESS < 2 \times EC_0$, or $\text{Max}(S(\Theta_0)) - \text{Min}(S(\Theta_0)) < 1.5 \times \sqrt{EC}$.

Set $R_0 \leftarrow 2 \times R_0$.

(c) Exit to level 3.

Output: R_0

Level 5²

Input: $\hat{\beta}$, quadratic fit to the objective function

1. Calculate vector b and matrix B such that the quadratic fit has the form
$$\hat{y} = b_0 + X^T \times b + X^T \times B \times X.$$
2. Find matrices M and Λ such that $M^T \times B = \Lambda$.
3. Calculate the minimum $\Theta \leftarrow -1/2B^{-1} \times b$.
4. If $\sqrt{\|\Theta\|^2} > 1$, then
 - (a) Set $\text{Min}(|\Theta_1|, \dots, |\Theta_n|) \leftarrow 0$.
 - (b) Repeat (a) until $\sqrt{\|\Theta\|^2} \leq 1$.
5. Calculate the rescaling matrix $R_0 \leftarrow \text{Diag}(|\lambda_i|^{-1/2})$.
6. Set $T_0 \leftarrow M$.

Output: Θ_0, R_0, T_0

Evaluation

Input: Θ, R, T, EC

1. Set $i \leftarrow 0$
2. Evaluate S 10 times at $\Theta_0 = T \times R \times D + \Theta$.
3. Increment i by 7.
4. Calculate the sample mean \bar{S} and the sample variance V of all i evaluations.

²Level 5 is based on Box and Draper [4].

5. If $V/EC > C_{i-1}$, where C_{i-1} is the 95th percentile of the χ^2_{i-1} distribution, then go to (2).

6. Else exit.

Output \bar{S}

Using the Lawera–Thompson variant of the Box–Hunter algorithm on generated tumor data (150 patients) with $\alpha = 0.31$, $\lambda = 0.003$, $a = 1.7 \times 10^{-10}$, $b = 2.3a \times 10^{-9}$, using a SIMEST sample size of 1500, with starting value $0(0.5, 0.005, 4 \times 10^{-10}, 10^{-9})$, and using 10 bins, we converged to $(0.31, 0.0032, 2 \times 10^{-10}, 2.3 \times 10^{-9})$. Computations were carried out on a Levco desktop parallel processor with 16 CPUs. Subsequent availability of very fast and inexpensive serial machines has caused us, temporarily, to suspend the parallel investigation. It is clear, however, that the Box–Hunter paradigm, requiring minimal handshaking between CPUs, is a natural candidate for parallelization.

11.4 Simulated Annealing

The algorithm of Nelder–Mead and that of Box–Hunter will generally not stall at a local minimum with the same degree of risk as a Newton's method based approach. Nevertheless, it sometimes happens that stalling local minima do occur with Nelder–Mead. How to get around this problem?

Naturally, there is no easy answer. Practically speaking, there is no general way to make sure that a minimum is global without doing a search over the entire feasible region. We might well converge, using Nelder–Mead to point **A** in Figure 11.5. We need some way to make sure that we really have arrived at the global minimum.

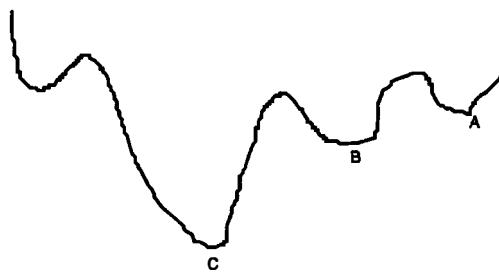


Figure 11.5. The problem with local minima.

In the example in Figure 11.5, we might argue that we need something to kick us away from **A**. Then, we can see if, say using Nelder Mead, we move to a new candidate for the global minimum, say **B**, or fall back to **A**. From the picture, it is clear that it is a critical matter just how far we move away from **A** which determines whether we will progress on to **B** and thence to

C, or whether we will fall back into A. Mysterious analogies to Boltzmann energy levels are probably not very helpful, but we shall mention the idea, since users of *simulated annealing* generally do.

In cooling a molten metal too quickly, one may not reach a level of minimum energy (and hence apparently of crystalline stability). When this happens, it sometimes happens that a decision is made to reheat the metal (although not to the molten state, necessarily) and then cool it down again, slowly, hoping to move to another lower-energy state. The probability of moving from energy state E_1 to another energy state E_2 , when $E_2 - E_1 = \Delta E > 0$, is given by $\exp(-\Delta E/(kT))$ where k is the *Boltzmann constant*. The analogy is relatively meaningless, and generally the $1/(kT)$ is simply replaced by a finagle factor. The finagle factor will determine how far we kick away from the apparent minimum. As time progresses, we may well decide to "stop the kicking."

Let us suppose that we have used Nelder–Mead to get to a minimum point x_0 . Using this as our starting point, we can use an algorithm suggested by Bohachevsky, Johnson, and Stein [1, 12].

BJS "General" Simulated Annealing Algorithm

1. Set $f_0 = f(x_0)$. If $|f_0 - f_m| < \epsilon$, stop.
2. Generate n independent standard normal variates Y_1, Y_2, \dots, Y_n . Let $U_i = Y_i/(Y_1^2 + Y_2^2 + \dots + Y_n^2)^{1/2}$ for $i = 1, 2, \dots, n$.
3. Set $x^* = x_0 + (\Delta r)U$.
4. If x^* is not in the feasible set, return to step 2, otherwise, set $f_1 = f(x^*)$ and $\Delta f = f_1 - f_0$.
5. If $f_1 \leq f_0$, set $x_0 = x^*$ and $f_0 = f_1$. If $|f_0 - f_m| < \epsilon$, stop. Otherwise, go to step 2.
6. If $f_1 > f_0$, set $p = \exp(-\beta f_0^g \Delta f)$.
7. Generate a uniform $\mathcal{U}(0,1)$ random variate V . If $V \geq p$, go to step 2. If $V < p$, set $x_0 = x^*$, $f_0 = f_1$ and go to step 2.

It is clear that the above algorithm contains a fair amount of things that are a bit arbitrary. These include

- Step size Δr

- g
- β
- f_m assumed value of the global minimum

Before the advent of simulated annealing, investigators tried to seek pathways to the global optimum by starting at a random selection of starting points. Such a *multistart* approach is still a good idea. Taking a set of local minima obtained from different starting points, one might try a number of strategies of starting from each of the local minima and conducting a random search in hyperspheres around each to see whether better minima might be obtained, and so on. Any simulated annealing approach will be, effectively, a random search on a set much smaller than the entire feasible region of the parameter space. We should despair, in general, of coming up with a fool-proof method for finding a global optimum that will work with any and all continuous functions. Much of the supposed success of simulated annealing, as opposed to the kind of multistart algorithm, is probably a result of the very fast computers that simulated annealers tended to have available.

11.5 Exploration and Estimation in High Dimensions

The power of the modern digital computer enables us realistically to carry out analysis for data of higher dimensionality. Since the important introduction of exploratory data analysis in the 1970s, a great deal of effort has been expended in creating computer algorithms for visual analysis of data. One major advantage of EDA compared to classical procedures is a diminished dependency on assumptions of normality. However, for the higher-dimensional situation, visualization has serious deficiencies, because it tends to involve projection into two or three dimensions.

What are typical structures for data in high dimensions? This is a question whose answer is only very imperfectly understood at the present time. Some possible candidates are:

1. Gaussian-like structure in all dimensions.
2. High signal-to-noise ratio in only one, two, or three dimensions, with only noise appearing in the others. Significant departures from Gaussianity.
3. System of solar systems. That is, clusters of structure about modes of high density, with mostly empty space away from the local modes.

4. High signal-to-noise ratio along curved manifolds. Again the astronomical analogy is tempting, one appearance being similar to that of spiral nebulae.

For structure 1, classical analytical tools are likely to prove sufficient. For structure 2, EDA techniques, including nonparametric function estimation and other nonparametric procedures will generally suffice. Since human beings manage to cope, more or less, using procedures which are no more than three- or four-dimensional, it might be tempting to assume that structure 2 is somehow a natural universal rule. Such an assumption would be incredibly anthropomorphic, and we do not choose, at this juncture, to make it. For structure 3, the technique investigated by Thompson and his students [2, 6-8] is the finding of modes, utilizing these as base camps for further investigation. For structure 4, very little successful work has been done. Yet the presence of such phenomena as diverse in size as spiral nebulae and DNA shows that such structures are naturally occurring. One way in which the astronomical analogy is deceptively simple is that astronomical problems are generally concerned with relatively low dimensionality. By the time we get past four dimensions, we really are in *terra incognita* insofar as the statistical literature is concerned. One hears a great deal about the "curse of dimensionality." The difficulty of dealing with higher-dimensional non-Gaussian data is currently a reality. However, for higher-dimensional Gaussian data, knowledge of data in additional dimensions provides additional information. So may it also be for non-Gaussian data, if we understood the underlying structure.

Here, we are concerned mainly with structure 3. Mode finding is based on the mean update algorithm (MUA) [2, 6, 7, 18]:

Mean Update Algorithm

Let $\hat{\mu}_1$ be the initial guess

Let m be a fixed parameter;

$i = 1$;

Repeat until $\hat{\mu}_{i+1} = \hat{\mu}_i$;

Begin

Find the sample points $\{X_1, X_2, \dots, X_m\}$ which are closest to $\hat{\mu}_i$;

Let $\hat{\mu}_{i+1} = \frac{1}{m} \sum_{j=1}^m X_j$;

$i = i + 1$;

end.

Let us consider a sample from a bivariate distribution centered at (0,0). The human eye easily picks the (0,0) point as a promising candidate for the "location" of the distribution. Such a Gestaltic visualization analysis is not as usable in higher dimensions. We will be advocating such an automated technique as the mean update algorithm. Let us examine Figure 11.6. Suppose that we have only one dimension of data. Starting at the projection of $\mathbf{0}$ on the x -axis, let us find the two nearest neighbors on the x -axis.

Taking the average of these, brings us to the 1 on the x -axis. And there the algorithm stalls, at quite a distance from the origin.

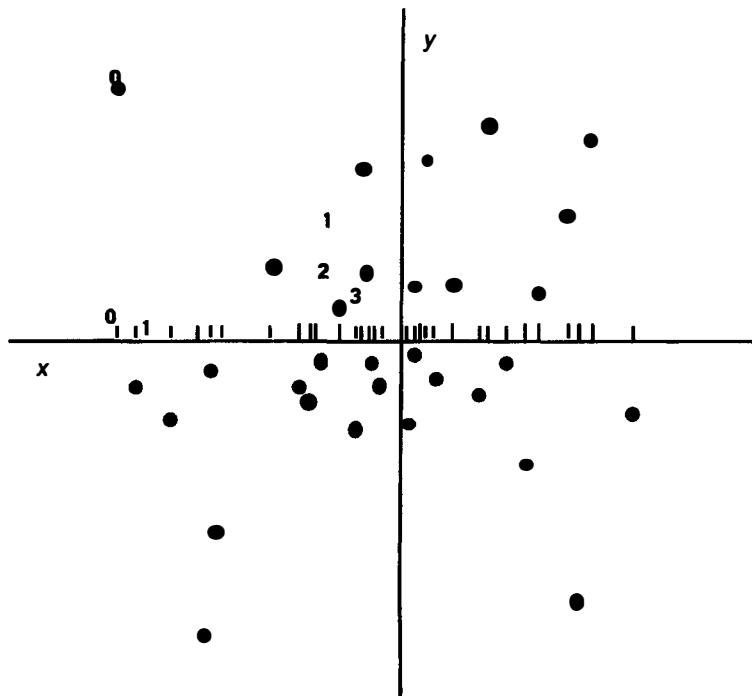


Figure 11.6. Mean update estimation of mode.

However, if we use the full two dimensional data, we note that the algorithm does not stall until point 3, a good deal closer to the origin. So increased dimensionality need not be a curse. Here, we note it to be a blessing.

Let us take this observation further. Suppose we are seeking the location of the minor mode in a data set which (unbeknownst to us) turns out to be

$$f(x) = .3\mathcal{N}(x; .051, \mathbf{I}) + .7\mathcal{N}(x; 2.4471, \mathbf{I}). \quad (11.8)$$

If we have a sample of size 100 from this density and use the mean update algorithm, we can measure the effectiveness of the MUA with increasing dimensionality using the criterion function

$$\text{MSE}(\hat{\mu}) = \frac{1}{p} \sum_{j=1}^p (\hat{\mu}_j - \mu_j)^2. \quad (11.9)$$

Below we consider numerical averaging over 25 simulations, each of size 100.

Table 11.2. Mean Square Errors

<i>p</i>	<i>m</i>	MSE
1	20	0.6371
3	20	0.2856
5	20	0.0735
10	20	0.0612
15	20	0.0520

We note in Table 11.2, how, as the dimensionality increases, essentially all of the 20 nearest neighbors come from the minor mode, approaching the idealized MSE of 0.05 as p goes to ∞ . Subsequent work [6] has shown that for multiple modes in dimensions five and over, the MUA appears to find, automatically, which points to associate with each mode, so that even for mixtures of rather taily distributions such as $T(3)$, we come close to the idealized MSE for the location of each mode, namely, $1/(np)$ where n is the total sample size and p is the proportion of the data coming from the mode. So far from being a curse, an increasing dimensionality can be an enormous blessing. We really have no very good insights yet as to what happens in, say, 8-space. This examination of higher-dimensional data is likely to be one of the big deals in statistical analysis for the next 50 years. The examination of data in higher-dimensions is made possible by the modern computer. If we force ourselves, as is currently fashionable, to deal with higher-dimensional data by visualization techniques (and hence projections into 3-space) we pay an enormous price and, quite possibly, miss out on the benefits of high-dimensional examination of data.

The analogy we shall employ is that of moving through space until we find a “center” of locally high density. We continue the process until we have found the local modes for the data set. These can be used as centers for local density estimation, possibly nonparametric, possibly parametric (e.g., locally Gaussian). It turns out, as we shall see, that finding local modes in high dimensions can be achieved effectively with sample sizes orders of magnitude below those generally considered necessary for density estimation in high dimensions [14–16]. Moreover, as a practical matter, once we have found the modes in a data set, we will have frequently gleaned the most important information in the data, rather like the mean in a one-dimensional data set.

Let us suppose that we have, using each data point from the data set of size n as a starting point, found mm apparent local modes. As a second step, let us develop an algorithm for consolidating the apparent local modes to something more representative of the underlying distribution. There are many ways to carry out the aggregation part of the algorithm. This is only one of the possibilities.

Take two of the local modes, say M_1 and M_2 . Examine the volume $V_{1,m}$ required to get, say, m nearest neighbors of M_1 and $V_{2,m}$ required to get,

say, m nearest neighbors of M_2 . Standing at the midpoint between M_1 and M_2 , say, $M_{1,2}$, draw a sphere of volume $V_{12,m} = V_{1,m} + V_{2,m}$. Suppose that the number of distinct points in the pooled clouds is m_{12} . Suppose that the hypersphere centered at $M_{1,2}$ has a density as high as that in the other two clouds. Let the number of points falling inside the hypersphere centered at $M_{1,2}$ be $n_{1,2}$. Then if the number of data points falling inside that hypersphere is greater than the 5th percentile of a binomial variate of size m_{12} and with $p = 0.5$, we perform a condensation step by replacing the two modes M_1 and M_2 by $M_{1,2}$ as shown in Figure 11.7.

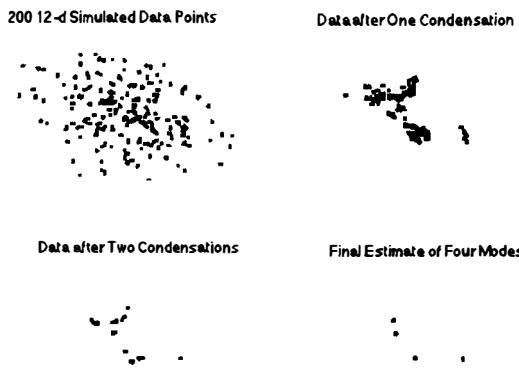


Figure 11.7. Condensation progression using MUA.

To examine the progression of the condensation algorithm, we simulate 200 data points from a mixture of four 12-dimensional normal distributions with mixture weights .40, .24, .19, and .17. The four modes are well estimated both in terms of numerosity and location even for such a small data set.

Let us apply Elliott's version [6, 7] of the MUA to the much-studied Fisher–Anderson iris data. This is a database of three varieties of iris with 50 observations from each of the varieties. The algorithm found four (see Table 11.4) rather than the hoped-for three clusters shown in Table 11.3.

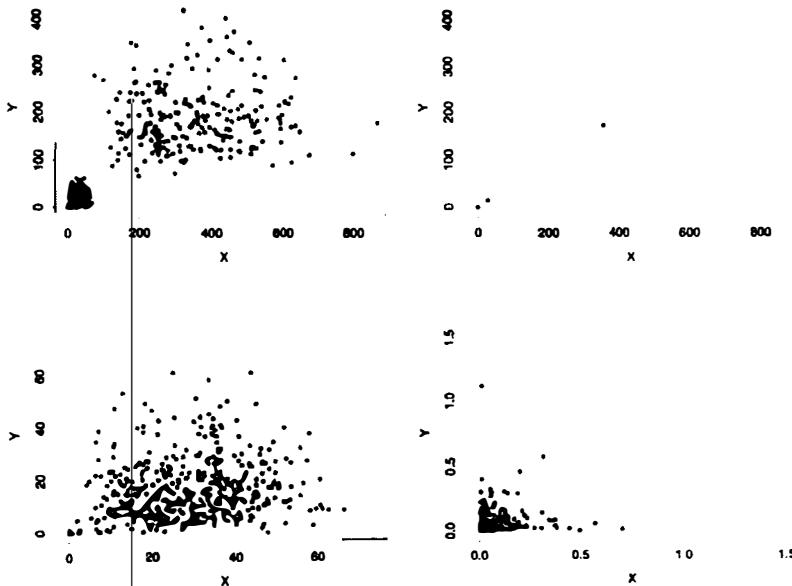
Table 11.3. Fisher–Anderson Iris Data.

Species	Sepal Length	Sepal Width	Petal Length	Petal Width
<i>Setosa</i>	5.006	3.428	1.462	0.246
<i>Versicolor</i>	5.936	2.770	4.260	1.326
<i>Virginica</i>	6.588	2.974	5.552	2.026

Table 11.4. Estimated Modes Based on 150 Observations

Species	Sepal Length	Sepal Width	Pet. Lgth	Pet. Width
<i>Setosa</i>	4.992	3.411	1.462	.225
<i>Versicolor</i>	5.642	2.696	4.101	1.267
<i>Virginica</i>	6.762	3.067	5.589	2.218
<i>Versi/Virgin.</i>	6.249	2.893	4.837	1.591

At first, we feared that some fundamental flaw had crept into the algorithm, which had always performed quite predictably on simulated data. Later, it seemed plausible, based on the fact *Virginica* and *Versicolor* always spontaneously hybridize and that it is almost impossible not to have this hybrid present, to believe our eyes. The fourth mode had occurred, in fact, almost precisely at the mean of the *Virginica* and *Versicolor* modes. This was a rather surprising result, apparently unnoticed in the some fifty years since the Fisher–Anderson iris data became something of a test bed for measuring the effectiveness of discrimination and clustering algorithms.

**Figure 11.8. Mode finding when scales of underlying mixtures are very different.**

In another application, Elliott and Thompson [6] have examined a four dimensional ballistics data set of size 944 kindly provided by Malcolm Taylor of the Army Research Laboratory. Consider the two-dimensional projections displayed in Figure 11.8. Our algorithm was able to find modes from

overlapping subpopulations whose scales differed by nearly 1000. We see in the top left quadrant of Figure 11.8 a two-dimensional projection of the data set. The top right quadrant gives the three estimated modes. In the lower left quadrant, we have zoomed in on the cluster in the lower left of the data set. In the lower right quadrant, we have zoomed in to a scale 10^{-3} of that used in the display of the raw data. As we have seen, even in data sets of dimensionality as low as four, there seems to be appearing a big bonus for the extra dimension(s) past three for finding modes.

Mean update algorithms show great promise for exploratory purposes. The problem of nonparametric function estimation is one to which some of us at Rice University have given some attention for a number of years. Our foray into the higher dimensions has produced a number of surprises. The notion that increasing dimensionality is a “curse” seems only to be true if we insist on graphical approaches. Our multidimensional mode-finding algorithm dramatically improves with increasing dimensionality.

Problems

11.1. In optimization examples, perhaps the easiest problem is that of finding the minimum of the dot product. Consider finding the minimum of

$$J_1(\Theta) = \Theta_1^2 + \Theta_2^2 + \Theta_3^2 + 2\epsilon$$

where ϵ is $\mathcal{N}(0, 1)$. Examine the performance of both the Nelder–Mead and Box–Hunter algorithms.

11.2. A somewhat more difficult minimization case study is that of the Rosenbrock function with additive Gaussian noise

$$J_2(\Theta) = 100(\Theta_1^2 - \Theta_2)^2 + (1 - \Theta_1)^2 + 1 + \epsilon,$$

where ϵ is $\mathcal{N}(0, 1)$. Examine the performance of both the Nelder–Mead and Box–Hunter algorithms.

11.3. Returning to the problem in Section 5.4.1, generate a set of times of discovery of secondary tumor (time measured in months past discovery and removal of primary) of 400 patients with $a = .17 \times 10^{-9}$, $b = .23 \times 10^{-8}$, $\alpha = .31$, and $\lambda = .0030$. Using SIMEST, see if you can recover the true parameter values from various starting values, using the Box–Hunter algorithm.

11.4. Consider the density function

$$f(x) = 0.5\mathcal{N}(x; 0.551, \mathbf{I}) + 0.3\mathcal{N}(x; 21, \mathbf{I}) + 0.2\mathcal{N}(x; 21, \mathbf{I})$$

Generate random samples of size 100 for dimensions 2, 3, 4, 5, and 10. Examine the efficacy of the MUA in finding the centers of the three distributions.

References

- [1] Bohachevsky, I. G., Johnson, M. E., and Stein, M. L. (1988). "Generalized simulated annealing for function optimization," *Technometrics*, 209-217.
- [2] Boswell, S. B. (1983). *Nonparametric Mode Estimation for Higher Dimensional Densities*. Doctoral dissertation. Houston: Rice University.
- [3] Box, G. E. P. and Draper, N. R. (1969). *Evolutionary Operation*. New York: John Wiley & Sons.
- [4] Box, G.E.P. and Draper, N. R. (1989). *Empirical Model-Building and Response Surfaces*. New York: John Wiley & Sons.
- [5] Box, G.E.P. and Hunter, J. S. (1957). "Multifactor experimental designs for exploring response surfaces," *Ann. Statist.*, 28, 195-241.
- [6] Elliott, M. N. and Thompson, J. R. (1993). "The nonparametric estimation of probability densities in ballistics research," *Proceedings of the Twenty-Sixth Conference of the Design of Experiments in Army Research Development and Testing*. Research Triangle Park, N.C.: Army Research Office, 309-326.
- [7] Elliott, M. N. (1995). *An Automatic Algorithm For The Estimation of Mode Location and Numerosity in General Multidimensional Data*. Doctoral dissertation. Houston: Rice University.
- [8] Fwu, C., Tapia, R. A., and Thompson, J. R. (1981). "The nonparametric estimation of probability densities in ballistics research," *Proceedings of the Twenty-Sixth Conference of the Design of Experiments in Army Research Development and Testing*. Research Triangle Park, N.C.: Army Research Office, 309-326.
- [9] Lawera, M. and Thompson, J. R. (1993). "A parallelized, simulation based algorithm for parameter estimation," *Proceedings of the Thirty-Eighth Conference on the Design of Experiments in Army Research Development and Testing*, B. Bodt, ed. Research Triangle Park, N.C.: Army Research Office, 321-341.
- [10] Nelder, J. A. and Mead, R. (1965). "A simplex method for exploring response surfaces," *Comput. J.*, 308-313.
- [11] Poisson, S. D. (1837). *Recherches sur la probabilité des jugements en matière criminelle et en matière civile, précédées des règles générales du calcul des probabilités*. Paris.
- [12] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P.

(1992). *Numerical Recipes in FORTRAN*. New York: Cambridge University Press, 443–448.

[13] Schwalb, O. (1999). *Practical and Effective Methods of Simulation Based Parameter Estimation for Multidimensional Data*. Doctoral dissertation. Houston: Rice University.

[14] Scott, D. W. (1992). *Multivariate Density Estimation*. New York: John Wiley & Sons, 199.

[15] Scott, D. W. and Thompson, J. R. (1983). “Probability density estimation in higher dimensions,” in *Computer Science and Statistics*, J. Gentle, ed. Amsterdam: North Holland, 173–179.

[16] Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. London: Chapman & Hall.

[17] Thompson, J. R. and Koronacki, J. (1993). *Statistical Process Control for Quality Improvement*. New York: Chapman & Hall, 238–245.

[18] Thompson, J. R. and Tapia, R. A. (1990). *Nonparametric Function Estimation, Modeling, and Simulation*. Philadelphia: SIAM, 167–168.

[19] Wald, A. (1949). “Note on the consistency of the maximum likelihood estimate,” *Ann. Math. Statist.*, 595–601.

Chapter 12

Utility and Group Preference

12.1 Introduction

If our analysis consisted simply in trying to see what the dollar value of a stock might be six months from today, that would be difficult enough. But analysis of the market is complicated by the fact that it is not simply the dollar value that is of interest. A thousand dollars is not necessarily of the same subjective value to investor A and investor B. It is actually this difference in personal utilities which helps create markets.

Narratives concerning the differing values of the same property to individuals of differing means go back into antiquity. Around 1035 B.C., the Prophet Nathan told about a rich man with many flocks who slaughtered the sole lamb of a poor man rather than kill one of his own (2 Samuel 1:12). King David was enraged to hear of the deed and promised harsh justice to the evil rich man. One lamb means more to a poor man than to a man with many lambs. (Of course, then Nathan dropped the punch line which had to do with David having had Uriah killed, so that David could gain access to Uriah's wife Bathsheba. The parable was about King David himself.) In the New Testament Jesus tells of a poor widow whose gift to the Temple of two mites, a small multiple of lowest currency in the realm, had moral value more than the magnificent gifts of the very wealthy, since the widow had given away everything she had.

All this more or less resonates with us as a matter of common sense. We understand that the gain or loss of a small amount of property means much more to a poor person than it does to a wealthy one. Although the effect of present wealth on the utility of the gain of a certain amount of money

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

had been clear for millenia, it seems that no attempt had been made until 1738 A.D. to come up with a quantitative measure of the relationship between present wealth and marginal gain. In that year, the Swiss proto-statistician Daniel Bernoulli published "Exposition of a New Theory on the Measurement of Risk" [3]. In this paper Bernoulli laid the basis of utility theory. Before starting our discussion of Bernoulli's paper, we should note that utility theory today in market modeling does not have quite the relevance one might have supposed in the golden age of utility theory, which ended some time before 1970. To many, a dollar is a dollar, regardless of one's wealth. That is unfortunate, for as we shall emphasize repeatedly in this text, markets are made by diverse individuals who view the same commodity or security quite differently from one another. Prices are most efficiently determined by a jostling of buyers and sellers who come together with different notions of utility and arrive at a trading price by agreement.

When governments or agencies attempt to enforce fair pricing, disaster is generally the result, even when the government is really high-minded and has worthy goals. As an extreme example, during the general uprising of the population of Warsaw against the Nazis in August of 1944, things were going rather badly for the Poles. Even clean water was all but impossible to obtain. Some industrious peasants from the suburbs who had access to water carrier wagons loaded up from the family wells and, risking life and limb, delivered water to the freedom fighters for a price well above the peace time price of fresh water. What was a fair price for water delivered to patriots fighting to the death for the freedom of the nation? The freedom fighter high command decided it was zero and had the water vendors shot. Of course this marked the end of fresh water, but at least a politically correct fair price for water for freedom fighters had been enforced. Fresh water was both free and unavailable. "Nothin' ain't worth nothin', but it's free."

12.2 The St. Petersburg Paradox

Many are familiar with the following apparently paradoxical question:

How much should you be willing to pay to play a game in which a coin is repeatedly tossed until the first heads? If the first heads appears on the first toss, you will receive $2^1 = 2$ dollars. If the first heads appears on the second toss, you will receive $2^2 = 4$ dollars. If the first heads appears on the k th toss, you receive 2^k dollars. The game terminates on the round where the first heads is obtained.

Now, on the average the expectation of the pay-off in this game is

$$V = \sum_{k=1}^{\infty} 2^k \left(\frac{1}{2}\right)^k = 1 + 1 + \dots = \infty, \quad (12.1)$$

The expected payoff in the game is infinity. Would anybody pay a million dollars to play this game? Very likely, the answer is negative. Possibly somebody who was already incredibly rich might do it, for the chances of tossing 19 straight tails is small (one in $2^{19} = 524,288$). A poor person might be unwilling to pay more than two dollars, the minimum possible pay-off of the game. There is, it would appear, a relationship between one's wealth and the amount one would pay to play this paradoxical game. But things are more complicated than that. It really is the case that simply looking at the expected value of the game does not tell the whole story unless the game is replayed a very large number of times.

As we shall discuss later on, it is customary to talk about the expected value of an investment and also about its volatility (standard deviation). But the fact is that these two numbers generally will not give an investor all the information to determine whether an investment is attractive or not. Now, for the St. Petersburg game, we know that the probability that the player will realize at least $(\$2)^{k+1}$ is given (for $k = 1, 2, \dots$) by¹

$$1 - \sum_{j=1}^k \left(\frac{1}{2}\right)^j = \left(\frac{1}{2}\right)^k. \quad (12.2)$$

The player cannot make less than \$2. But what is the probability that he will make at least, say, \$1024? From (12.2), the answer is $1/2^9 = 1/512 = .001953$. This would be roughly two chances in a thousand. Probably, the player would dismiss such an event as being very unlikely. In Figure 12.1, we give a profile showing the probabilities that the player will make at least various amounts of dollars.

The probability profile gives a reasonable insight as to why the St. Petersburg game is not worth the expectation of payoff, in this case ∞ .² We believe that, for most players, most of the time, it is advisable to look at the entire probability profile when deciding whether an investment suits the investor. Most investors will not be as impressed with the fact that the expected value of the game is infinity as they would with the fact that in only two chances out of a thousand will the game produce winnings in excess of \$1000. Looking at the overall picture, half the time, a player will make at least \$4. He will never make less than \$2. One-fourth of the time, he will make at least \$8. One-eighth of the time he will make at least \$16, etc. In deciding how much he will wager to play the game, the player should have the entire probability profile at his disposal, not simply the mean and standard deviation. At the end of the day, the player must decide how much he is willing to pay to play. In other words, he must combine the entire probability profile into his decision, "Yea or nay." It is tempting to

¹Here we are using the fact that a series of the form $1 + r + r^2 + r^3 + \dots + r^n = (1 - r^n)/(1 - r)$ if r is greater than 0 and less than 1.

²It is interesting to note, of course, that no gambling house exists which could pay ∞ !

create a *deus ex machina* which will automatically decide how much one should pay to play the game. The expected value is such a rule, and we have seen that it does not work.

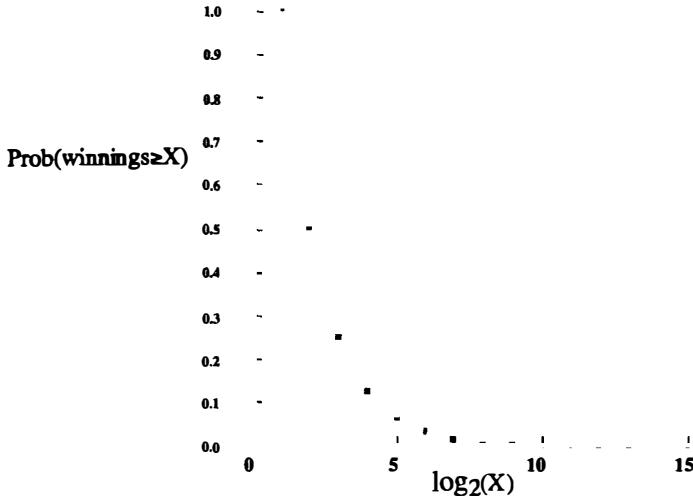


Figure 12.1. Probability profile of $\log_2(\text{winnings})$.

The situation where the *utility function* U is the dollar value of game payoff is a special case of the class of reasonable utilities. Here, the utility is *linear* in dollar pay-off, i.e., the relation between utility U and payoff X is given by

$$U(X) = a + BX. \quad (12.3)$$

We note that if we used any function which grows more slowly than X , then, for the coin flipping example under consideration,

$$E(U(X)) = V = \sum_{k=1}^{k=\infty} U(2^k) \left(\frac{1}{2}\right)^k \quad (12.4)$$

is finite. For example, suppose we consider $U(X) = \sqrt{X}$, then we have

$$E(U(X)) = \sum_{k=1}^{k=\infty} \sqrt{2^k} \left(\frac{1}{2}\right)^k \quad (12.5)$$

$$= \sum_{k=1}^{k=\infty} 2^{-k/2} \quad (12.6)$$

$$= \frac{1}{\sqrt{2}} \left[1 + \frac{1}{\sqrt{2}} + \frac{1}{2} + \frac{1}{2^{3/2}} + \dots \right] \quad (12.7)$$

$$= \frac{1}{\sqrt{2}} \frac{1}{1 - 1/\sqrt{2}} \quad (12.8)$$

$$= \frac{1}{\sqrt{2} - 1}$$

Another popular utility function is the logarithm $U(X) = \log(X)$. Here, for the coin-flipping problem ³

$$\begin{aligned}
 E(U(X)) &= \sum_{k=1}^{k=\infty} \log((2^k)) \left(\frac{1}{2}\right)^k & (12.9) \\
 &= \log(2) \sum_{k=1}^{k=\infty} k \left(\frac{1}{2}\right)^k \\
 &= 4 \log(2).
 \end{aligned}$$

Again, it should be emphasized that for most investors, looking at the expected utility will be at best a poor substitute for looking at the entire probability profile of the utilities. We consider such a profile in Figure 12.2 where probability of making at least a value of utiles as those given on the abscissa is plotted for the case where $\log_e(\text{winnings})$ is the utility. We note that we have the same probability masses as in Figure 12.1. Only the abscissa axis has changed, since we are looking at $\log_e(\text{winnings})$ as opposed to $\log_2(\text{winnings})$ (where $e = 2.7183$). Most investors could work just as well with Figure 12.1 as with Figure 12.2, even if their utility function were \log_e . We really can get into trouble if we do not look at the entire probability profile (which could appropriately also be referred to as a *risk profile*). If we decide to make our decision based on one summary number, such as the expected value of the utility, we can quite easily lose our grasp of what the numbers are telling us.

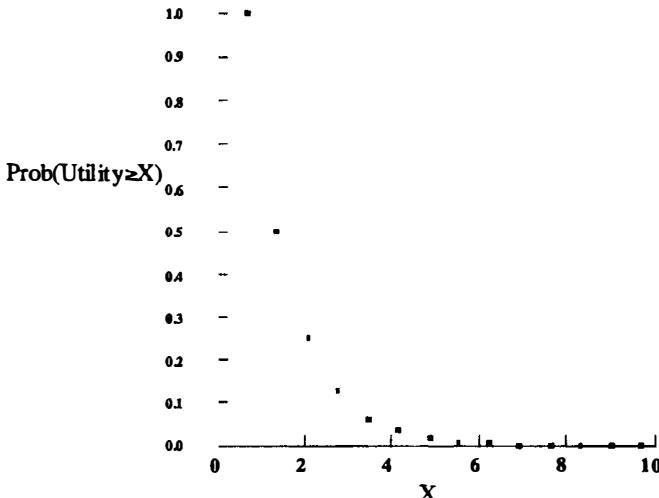


Figure 12.2. Probability profile of $\log(\text{winnings})$.

Generally speaking, reasonable candidates for utility should be nondecreasing in capital. That is, by increasing one's capital, one's utility does

³Here we use the fact that $1/(1 - y)^2 = \sum_{k=1}^{k=\infty} ky^k$.

not decrease. There are, of course, exceptions to this rather reasonable assumption. As with most exceptions to rationality, these are generally state imposed. For example, during the Socialist government of Olaf Palme in Sweden, graduated tax rates actually got over 100%. In Lenin's war to break the power of free agriculture in Ukraine, "kulaks" (*kulak* means "fist") were defined to be peasants with over a certain modest amount of holdings. If one had a bit less than the boundary value, one was (for the time) left alone. But over the value, one was shot. Similarly, in the Red controlled areas of Spain during the Civil War (1936–1939), there was a critical boundary between who was a peasant and who was a "landlord."

The marginal utility of a new increment or decrement of wealth is clearly a personal matter related to one's wealth. The loss of ten thousand dollars is trivial to a wealthy person. To someone in the lower income brackets, it may be ruinous.

"A billion here, a billion there," can be easily be doled out by an American politician at the national level. An incremental dollar has utility dependent, somehow, on the assets of the person(s) considered. This rather obvious fact can be used for many purposes, including graduated taxation. Here we simply wish to consider the matter from the standpoint of its practical implications.

One can argue, as did Bernoulli, that the marginal increase in a person's wealth by a profit should be measured as a ratio of the new profit to the assets already in hand before the profit was realized. Expressing this in symbols, where U is the utility, k is a constant of proportionality, X is the base amount of wealth, and ΔX is the change in that wealth, one can write⁴

$$\Delta U = k \frac{\Delta X}{X}. \quad (12.10)$$

This, then, says that an increase in wealth of one dollar changes the utility of a person with beginning wealth of \$100 the same as an increase in wealth of \$1,000 has for a person with beginning wealth of \$100,000. This "law" of Bernoulli's is, of course, not really a law but rather an assumption with shortcomings. For example, if both of these hypothetical persons have a child being held for a ransom of \$101,000, then the wealthier individual can buy his child's freedom, whereas the poorer one is as far away from achieving the goal with \$101 as with \$100. On the other hand, if the ransom is \$101, then the poorer person's utility goes up much more with the addition of one dollar than that of the richer one with the addition of \$1,000. These are both examples of "step function" utilities, and Bernoulli wanted to look at a smooth utility. Objections can be raised that utility functions which have critical jumps up or down are not realistic. The thousands of bankruptcies experienced in the United States yearly would seem to be an example of step function realities.

⁴If the initial capital is X , then (2.9) is satisfied by $U(X, \Delta X) = \log[(X + \Delta X)/X]$.

Having noted that Bernoulli discovered an insight rather than a law, we must concede that his insight was valuable. Generally speaking, when starting to understand a new concept, it is good to try and reason from an example, even a hypothetical one. Bernoulli gave a hypothetical example based on Caius, a fictitious merchant of St. Petersburg, Russia, who was contemplating whether he should take insurance on a shipment from Amsterdam to St. Petersburg. The shipment, upon delivery, provides Caius with 10,000 rubles. But storms are such that, on the average, 5% will be lost at sea. The Amsterdam underwriters want a full covered policy payment of 800 rubles, or 8% of the profit if no mishap occurs. Should Caius buy the policy? His expected value if he does not is 9,500 rubles. The underwriters are clearly demanding a premium of 300 rubles above their expected payout. Is it worth it for Caius to purchase the policy or to "self-insure"? If we go by naive ruble values, then he should self-insure. But if Caius follows Bernoulli's advice, he should buy the policy if his expected utility for insuring is greater than that for not insuring. Let us suppose Caius's capital is X rubles. We shall, without loss of generality, take $k = 1$, since any k will change both the insured and self-insured options in the same way. Then Caius's expected utility of not insuring minus that of insuring is given by

$$f(X) = .95 \log \left[\frac{X + 10,000}{X} \right] - \log \left[\frac{X + 9,200}{X} \right]. \quad (12.11)$$

Setting $f(X) = 0$, we can easily solve the resulting equation using Newton's Method (see Appendix B at the end of this book).

$$X_{n+1} = X_n - \frac{f(X_n)}{f'(X_n)}. \quad (12.12)$$

Starting with 5000 rubles as our first guess, we arrive at the indifference value of 5042 rubles. If Caius has less than this amount, he should (according to Bernoulli), buy the 800 ruble policy. If he has more, he should self insure (i.e., not buy the insurance).

Next, let us ask the question as to how much the underwriter (insurer) should have in hand in order to sell the 10,000 ruble policy for the amount of 800 rubles. Let Y be the assets of the underwriter. Then he should sell the policy if his assets exceed the Y value in

$$g(Y) = .95 \log \left[\frac{Y + 800}{Y} \right] + .05 \log \left[\frac{Y - 9,200}{Y} \right]. \quad (12.13)$$

Again, using Newton's method, we find that if the underwriter has a stake of 14,242 rubles or more, Bernoulli's rule tells us that selling the policy will improve the underwriter's position. Going further, let us ask the question as to what is the minimum price the underwriter might reasonably sell the

10,000 ruble policy if the underwriter has capital of one million rubles. To achieve this, we simply solve

$$h(W) = 0.95 \log \left[\frac{10^6 + W}{10^6} \right] + 0.05 \log \left[\frac{10^6 + W - 10^4}{10^6} \right] \quad (12.14)$$

The solution here tells us that the underwriter with capital of 1,000,000 rubles might reasonably sell the 10,000 ruble policy for 502.4 rubles. But Caius, if he has capital of 5,042 rubles, say, would find it reasonable to pay up to 800 rubles. Thus, the underwriter is in a position to get more than his own indifference value for the trade (502.4). Naturally, the underwriter is largely looking at things from his own standpoint, rather than that of potential clients. He is only interested in the amount of the policy, the risk to the underwriter, and charging whatever rate the market will bear. If the underwriter has no competition, he must remember that there is always competition from the merchant himself who can decide to self insure. This is simply a manifestation of substitutability of one service by another. Still, there is the likelihood that the presence of a second underwriter in the Amsterdam to St. Petersburg run will drive prices downward. There is a spread of several hundred rubles where the purchase of the policy is a good deal, from a utility standpoint, for both underwriter and the insured merchant.

Note that in the above example, if the underwriter sells Caius a policy for 650 rubles, then Caius has a good deal. He has the policy for less than his utility function would make him willing to pay. And the underwriter has a good deal also, for he is well above the minimum rate his utility function dictates. Here is an example of the reason transactions take place, for the deal is good for both parties from the standpoints of their respective utilities. In a true free trade situation where there are a number of merchants and several underwriters, there will be a jostling back and forth of rates. The change in riskiness of the transit due to weather, war, pirates, etc., will be a major driver in the change of rates. They will never stop their fluctuation, though at any given point in time, the cost of a 10,000 ruble policy will be similar from all of the underwriters. It is the difference in utility functions (driven in part by wealth status) as well as the difference in personal views as to the value of a commodity or service that cause markets to exist at all. If Caius were willing to buy the policy for no more than 800 rubles and the underwriter were only willing to sell it for 850 rubles, then no transaction would take place. By jostling and haggling, any given selling price will be in the "comfort intervals" of both buyer and seller. The buyer would always have been willing to have paid a bit more than he did, the seller always to have taken a bit less.

Human nature being what it is, there will be an attempt by the underwriters to combine together to set the rates at an unnaturally high level. They may even decide "for the good of the public" to have the Czar set the rates. These will tend to be rates determined by the St. Petersburg

Association of Underwriters, and thus on the high side, On the other hand, the Patriotic Association of Merchants will try and get the rates lowered, particularly for favored clients. But the creation of fixed rates will indeed "stabilize" the market by fixing the rates. Statist intervention, then, is really the only means of arriving at a "stable" market. Otherwise, the rates will fluctuate, and there will always be a market opportunity for an insurance agent continually to negotiate with all the underwriters to obtain better deals for the merchants the agent represents.

The presence of variation in market prices over time is woefully misunderstood. So far from being evidence of chaos in the market, variation represents a stabilizing continuum of adjustments to the ever changing realities of the market, including the goals, expectations and situations of the persons and institutions taking part in the market. One example of the destabilizing effects of statist efforts to "stabilize" markets is rationing. Another is the fixing of prices at arbitrary levels. In "People's" Poland before the Communists lost power on June 4, 1989, the retail price of milk was set by the government at one generally below the cost of production. There was little collectivization in Poland, so the state was forced to buy milk from independent farmers. It could, naturally, have paid the farmers something greater than the cost of production by using state funds. It chose, rather, to force each farmer to deliver a certain quantity of milk to the state at a price which was generally below the cost of production. Naturally, the milk sold to the state had a water content considerably in excess of that of milk which comes straight from the cow. The price stabilized milk that was made available to the populace had some similarities with milk, but was something rather different, and *it varied greatly in quality from day to day, from location to location*. On the other hand, the state turned more or less a blind eye to farmers selling a portion of their milk on the black (aka free) market. Occasionally (once every several years) a "speculator" was shot, but this was pro forma. Even under Russian control, some realism was ever present in the European satellites. Black market milk was expensive, and, since it was produced and vended under irregular conditions, the quality was variable. Those who could afford to do so generally would strike a deal with one particular farmer for deliveries of milk on a regular basis. Most city dwellers, however, were stuck with price stabilized "milk" for the fifty years of Soviet occupation.

There is much to fault with Bernoulli's treatment of utility. First of all, we can observe that the conditions of trade he posed were somewhat strange. Generally speaking, Caius would have to buy the goods for shipment. A more accurate way to pose the problem, perhaps, would be one in which Caius has acquired goods in Amsterdam for which he paid, say, 6000 rubles. He can sell the goods in St. Petersburg for 10,000 rubles. When he buys his insurance policy, he may well have to settle for insuring at the amount of purchase. Caius may own his own ship, in which case, insuring the value of the ship is another consideration. And so on. But Bernoulli

has revealed a most important point: the marginal value of a dollar varies depending on the financial status of the person involved. Such differences in utility from person to person should not pass unnoticed. It is one (though not the only one) reason that markets exist at all. Another problem with Bernoulli's treatment is that everything is based on the expected value of the utility. In order for an investment to have low risk, we need to know that the probability of a large loss is small. The risk profile cannot be captured by the expected utility or any other single number.

12.3 von Neumann–Morgenstern Utility

It is generally not a very good idea to assume a utility function for a particular individual concerning a set of possible transactions contemplated by an individual. A financial advisor who simply assumed, say, a logarithmic or square root relation between various returns and the utility of a client is making an unnecessary assumption. Utility being a matter of personal assessment, it is frequently possible to come up with a series of questions which would enable us to extract the implied utility of the various choices. von Neumann and Morgenstern [6] have given us a set of axioms which should be satisfied by the preferences of a rational person. These are

1. **Transitivity.** If the subject is indifferent between outcomes A and B , and also between B and C , he must be indifferent between A and C . Symbolically, $A \asymp B$ and $B \asymp C \Rightarrow A \asymp C$.
2. **Continuity of preferences.** If A is preferred to B and B is preferred to no change, then there is a probability α ($0 < \alpha < 1$), such that the subject is indifferent between αA and B .
3. **Independence.** If A is preferred to B , then for any probability α ($0 < \alpha < 1$) $\alpha A + (1 - \alpha)B$ is preferred to B . (An equivalent axiom says that if $A \succ B$, then $\alpha A \succ \alpha B$.)
4. **Desire for high probability of success.** If A is preferred to no change, and if $\alpha_1 > \alpha_2$, then $\alpha_1 A$ is preferred to $\alpha_2 A$.
5. **Compound probabilities.** If one is indifferent between αA and B , and if $\alpha = \alpha_1 \alpha_2$, then one is indifferent between $\alpha_1 \alpha_2 A$ and B . In other words, if the outcomes of one risky event are other risky events, the subject should act only on the basis of final outcomes and their associated probabilities.

Let us now go through the largely psychometric exercise for determining a subject's utility function and his/her willingness to accept risk. First of all, notice that all of this discussion abandons the world where utility is linear in dollars. We shall talk of a new currency, called *utes*. This is, by the way, an interpolation rule. We will not feel very comfortable about

extrapolating outside the interval where our client can answer the questions we shall pose. We will start, then, with a range of dollar assets, say \$0 to \$1,000,000. We need to define a utility nondecreasing in dollars. It turns out that our hypothetical client has a utility function equal to the square root of the amount of dollars (he does not necessarily realize this, but the answers to our questions will reveal this to be the case). We need to define the utility at the two endpoints. So, we decide (ourselves, without consulting yet the client) that $U(\$0) = 0$ utiles and $U(\$1,000,000) = 1,000$ utiles.

Q. How much would you be willing to pay for a lottery ticket offering a 50-50 chance of \$1,000,000 or \$ 0?

A. The client responds “\$250,000.” (Of course, we recognize the expectation in utile scale as being $.5 \times 0$ utiles + $0.5 \times 1,000$ utiles = 500 utiles.)

Q. How much would you pay for a 50-50 chance of \$250,000 or nothing?

A. The client responds “\$62,500” (which we recognize to be the expectation on the utile scale, $.5 \times 0$ utiles + 0.5×500 utiles = 250 utiles).

Q. How much would you pay for a 50-50 chance of \$62,500 or nothing?

A. The client responds “\$15,625” (which we recognize to be the expectation on the utile scale, $.5 \times 0$ utiles + $.5 \times 250$ utiles = 125 utiles).

The above analysis is consistent with the

Utility Maxim of von Neumann and Morgenstern. The utility of a game (risky event) is not the utility of the expected value of the game but rather the expected value of the utilities associated with the outcomes of the game.

In Figure 12.3, we give a spline smoothed plot using the Q&A. We note that if all we had was the three questions and their answers, we would see a plot virtually indistinguishable from what we know the functional relationship is between utiles and dollars, namely

$$U(X) = \sqrt{X}.$$

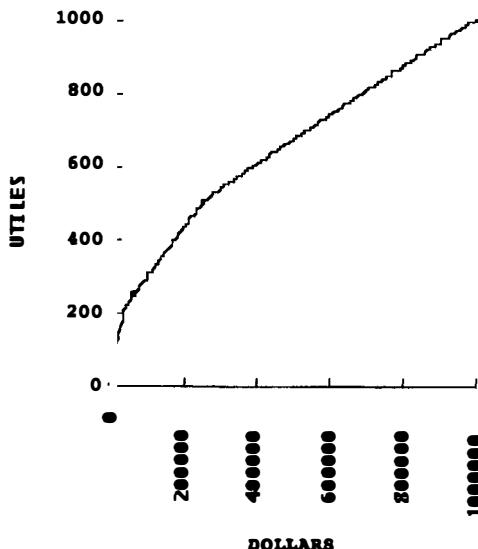


Figure 12.3. Empirically determined utility function.

Now, we recall that our client may not know that his utility is the square root of dollars. But his answers to our three questions give us a graph which is essentially equivalent to the square root.⁵

We can now ask the following question: How much is it worth to the client to receive the payoffs \$10,000 with probability .3, \$90,000 with probability 0.5, and \$490,000 with probability 0.2? Now, the naive answer would be that associated with the assumption that the utility of the client is simply the dollar amount:

$$E(X) = .3 \times 10,000 + 0.5 \times 90,000 + 0.2 \times 490,000 = 146,000.$$

But we have already determined that the utility function of the client is not dollars but the square root of dollars. This gives us:

$$E(U(X)) = .3 \times \sqrt{10,000} + .5 \times \sqrt{90,000} + .2 \times \sqrt{490,000} = 320 \text{ utiles.}$$

Going to Figure 12.3 (or recalling the $X = U^2$), we have that our client should consider a sure payment of $320^2 = \$102,400$ to be the value of the game to himself/herself.

Let us consider several scenarios each having expected utility value 500 utiles.

- A. A cash gift of $\$250,000 = \sqrt{250,000} = 500$ utiles. Here $E(X) = \$250,000$. The standard deviation of the dollar payout is $\sigma_A = \sqrt{(250,000 - 250,000)^2} = \0

⁵Economists derive the notion of "diminishing marginal utility" from functions of this sort.

- B. A game which pays \$90,000 with probability .5 and \$490,000 with probability .5. Here $E(U) = .5 \times \sqrt{90,000} + .5 \times \sqrt{490,000} = .5 \times 300 + .5 \times 700 = 500$ utiles. $E(X) = .5 \times \$90,000 + .5 \times \$490,000 = \$290,000$. $\sigma_B = \sqrt{.5(90,000 - 290,000)^2 + .5(490,000 - 290,000)^2} = \$200,000$.
- C. A game which pays \$90,000 with probability .5 and \$490,000 with probability .5. Here, $E(U) = .5 \times \sqrt{90,000} + .5 \times \sqrt{890,000} = 500$ utiles. $E(X) = .5 \times \$90,000 + .5 \times \$810,000 = \$400,000$. The standard deviation of the dollar output is given by $\sigma_C = \sqrt{.5(90,000 - 400,000)^2 + .5(810,000 - 400,000)^2} = \$400,000$.
- D. A game which pays \$1,000,000 with probability .5 and \$0 with probability .5. Here $E(U) = .5 \times 0 + .5 \times \sqrt{1,000,000} = 500$ utiles. $E(X) = .5 \times \$1,000,000 = \$500,000$. $\sigma_D = \sqrt{.5(1,000,000 - 500,000)^2 + .5(0 - 500,000)^2} = \$500,000$.

In Figure 12.3, we show the *indifference curve* for various games each having a value of 500 utiles.

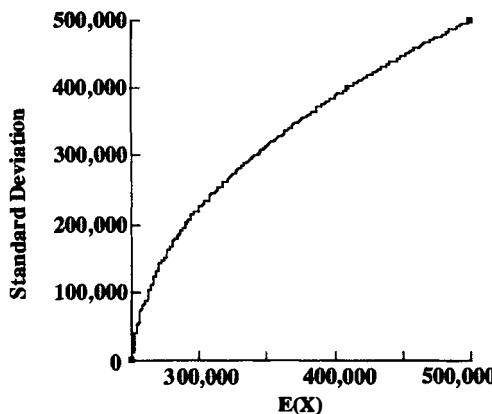


Figure 12.4. σ versus expected payoff with $E(U) = 500$ utiles.

We note the typical increase in expected payoff as the games become more risky. We note that the von Neumann-Morgenstern utility paradigm enables us to assess a rational choice system for a client who has no real notion of mathematical modeling, provided he or she can answer a few questions concerning indifference to choice between a sure thing and a set of particular games.

One major reason for the functioning of a market is that individuals and corporations will have different utilities for the same thing. Suppose that everybody valued a stock at \$100. What would be the incentive for the owner of such a stock to sell it at less than \$100? He might need money

for some other purpose. Perhaps he knows of another stock that is selling for \$99 which is worth to him \$110. In that case, he might be willing to sell his share of stock for \$99 because he could then use the money to make a net profit. But if all stocks, services and other properties were evaluated by all people the same, that would stifle the market.

Let us consider a bar of soap selling in a shop for one dollar. To the owner of the shop, the bar of soap is worth less than a dollar, or she would not sell it for a dollar. To the customer who buys the bar of soap, it is probably worth more than a dollar (to see that this is so, we need only ask whether the customer would buy the soap bar for \$1.01.) The price of sale, from the standpoints of the buyer and the seller, is in the interval between what the vendor values the soap and that which the buyer values the soap (assuming the former is less than the latter, for otherwise no voluntary transaction will occur). A transaction at any price within the interval will be *Pareto preferred* to no transaction, as one on both parties will be better off without any party being worse off. However, there is no "scientific" basis for selecting a price within the interval, because each price is *Pareto optimal* (a movement from any price to any other price would make one of the parties worse off). A price outside of the interval will only generate transactions through coercion, because one of the parties will be worse off than not transacting at all. For every value within that interval, we have achieved *Pareto efficiency*, i.e., the vendor is getting more than he values the bar of soap and the purchaser is buying it for less than he values it. The task of the merchant is to raise the price as high as possible without exceeding the valuation of a significant fraction of his customers. Any notion of *fair market price*, a one price for the same item, is achievable only in a state controlled situation. And such statist controlled prices generally produce absurdities.

In Soviet-controlled Poland, the state suppressed the price of bread to well below the cost of production. It made this work, in part, by forcing the farmers (Poland's farms were privately owned and operated, for the most part, even during the Russian occupation) to sell a certain fraction of their wheat below the cost of production. The price of pork, however, was not so artificially depressed. So some clever peasants became rich by buying stale bread and feeding it to their pigs. It was difficult for the authorities to overcome this strategy, since there was really no other good set of buyers for stale bread. The best that could be done was to raise the price of stale bread to nearly that of fresh bread. Ultimately, the only good solution (to escape embarrassment of the officials) was to raise the price of fresh bread as well. A market with set prices is very much like a cardiac patient who takes one medication for his heart plus three other medications to counteract bad side effects of the heart medication plus five additional medications to counteract the bad side effects of the three medications taken to counteract the effects of the primary heart medication.

12.4 Creating a "St. Petersburg Trust"

In the Enlightenment world of Daniel Bernoulli, neat and concise answers to virtually any problem were deemed possible if one only had the right insight. In our modern world, we understand that there is likely to be some arbitrariness in most simple solutions. All the notions of utility, at which we have looked, have some underlying assumption that if one only looks at the proper function of wealth, then decisions become rather clear, using a simple *stochastic* (probabilistic) model, relying on a few summary numbers, such as expected value. By constructing a utility function and then looking at its expected value in the light of a particular policy, Bernoulli thought he could capture the entire profile of risk and gain. Simply looking at a utility linear in money came up, in the case of the St. Petersburg game, with an absurd result, infinite gain. So, Bernoulli constructed other utility functions which had other than infinite expectation for the St. Petersburg game. He really did not, however, achieve his goal: reducing the entire prospectus of the gain to one scalar number. Moreover, market decisions these days are still based on dollars—not logarithms of dollars or square roots of dollars. This is surely evidence that utility theory has not lived up to the hope some have held for it.

The real essence of constructing a risk profile has to do with looking at the stochastic process that underlies the investment being considered. (For detailed information concerning stochastic processes, the reader is referred to the Appendix A at the end of this book.) We need to look at the probabilities of various results which might be obtained and decide, in the aggregate, whether the deal appears good. Of course, at the end of the day, we must make a decision whether to take the deal or turn it down. However, we believe that this is better done from the risk profile (time slices of the cumulative distribution function of the payoff) than from a one dimensional summary number (which, of course, expectation is).

In our postmodern world, in which few things are really simple, we should take a different view. It could be argued that Bernoulli was enchanted with the notion that a benevolent Providence had constructed the universe in structures selected to make them easy for human beings to understand. Or, he might have taken the nominalist (essentially, postmodern) view of William of Ockham that "truth" was a matter simply of fashion, so one might as well pick the simplest model that seemed, more or less, to work. After all, Bernoulli had no notion of fast computing, a primitive slide rule being the hottest computer available.

The view we shall take is that one should try to use any means available to get close to the truth (and we do not put quotes around the word), realizing we will generally use a model at variance with reality, but hopefully not too far from it. Bernoulli's consideration of a game with infinite expected payoff was unfortunate. In the first place, one should ask who the croupier would be for such a game. Beyond that, games with infinite

expectation do not correspond to any real world economic situation (again, we notice that no casino in the world can offer a payout of ∞ . Everyone places some maximum betting limit—perhaps high for “high rollers” but not ∞ .) Finally, as we show below, a game with an infinite expected value can be diddled in such a way that a clever player can come up with a very high gain with high probability.

Next, we will perform some computer simulations to try and understand better the St. Petersburg Paradox. First of all, we will follow the common policy of not using a utility function formally, but only implicitly. If we play the game for \$2 payoffs, we get a false sense of value. So let us make the payoffs to be in the \$2 million range. Suppose we take a reasonably well-to-do individual and ask how much he would pay to play the game using the profile in Figure 12.1. We note that the investor cannot make less than \$2 million. So that is an obvious floor. We recall that the expected payoff of the game is infinite, but the chance of making more than \$2 million is only 50%. The chance of making more than \$4 million is only 25%. The chance of making more than, say, \$5 million is also 25%. If the person’s total assets are, say, \$5 million, it is hard to believe he will be willing to pay the full \$5 million to play the game, even knowing that there is a \$2 million floor below which he cannot fall. Rather clearly, he would be willing to pay, say, \$2.1 million to play. For figures between, say \$2.4 million and \$5 million, we would see a variety of decisions made by the variety of possible players, depending upon their incomes, ages, psychology, and so on. Will the decision be made on the basis of expected winnings? Of course not. It will be made utilizing the entire information given in Figure 12.1.

Next, let us see whether the investor can structure the game somewhat differently. Suppose he notes the possibility of arranging things so that another investor and he form a trust and strike an agreement with the house that each will play the game for half the payoffs and the winnings will be pooled and divided by two. In this case, we note that both players may get heads first toss. In that case, the pooled winnings will be \$2,000,000. This occurs with probability $0.50 \times 0.50 = 0.25$. Suppose the first player gets heads first time, but the second gets tails first time immediately followed by heads. That would result in net winnings of \$3,000,000. Or it could be the first player gets tails first, then heads. Again, net winnings of \$3,000,000. The probability of one or the other of these is $2 \times 0.5 \times 0.25 = 0.25$. Next, they could both toss TH for a total winning of \$4,000,000. The probability of this is $0.25 \times 0.25 = 0.125$. So, then, the probability that the trust wins more than \$2,000,000 is $1 - 0.25 = 0.75$. The probability the trust wins more than \$4,000,000 is $1 - 0.25 - 0.25 - 0.125 = 0.375$. According to the original game in which only one player plays, these probabilities were .50 and 0.25, respectively. Clearly, the idea of playing according to each player playing for half the pay-offs of the original game rules is a good one. We can carry out computer simulations for games with varying numbers of players: one, ten, one hundred, one thousand. We show the probability profile in

Figure 12.5 where sums are in millions of dollars.

Here we see the advantage which can be achieved by pooling in the rather stylized St. Petersburg Paradox situation. With 1,000 investors participating, the probability of the trust winning a total in excess of nine million dollars is 90%. For the same game played by one investor, the probability of winning in excess of nine million dollars is less than 10%.

It would be a fine thing if there were some sort of way we might develop a portfolio or trust fund which mimicked the strategy above for dealing with Bernoulli's St. Petersburg scenario. Alas, the author is unable to point out such a strategy. Essentially, our "trust strategy" uses the fact that sample means converge, with increasing sample size, to the mean (expected value) of the population. In the St. Petersburg case, this value is infinite. Nature and the market do not tend to provide such opportunities.

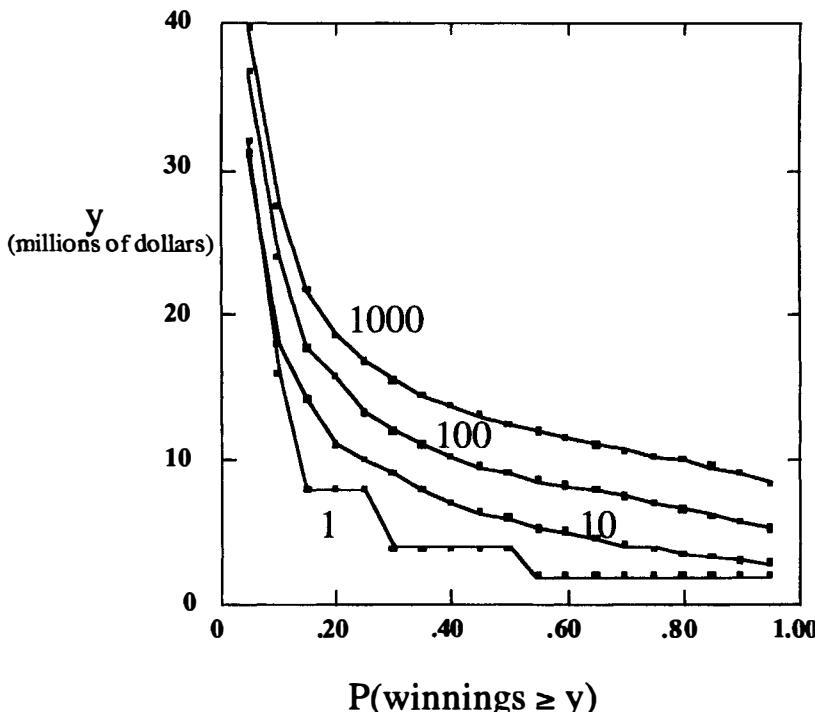


Figure 12.5. Probability profile for St. Petersburg trust.

There are other flaws with the St. Petersburg scenario of Bernoulli. For example, the coin throws are taken to occur with constant probability. But storms in the North Atlantic and Baltic are not of constant probability of occurrence. A heavy storm period will increase the risks to all shipping during that period. This happens also in markets. There will be periods of bull market growth across the market. And there will be periods of bear market declines. There may be safe harbors during the bear market times.

But very frequently, as in the 2000–2001 case, a bear market will adversely affect most securities.

Perhaps the bottom line in our analysis of the St. Petersburg Paradox is that by examining the full probability profile instead of simply looking at the expected value of utility, we can gain something of an understanding of what is going on. We replace arbitrary formalism and a simple answer with a more complicated, albeit more insightful analysis. Moreover, the real world may not be one of risk but rather varying degrees of hazy uncertainty. Finally, almost all people violate the axioms anyway. That is why they purchase insurance on one day and go to Las Vegas on the next!

12.5 Some Problems with Aggregate Choice Behavior

So far, we have been concerned with the utility–based choices made by individuals. We now look at the pooling of individual choices to obtain pooled choices of groups of individuals. A key part of our treatment will be to show Kenneth Arrow’s proof [1] that rational rules for preferences among individuals do not translate into the same rules for the group. We essentially follow the treatment given by Thompson [5] in *Empirical Model Building*.

To begin, we have a collection of n individuals $\{1, 2, \dots, n\} = G$. The task confronting the individuals is to rank their preferences amongst at least three decisions $D = \{a, b, c, \dots\}$. By $a \succ_1 b$, we mean that the first individual in the group prefers a to b . By $a \succ_G b$, we mean that the group as a whole prefers a to b , i.e., whatever the underlying mechanism used to obtain group consensus, the group picks a over b .

Suppose we have the following four axioms for pooling individual preferences into group decision making:

1. For a particular set of individual preferences, suppose the group prefers a to b ($a \succ_G b$). Then, suppose that some of the individuals change their preferences in such a way that preferences for a over b are unchanged or increased in a ’s favor, and that each individual’s preference between a and any alternative other than b are unchanged. Then, the group preference for a over b is maintained.
2. (Axiom of the Irrelevant Alternative.) Suppose that the group prefers a to b . Then, some of the individual preferences between alternatives other than a and b are changed, but the preferences between a and b are unchanged. Then, the group preference for a over b is maintained.
3. For any pair of alternatives a and b , there is some collection of individual preferences for which $a \succ_G b$.

4. (Axiom of Disallowing Dictators). No individual in the group has such influence that if he or she prefers a to b , and every other member of the group ranks b over a , then $a >_G b$.

Arrow's Impossibility Theorem. If Axioms 1–3 hold, then Axiom 4 cannot hold, assuming there are two or more decision makers and three or more possible decisions.⁶

Definition: Suppose we have $k \geq 3$ mutually exclusive decisions for which each of $n \geq 2$ voters have ordered preferences. Let $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$ represent the ordered preferences (*profile*) of each of the individual voters. Let J be a subset of the set of individual voters G and a and b be included among the set of decisions D . If, for all the individuals in J , $a P_i b$ (i.e., $a >_i b$), we then say that \mathcal{P} is *J-favored* for (a, b) . If for all the individuals not in J , $b P_i a$, then we say that J is *strictly J-favored*.

Definition: Suppose that the fact that all members in J favor a over b implies that $a F(\mathcal{P}) b$, where F is the group decision rule (aka “social utility function”) which integrates the preferences of the individual voters into a group decision for \mathcal{P} (i.e., we suppose that all the members in subset J deciding for a over b will cause the consensus of the entire group \mathcal{P} to prefer a over b .) Then we say that J is *decisive* for (a, b) .

Definition: A *minimal decisive set* J is a subset of G which is decisive for some (a, b) and which has the property that no subset of J is decisive for

⁶Before proving the theorem, we note that all the axioms appear quite reasonable at first glance, somewhat less so upon closer inspection. For example, even in North Korea, the Supreme Leader could not make policy opposed by everyone else on the Politburo. But then, it is not hard to consider examples where dictators are possible. For example, a person owning 51% of the stock of a company can make policy disagreed to by all the other shareholders. Then, again, looking at the Axiom of the Irrelevant Alternative, suppose there are two Republican candidates, one strongly pro life, the other strongly pro choice. 51% of the party strongly supports the pro life position, the other 49% strongly supports the pro choice position. If the vote were held between these two, the strongly pro life candidate would win the primary. Of course, this would be done in the light of the general election, where the Democrats would then be able to capture the votes of many pro choice Republicans. But then, a third feasible candidate appears, one who is mildly pro life. In all probability, this third candidate will get the nomination. Clearly, in bringing forward this example, we have used the fact that decisions are seldom stand alone. It is the knowledge of the general election which will cause some of the strongly pro life Republicans to consider voting for the mildly pro life candidate. But let us suppose the four axioms above are all satisfied. Then, Arrow shows that they cannot be all satisfied. It is simply an impossibility. Fundamentally, the reason for this fact has to do with the fact that although in one-dimensional space, it is easy to make orderings (for example, clearly $2 > 1$), it is not nonarbitrarily possible in higher dimensional spaces (for example, we cannot say that $(2,1) > (1,3)$ nor that $(1,3) > (2,1)$, unless we arbitrarily impose a one-dimensional structure onto the two-dimensional space; but if we use as the criterion the sum of squares of the components, then we can say that $(1,3) > (2,1)$, since $1^2 + 3^2 > 2^2 + 1^2$. Having said this, we shall go through *Arrow's Impossibility Theorem* in part to show how much more important are insight and conjecture than theorem proving ability in the acquisition of Nobel Prizes. The proof is rather easy. The conjecture (which is what a theorem is before it has been proved) is profound, and its consequences are important in the consideration of efficient market theory.

any other pair of decisions.

Lemma 12.1. Assume Axioms 1–4. Then J is decisive for (a, b) if and only if there is a set of preferences which are strictly J – favored for (a, b) and for which $aF(\mathcal{P})b$.

Proof. Suppose J is decisive for (a, b) . Then any strictly J – favored for (a, b) profile has $aF(\mathcal{P})b$.

Next, suppose there is a profile \mathcal{P} that is strictly J – favored for (a, b) and for which $aF(\mathcal{P})b$. Then, every voter in J prefers a to b . But since \mathcal{P} is strictly J – favored for (a, b) , we know that voters in $G - J$ all prefer b to a . Let \mathcal{P}' be some other J – favored for (a, b) profile. For all voters in J , a is preferred to b for this second profile. Thus, insofar as voters in J are concerned, they all prefer a to b for both \mathcal{P} and \mathcal{P}' . But for $G - J$, all the voters, following profile \mathcal{P} prefer b to a . Nevertheless, $aF(\mathcal{P}')b$. However, for profile \mathcal{P}' it is possible some voters in $G - J$ prefer a to b . By Axiom 1, then, it must be true that $aF(\mathcal{P}')b$. Hence, J is decisive for (a, b) .

Lemma 12.2. Assuming all the four axioms to be true, the entire group G is decisive for every (a, b) .

Proof. Suppose every voter in G prefers a to b , but that $bF(\mathcal{P})a$. By Axiom 3, there must be some profile \mathcal{P} such that $aF(\mathcal{P})b$. Now, if every voter in G prefers a to b , but the group decision is for b over a , then changing some of the voters' preferences to b over a can only (by Axiom 1) strengthen the group resolve to prefer b over a . That would contradict Axiom 3, for then there would be no profile \mathcal{P} such that $aF(\mathcal{P})b$. Hence, G is decisive for every (a, b) .

Next, let J be a minimal decisive set. We know there is a decisive set, since the preceding lemma proved that G is decisive for every (a, b) . So, we can remove individuals from G until one more removal would no longer give a decisive set. Pick one of the voters j from the minimal decisive set J . We shall prove that j must be a dictator, contrary to Axiom 4.

Suppose J is decisive for (a, b) . Pick another decision c which is neither a nor b . Consider the profile shown in Table 2.1.

Table 12.1.

P_i for i in $J - j$	P_i for i not in J	P_j
c	b	a
a	c	b
b	a	c
$D - \{a, b, c\}$	$D - \{a, b, c\}$	$D - \{a, b, c\}$

By construction, J is decisive for (a, b) . Thus $aF(\mathcal{P})b$. We note that \mathcal{P} is strictly $J - j$ favored for (c, b) . Thus, if $cF(\mathcal{P})b$, then by Lemma 2.1, $J - j$

would be decisive for (c, b) contrary to the assumption that J is a minimal decisive set. Thus, c is not favored over b by the group, and a is favored over b by the group.

Consequently, we have two possible scenarios for the preference of the group: either a is preferred to b is preferred to c or a is preferred to b and c and the group ties b and c . Then, in both cases, we have $aF(\mathcal{P})b$. But j is the only voter who prefers a to c . Thus, by Lemma 12.1, j is decisive for (a, c) . Thus j cannot be a proper subset of j , i.e., $j = J$. So far we have shown that j is decisive for (a, c) for any $c \neq a$.

Next, we shall establish that j is decisive for (d, c) for any d, c not equal to a . Consider the profile in Table 12.2.

Table 12.2.

P_j	P_i for $i \neq j$
d	c
a	d
c	a
$D - \{a, c, d\}$	$D - \{a, c, d\}$

Note that the entire set G is decisive for any pair of decisions as we have proved in Lemma 12.2. Hence, $dF(\mathcal{P})a$. Thus the group as a whole ranks d over a and ranks a over c . Thus, $dF(\mathcal{P})c$. Therefore, by Lemma 2.1, j is decisive for (d, c) .

Finally, we shall demonstrate that j is decisive for (d, a) whenever $d \neq a$. Consider the profile in Table 12.3.

Table 12.3.

P_j	P_i for $i \neq j$
d	c
c	a
a	d
$D - \{a, c, d\}$	$D - \{a, c, d\}$

As j is decisive for (d, c) , we have $dF(\mathcal{P})a$. But j is the only individual preferring d to a . Hence, j is a dictator, contrary to Axiom 4, and the theorem is proved!

12.6 Jeffersonian Realities

Arrow's Impossibility Theorem was perceived intuitively by earlier social scientists. Vilfredo Pareto, for example, whose father had labored his entire life to bring forth an enlightened Jeffersonian democratic system to Italy, and who himself believed in the feasibility of such a system until well into middle age, finally opined that all social systems would naturally be controlled not by an orderly pooling of individual preferences, but rather by a circle of elites. In his youth, Pareto assumed that the twin pillars of

Jeffersonian government and Adam Smith policies toward free trade would bring economies and governments into a state of efficiency and optimality. Let us look at Arrow's result in the context of Jeffersonian performance as opposed to Jeffersonian ideals. In 1792, the Congress was confronted with the task of deciding how many members of the House of Representatives would be allocated to each state. Alexander Hamilton, the alleged opponent of states' rights, proposed the following rule:

Hamilton's Rule. Pick the size of the House = n . Divide the voting population N_j of the j th state by the total population N to give a ratio r_j . Multiply this ratio by n to give the quota q_j of seats for the j th state. If this quota is less than one, give the state one seat. Give each state the number of representatives equal to the integer part of its quota. Then rank the remainders of the quotas in descending order. Proceeding down the list, give one additional seat to each state until the size of the House n has been equaled.

The method of Hamilton has firmly embodied in it the notion of the state as the basic entity of indirect democracy. Once the number of Representatives had been arrived at by a comparison of the populations of the several states, the congressional districts could be apportioned by the state legislatures within the states. But the indivisible unit of comparison was that of state population. If one conducts a poll of educated Americans and asks how seats in the House of Representatives are apportioned amongst the several states, much the most popular rule given is that of Hamilton. It is a very intuitive rule. Furthermore, if we let a_j be the ultimate allocation of seats to each state, then Hamilton's Rule minimizes

$$\sum_{j=1}^k |a_j - q_j|, \quad (12.15)$$

(i.e., Hamilton's Rule minimizes the sum of the discrepancies between the allocations obtainable without consideration of states and those with the notion of noncrossover of state boundaries to obtain districts.)⁷

⁷Census figures are from Balinski and Young [2].

Table 12.4. Method of Hamilton.

State	Population	Quota	Hamiltonian Allocation	Voters per Seat
Connecticut	236,841	7.860	8	29,605
Delaware	55,540	1.843	2	27,770
Georgia	70,835	2.351	2	35,417
Kentucky	68,705	2.280	2	34,353
Maryland	278,514	9.243	9	30,946
Massachusetts	475,327	15.774	16	29,708
New Hampshire	141,822	4.707	5	28,364
New Jersey	179,570	5.959	6	29,928
New York	331,589	11.004	11	30,144
North Carolina	353,523	11.732	12	29,460
Pennsylvania	432,879	14.366	14	30,919
Rhode Island	68,446	2.271	2	34,223
South Carolina	206,236	6.844	7	29,462
Vermont	85,533	2.839	3	28,511
Virginia	630,560	20.926	21	30,027
Total	3,615,920	120	120	

It is interesting to note that the first Congressional Senate and House passed a bill embodying the method of Hamilton and the suggested size of 120 seats. That the bill was vetoed by President George Washington and a subsequent method, that of Thomas Jefferson, was ultimately passed and signed into law is an interesting exercise in Realpolitik. The most advantaged state by the Hamiltonian rule was Delaware, which received a seat for every 27,770 of its citizens. The most disadvantaged was Georgia which received a seat for every 35,417 of its citizens. Jefferson's Virginia was treated about the same as Hamilton's New York with a representative for every 30,027 and 30,144 citizens, respectively. The discrepancy between the most favored state and the least favored was around 28%, a large number of which the supporters of the bill were well aware. The allocation proposed by Hamilton in 1792 did not particularly favor small states. In general, however, if we assume that a state's likelihood of being rounded up or down is independent of its size, the method of Hamilton will favor somewhat the smaller states if our consideration is the number of voters per seat. But Hamilton, who was from one of the larger states and who is generally regarded as favoring a strong centralized government which de-emphasized the power of the states, is here to be seen as the clear principled champion of states' rights and was apparently willing to give some advantage to the smaller states as being consistent with, and an extension of, the notion that each state was to have at least one Representative. Now let us consider the position of Thomas Jefferson, the legendary defender of states' rights. Jefferson was, of course, from the largest of the states, Virginia. He was loathe to see a system instituted until it had been properly

manipulated to enhance, to the maximum degree possible, the influence of Virginia. Unfortunately for Jefferson, one of the best scientific and mathematical minds in the United States who undoubtedly recognized at least an imprecisely stated version of (12.15), the result in (12.15) guaranteed that there was no other way than Hamilton's to come up with a reasonable allocation rule fully consistent with the notion of states' rights. Given a choice between states' rights and an enhancement of the power of Virginia, Jefferson came up with a rule which would help Virginia, even at some cost to his own principles. Jefferson arrived at his method of allocation by departing from the states as the indivisible political units. We consider the method of Jefferson as follows:

Jefferson's Rule. Pick the size of the House = n . Find a divisor d so that the integer parts of the quotients of the states when divided by d sum to n . Then assign to each state the integer part of N_j/d .

We note that the notion of a divisor d is an entity which points toward House allocation which could occur if state boundaries did not stand in the way of a national assembly without the hindrance of state boundaries. Let us note the effect of Jefferson's method using the same census figures as in Table 12.4.

We note that the discrepancy in the number of voters per representative varies more with Jefferson's method than with Hamilton's—94% versus 28%. In the first exercise of the Presidential veto, Washington, persuaded by Jefferson, killed the bill embodying the method of Hamilton, paving the way for the use of Jefferson's method using a divisor of 33,000 and a total House size of 105. Let us examine the differences between the method of Hamilton and that of Jefferson.

The only practical difference between the two allocation systems is to take away one of Delaware's two seats and give it to Virginia. The difference between the maximum and the minimum number of voters per seat is not diminished using the Jeffersonian method which turns out to give a relative inequity of 88%; for the Hamiltonian method the difference is a more modest 57%. The method of Jefferson favors the larger states pure and simple. Jefferson essentially presented George Washington and Congress with a black box and the message that to use Hamilton's Rule would be unsophisticated, whereas Jefferson's Rule was somehow very politically correct. It worked. Congress approved the method of Jefferson, and this method was in use until after the census of 1850 at which time the method of Hamilton was installed and kept in use until it was modified by a Democratic Congress in 1941 in favor of yet another scheme.

Table 12.5. Method of Jefferson (divisor of 27,500).

State	Population	Quotient	Jeffersonian Allocation	Voters Seat
CN	236,841	8.310	8	29,605
DE	55,540	1.949	1	55,540
GA	70,835	2.485	2	35,417
KY	68,705	2.411	2	34,353
MD	278,514	9.772	9	30,946
MA	475,327	16.678	16	29,708
NH	141,822	4.976	4	35,456
NJ	179,570	6.301	6	29,928
NY	331,589	11.635	11	30,144
NC	353,523	12.404	12	29,460
PA	432,879	15.189	15	28,859
RI	68,446	2.402	2	34,223
SC	206,236	7.236	7	29,462
VE	85,533	3.001	3	28,511
VA	630,560	22.125	22	28,662
Total	3,615,920	126.88	120	

We have in the 1792 controversy a clear example of the influence on consensus of one individual who passionately and cleverly advances a policy about which his colleagues have little concern and less understanding. Jefferson was the best mathematician involved in the Congressional discussions, and he sold his colleagues on a plan in the fairness of which one must doubt he truly believed.

Table 12.6. Allocations of Hamilton and Jefferson.

State	Population	H	J	Voters/Seat	Voters/Seat
				Hamilton	Jefferson
CN	236,841	7	7	33,834	33,834
DE	55,540	2	1	27,220	55,440
GA	70,835	2	2	35,417	35,417
KY	68,705	2	2	34,353	34,353
MD	278,514	8	8	34,814	34,814
MA	475,327	14	14	33,952	33,952
NH	141,822	4	4	35,455	35,455
NJ	179,570	5	5	35,914	35,914
NY	331,589	10	10	33,159	33,159
NC	353,523	10	10	35,352	35,352
PA	432,879	13	13	33,298	33,298
RI	68,446	2	2	34,223	34,223
SC	206,236	7	7	29,462	29,462
VE	85,533	2	2	42,766	42,776
VA	630,560	18	19	35,031	33,187
Total	3,615,920	106	106		

Naturally, as the large state bias of the method of Jefferson began to be understood, it was inevitable that someone would suggest a plan that, somewhat symmetrically to Jefferson's method, would favor the small states. We have such a scheme proposed by John Quincy Adams.

John Quincy Adams's Rule. Pick the size of the House = n . Find a divisor d so that the integer parts of the quotients (when divided by d) plus 1 for each of the states sum to n . Then assign to each state the integer part of $N_j/d + 1$.

The plan of Adams gives the same kind of advantage to the small states that that of Jefferson gives to the large states. Needless to say, it has never been used in this country or any other (although amazingly, Jefferson's has). It is interesting to note that Adams, instead of saying, "I see what's going on. Let's go to Hamilton's Rule," tried to do for the small states the same thing Jefferson had done for the large states.

It is interesting to note that Daniel Webster attempted to come up with a plan which was intermediate to that of Jefferson's and that of Adams. He noted that whereas Jefferson rounded the quotient down to the next smallest integer, Adams rounded up to the next largest integer. Webster, who was a man of incredible intuition, suggested that fractions above .5 be rounded upward, those below 0.5 be rounded downward. From the 1830's until 1850 there was very active discussion about the unfairness of the method of Jefferson and a search for alternatives. It was finally decided to pick Hamilton's method, but Webster's was almost selected and it was a contender as recently as 1941. As it turns out, there is very little practical difference between the method of Hamilton and that of Webster. Both methods would have given identical allocations from the beginning of the Republic until 1900. Since that time, the differences between the two methods usually involve one seat per census. The method of Hamilton was replaced in 1941 by one advocated by Edward Huntington, Professor of Mathematics at Harvard. Huntington, instead of having the division point of fractions to be rounded up and rounded down to one half, advocated that if the size of the quotient of a state were denoted by N_j/d then the dividing point below which rounding down would be indicated would be the geometric mean $\sqrt{[N_j/d]([N_j/d] + 1)}$, where $[.]$ denotes "integer part of." One might say that such a method violates the notion that such methods should be kept simple. Furthermore, the rounding boundaries do increase slightly as the size of the state increases, giving an apparent advantage to the smaller states. At the last minute, the more popular method of Webster was rejected in favor of that of Huntington, since its application using the 1940 census would give a seat to Democratic Arkansas rather than to Republican Michigan. The Huntington method is in use to this day, though not one American in a thousand is aware of the fact. And indeed, it is not a very important issue whether we use the method of Hamilton or that of Webster or that of Huntington or even that of Jefferson or that of

Adams. Not one significant piece of legislation would have changed during the course of the Republic if any one of them were chosen. The subject of apportionment possibly receives more attention than practicality warrants.

But looking over the history of the apportionment rule should give pause to those among us who believe in the pristine purity of the Founding Fathers and the idea that there was a time when Platonic philosophers ruled the land with no thought except virtue and fair play. Jefferson, the noblest and wisest and purest of them all, was working for the advantage for his state under the guise of fairness and sophistication. And George Washington, the Father of his Country, was tricked into using the first veto in the history of the Republic to stop a good rule and put one of lesser quality in its place. All this, in the name of being a good Enlightenment chief of state.

Arrow [1] proved mathematically what Machiavelli had observed and Pareto had described as being part of a general taxonomy of political behavior: in economics, politics and society generally: important public policy decisions are not made as the orderly aggregation of collective wisdom. Group decisions are made using mechanisms we do not clearly understand (and, as Pareto urged, we really should try without passion to learn these mechanisms). It appears that insiders have a great deal to do with these decisions. Sometimes, as when Robert Rubin, Secretary of the Treasury, rushed in to secure Mexican loans made by American financial institutions, including that of which he had been boss, Goldman-Sachs, one may raise an eyebrow. Of course, he could respond that he was just following an example set by arguably the smartest and most idealistic President in the history of the Republic. Then again, we have the Federal Reserve Board in recent years presenting the stock market with frequent "Gotcha!" type surprises. What sort of efficiency is possible in a market where interest rates are changed at the whim of Alan Greenspan and his fellows? Then there is the matter of anti-trust law. Should a CEO be careful lest his company be so successful that it is dismembered by the Attorney General?

Some years ago, the author was consultant to a firm that had a well-designed plan to raise chickens in Yucatan according to the notions of modern poultry husbandry. In order to purchase buildings, poultry, land, and equipment, the firm had to convert millions of dollars of cash from dollars into pesos. It did so on the Friday before the deals were to be paid for on Monday. During the week-end, the President of Mexico devalued the peso hugely, wiping out the cash reserves of the firm. Naturally, on Monday, all the Mexican parties from whom land, structures, equipment, poultry, etc., were to be purchased, changed their peso prices hugely upward to reflect the devaluation. (It turns out that prior to the devaluation, the President of Mexico had leveraged his own cash assets to purchase a number of villas in Mexico.) Such surprises are hardly helpful to facilitating market efficiency. One may well ask how an investor can cope with such inefficiency producing spikes in the market. Pareto's answer is that one had better learn how to do precisely that, for such is the real world. There

is indeed a tendency towards efficiency in most markets. But this is only part of the market mechanism. In our study, we model an efficient market tendency and then look at what happens when varieties of contaminating mechanisms are superimposed upon it.

12.7 Conclusions

In this chapter on utility, we have started with Enlightenment assurance as to the way rational people should make choices. The buttressing of Bernoulli by von Neumann and Morgenstern seemed very promising indeed. But then we ran into Arrow and looked back to Pareto and (shudder) Thomas Jefferson. And we became less confident about the orderliness of the way aggregate decisions are made. In *Mind and Society*, Pareto tells us that in reality many decisions are made with results which are the equal in consequence of any Adam Smith might have made. The market may not be efficient; politicians may not be clones of Lucius Quintus Cincinnatus; but in some societies (and Pareto always had great hopes for the United States), the market moves toward efficiency, and political decisions are frequently Cincinnatus-like in their consequences. The West and its economies are not at all chaotic. There is a tendency toward efficiency. But it is ridiculous if we assume efficient markets as an iron law, and then fudge when our illusions are challenged by reality. In this book, we shall assume a more or less efficient driver for a given market situation with departures therefrom taken cognizance of empirically by perturbations to the model.

Problems

12.1. Sempronius owns goods at home worth a total of 4,000 ducats and in addition possesses 8000 ducats of commodities in foreign countries from where they can be transported only by sea. However, our daily experience teaches us that of ten ships, one perishes.

- (a) What is Sempronius' expectation of the commodities?
- (b) By how much would his expectation improve if he trusted them equally to two ships?
- (c) What is the limit of his expectation as he trusted them to increasing numbers of ships?

12.2. Let us use Bernoulli's logarithmic utility function

$$U(X) = \log\left(\frac{X + \Theta}{\Theta}\right)$$

where Θ is one's initial wealth. Would it be rational to play a game where there is a finite probability of losing all one's wealth? Why? How might this result be rationalized? Why, in modern first world society, might it be said that is impossible to lose all one's wealth?

12.3. Suppose a woman were offered either a certain \$230 or a 50–50 chance of \$400 or \$100. Which option should be taken if:

- (a) She possesses a square root utility function,
- (b) A Bernoulli utility function with initial wealth of \$1,000,
- (c) The same as 2, with initial wealth of \$100?

12.4. Consider the section on von Neumann–Morgenstern utility.

(a) Determine the certain sum for which the subject in the example would relinquish the following opportunity:

W	Probability
40,000	.4
160,000	.4
250,000	.2

(b) Explain how von Neumann–Morgenstern Axiom 5 treats the utility derived from gambling itself (i.e., deals with the existence of casinos).

(c) Compare the mean and standard deviation of outcomes in the example and in (a).

(i) What are the mean and standard deviation in each case of the certain dollar value for which the subject would be indifferent?

(ii) We now have two sets of (σ, μ) for each of the two opportunities. Which of the two games (example or (a)) would the subject prefer to play?

(iii) For what probability α would he be indifferent between $\alpha \times$ (game he prefers) and the game he does not prefer?

12.5. For this problem, assume $U = \sqrt{W}$ and $U(0) = 0$ and $U(1,000,000) = 1,000$ utiles.

(a) For what $0 \leq \alpha \leq 1$ would the subject be indifferent between $\alpha \times \$1,000,000$ and a certain $\$250,000$?

(b) For what α would the subject be indifferent between $\alpha \times \$500,000$ and $(1 - \alpha) \times \$200,000$?

12.6. Consider the *Last Shall be First Rule* for establishing group preferences:

- Step 1. Select as the group choice, the candidate who was ranked number one by a majority of the decision makers, if such a candidate exists.
- Step 2. If no candidate received a majority of number one rankings, then take all the preference rankings of voters who voted for the candidate with the smallest number of first choice preferences and treat them as if their first choice candidate is simply deleted from their preference lists and all other candidates on their lists moved up one rung on the preference ladder.
- Step 3. Go to Step 1.

Which of Arrow's Axioms does this rule fail to satisfy? Discuss its advantages and disadvantages in three settings:

- (a) A mayoral election in an American city of over 10,000 voters
- (b) An election for a position on a university committee
- (c) A decision making process for picking one of four potential new products by a board of directors (with 20 members) of a company.

12.7. Consider the *Borda Count Rule*

- Step 1. Rank the preferences amongst k choices for voter i , one of n voters.
- Step 2. For each preference a , count the number of choices below a in the preferences of voter i . This gives us $B_i(a)$.
- Step 3. For each choice sum the Borda counts, e.g., $\mathbf{B}(a) = \sum_i^n B_i(a)$.
- Step 4. The group chooses the preference with the highest Borda count sum

A group decision rule defined on the set of all profiles (preferences) on the set of decisions is said to be *Pareto optimal* if for every a and b in the

set of decisions whenever a is ranked over b in every ranking of a profile, then a is ranked over b in the corresponding group ranking.

Prove or disprove the following:

The Borda count rule is Pareto optimal.

12.8. Consider the following *Plurality Rule*: rank a over b for the group if and only if a receives more first place votes than b . Is this rule Pareto optimal?

References

- [1] Arrow, K. J. (1950). "A difficulty in the concept of social welfare, *Journal of Political Economy*, v. 58, pp. 328–346.
- [2] Balinski, M. L. and Young, H. P. (1982). *Fair Representation: Meeting the Ideal of One Man, One Vote*. New Haven: Yale University Press.
- [3] Bernoulli, D., "Exposition of a new theory on the measurement of risk" (Louise Sommer, trans.), *Econometrica*, January 1954, pp.23-36 (first published 1738).
- [4] Pareto, V. (1935). *Mind and Society*. New York: Harcourt, Brace, and Company.
- [5] Thompson, J. R. (1989). *Empirical Model Building*. New York: John Wiley & Sons, pp. 158–177.
- [6] von Neumann, J., and Morgenstern, O. (1944) *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- [7] Williams, E. E. and Findlay, M. C. (1974). *Investment Analysis*. Englewood Cliffs: Prentice Hall.

Chapter 13

A Primer in Sampling

13.1 Introduction

Typically, when we talk about sampling, we are randomly selecting a relatively few items from a much larger population in order to make inferences about that much larger population. If we fail to pick a subsample of the population which is typical of the population under consideration, nothing can save us from false conclusions. Three of the greatest statisticians of all time, Cochran, Mosteller, and Tukey [2] in examining Kinsey's [5] surveys on the sexual behavior of the American male were hoodwinked by failing to uncover the fact that a large fraction of Kinsey's sample consisted of male prostitutes and prison inmates. Had they probed more deeply into Kinsey's sampling methods, they could have spared the world years of false science. In effect, while being somewhat critical, they gave Kinsey a pass. It serves as an example to all persons examining sampling studies: be very careful to vet the sampler and his/her sampling strategy.

The increasing uses of sampling are quite significant. For example, two distinguished law professors, Walker and Monahan [7], have noted the advantages in time and cost if sampling techniques were used for assessing damages in class action lawsuits. This would be bad news for trial lawyers, but good news for society as a whole.

The notion of an opinion poll about various candidates comes immediately to mind. Most of the time, these opinion polls are amazingly accurate in assessing the current opinion of the electorate even though the sample is a tiny fraction, less than one in ten thousand, say, of the electorate. The polls may tell a presidential candidate that he or she is so far behind in a particular state, say less than 35% support, that it is simply pointless to spend significant marketing resources for campaigning in that state. Or it may tell the candidate that his or her support is so great, say over 65% that

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

it really is not necessary to expend significant resources in that state. Typically, based on opinion polling, the candidates will expend most of their resources in battleground states, states in which no candidate can be at all certain about the results. Consequently, in the presidential election of 2008, neither McCain nor Obama spent significant resources campaigning in California or Texas. In California, the polls showed Obama to be the easy winner. In Texas, the polls showed McCain to be comfortably ahead. Both candidates spent considerable resources in Indiana, Ohio, North Carolina, Virginia and Florida, for in these states neither candidate had polls showing over 55% strength.

Such neglect by the candidates of vast segments of the population is one argument used against the Electoral College, where one candidate receives all the electoral votes of the state, and the other receives none. On the other hand, a system whereby the President is selected on the basis of the aggregate popular vote without regard to state boundaries, many states would become "flyover" states. Little time would be spent campaigning in New Hampshire or Montana or Wyoming, for there are not so many voters in these states.

Let us next consider two states of very different population sizes, Wyoming and Texas. Wyoming has a population of roughly 500,000 whereas Texas has around 20,000,000. In other words, Texas has a population around 40 times that of Wyoming.

Consider a state of size N where there are M persons who prefer Obama. Let X be the number of voters in a sample of size n who prefer Obama. We wish to estimate the proportion of voters in the state who prefer Obama. In Appendix A, we have shown in (A.35) that

$$E(X) = \mu = np \quad (13.1)$$

where $p = M/N$ and $q = 1 - p$.

In (A.38), we have shown that

$$E[(X - \mu)^2] = \sigma^2 = npq(N - n)/(N - 1). \quad (13.2)$$

Generally, we will have a much larger population than the sample size. For this reason, we may frequently write the variance as

$$Var(X) = npq \quad (13.3)$$

We note that this gives us a conservative estimate for the variance of X , (i.e., a slightly inflated one). Again, if we do not have a rough idea as to the value of p , we might use $p = 0.5$, which will give an inflated estimate for the variance.

Let us suppose we take a random sample of size 400 from the population of Wyoming and find that 180 favor Obama. Based on our sample, the natural estimator for p is $\hat{p} = 180/400$. What is a natural two-tailed 95% confidence interval for p ? The obvious interval will be

$$\hat{p} \pm 1.96 \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}} = 0.45 \pm 0.04875. \quad (13.4)$$

Now, if we should carry out the same survey in Texas, and in a sample of size 400 received again 180 votes for Obama, we would obtain precisely the same 95% confidence interval. Many persons confronted with this fact for the first time are surprised that the size of the state does not enter into the formula.

On the other hand, suppose the question was to obtain a 95% confidence interval on the numbers of persons favoring Obama in the two states. Let us call these two numbers M_W and M_T .

Clearly then, we need to multiply by proportions confidence interval by the populations of the two separate states. This would give us

$$\begin{aligned} M_W &= 500,000\hat{p} \pm 500,000 \times 1.96 \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}} \\ &= 225,000 \pm .04875 \times 500,000 \\ &= 225,000 \pm 24,375. \end{aligned} \quad (13.5)$$

For Texas we have

$$\begin{aligned} M_T &= 20,000,000\hat{p} \pm 20,000,000 \times 1.96 \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}} \\ &= 9,000,000 \pm .04875 \times 20,000,000 \\ &= 9,000,000 \pm 975,000. \end{aligned} \quad (13.6)$$

Next, let us consider the situation where we have 800 samplings to use in Wyoming and Texas. Let us suppose that our goal is to minimize the variance of the sum of the estimated Obama voters. Our function to be minimized (assuming the polls are independent) for the two states is

$$V = \text{Var}(M_W) + \text{Var}(M_T) = 500,000 \frac{p_W(1 - p_W)}{n_W} + 20,000,000 \frac{p_T(1 - p_T)}{(800 - n_W)}. \quad (13.7)$$

Of course, if we actually knew p_W and p_T , we would not have to take a survey in the first place. Let us suppose, however, that we have, perhaps on the basis of earlier surveys, have rough a priori estimates of p_W and p_T —say, \hat{p}_W and \hat{p}_T .

Now, both Wyoming and Texas were not very strong for Obama. Let us assume that both states had approximately the same p for the proportion of Obama voters.

Then,

$$V = \left[\frac{500,000}{n_W} + \frac{20,000,000}{800 - n_W} \right] \frac{p}{1 - p}. \quad (13.8)$$

Taking the partial derivative with respect to n_W and setting the resulting equation equal to zero, we have

$$\frac{\partial V}{\partial n_W} = \frac{-500,000}{n_W^2} + \frac{20,000,000}{(800 - n_W)^2} = 0 \quad (13.9)$$

This gives

$$39n_W^2 + 1,600n_W - 640,000 = 0 \quad (13.10)$$

with the solution $n_W = 109$, and $n_T = 691$.

13.1.1 Tracking Polls

Frequently, when looking at public opinion polling, there may be some advantage to taking an exponentially weighted moving average approach. Below, we use weight $1 - \alpha$ for today's sample proportion favoring a candidate, and a weight α for the past sampled proportion. So, starting with day 1 for the first day of the tracking poll, we have

$$P_1 = \hat{p}_1. \quad (13.11)$$

Then, for the second day

$$P_2 = \alpha \hat{p}_2 + (1 - \alpha)P_1. \quad (13.12)$$

Continuing in this fashion, we have by day N

$$P_N = \alpha \hat{p}_{N-1} + (1 - \alpha)P_{N-1}. \quad (13.13)$$

The tracking poll is really an exploratory device. During a campaign the day to day and week to week values of p are likely to be changing somewhat, sometimes a great deal. By letting past week estimates influence current week estimates of the proportion favoring a candidate the sample size, though the process has changed, the sample size has effectively increased. If one wishes to create a sort of rough confidence interval on day N

$$P = p_N \pm 1.96 \sqrt{Var P_N}. \quad (13.14)$$

We demonstrate a few steps of the temporally indexed estimation process when $\alpha = 0.4$.

Table 13.1. Weekly Estimates for p .

j	n_j	\hat{p}_j	P_j	$Var(\hat{p}_j)$	$Var(P_j)$
1	100	0.400	0.40	0.0024	0.0024
2	105	0.42	0.408	0.0023	0.00123
3	110	0.43	0.420	0.00223	0.00080
4	85	0.428	0.420	0.00290	0.000752
5	100	0.48	0.4488	0.00175	0.000671

It would appear that the estimates of the proportion for Obama have not changed very much during the five week tracking period. This would be expected if this were a rather solid Republican state in which little had been spent in advertising either candidate. However, if we chose to look at a 95% confidence interval about P_6 , we would find the upper bound to be $0.4488 + 0.0507 = 0.4995$. Essentially, it would look like Wyoming is moving Obama's way. Perhaps he should not count Wyoming out.

13.2 Stratification

Suppose we wish to estimate a total of random variables in m containers. For example, we might want to estimate the total of voters favoring a candidate for President in m states. Or, we could be carrying out a survey to determine an estimate for the total value of items in m warehouses. This is a common inventory survey, which might be required for the payment of *ad valorem* taxes.

Both problems can be dealt with in the same urn model fashion. Let us start then with the situation where the N_1 variables in the first urn have mean μ_1 and variance $(\sigma_1)^2$. Similarly, the N_2 variables in the second urn have mean μ_2 and variance σ_2^2 . Continuing on in this fashion, we profess to the m th urn with mean μ and variance σ^2 .

We wish to estimate

$$T = N_1\mu_1 + N_2\mu_2 + \dots + N_m\mu_m. \quad (13.15)$$

One natural procedure is to take a sample of size n_1 in the first urn, n_2 in the second, and so on, until we have n_m in the m th urn. Using these samples, our natural estimator is

$$\hat{T} = N_1\bar{X}_1 + N_2\bar{X}_2 + \dots + N_m\bar{X}_m. \quad (13.16)$$

It is a fair assumption (usually) that the variables in each urn are independent of the variables in the other urns and each other. Using squared deviation from the true value of T as the loss function to be minimized, we try to minimize:

$$\begin{aligned} Var(\hat{T}) &= N_1^2 Var(\bar{X}_1) + N_2^2 Var(\bar{X}_2) + \dots + N_m^2 Var(\bar{X}_m) \quad (13.17) \\ &= N_1^2 \frac{\sigma_1^2}{n_1} + N_2^2 \frac{\sigma_2^2}{n_2} + \dots + N_m^2 \frac{\sigma_m^2}{n_m}. \end{aligned}$$

(Note how we have used independence to eliminate cross product terms.)

Now, the question becomes: If we have a total sample size across urns of

$$n_1 + n_2 + \dots + n_m = n, \quad (13.18)$$

how shall we optimally allocate the samples across the m urns? Using the method of Lagrange for optimizing T subject to the constraint in (13.18), we have

$$-N_1^2 \frac{\sigma_1^2}{n_1^2} + \lambda \frac{\partial n}{\partial n_1} = -N_1^2 \frac{\sigma_1^2}{n_1^2} + \lambda = 0 \quad (13.19)$$

$$\begin{aligned} -N_2^2 \frac{\sigma_2^2}{n_1^2} + \lambda &= 0 \\ \dots \\ -N_m^2 \frac{\sigma_m^2}{n_1^2} + \lambda &= 0 \end{aligned}$$

Isolating λ and using the fact that things equal to λ are equal to each other, we have

$$N_1^2 \frac{\sigma_1^2}{n_1^2} = N_2^2 \frac{\sigma_2^2}{n_1^2} = \dots = N_m^2 \frac{\sigma_m^2}{n_1^2}. \quad (13.20)$$

This gives

$$\frac{n_2^2}{n_1^2} = \frac{\sigma_2^2}{\sigma_1^2} \frac{N_2^2}{N_1^2}. \quad (13.21)$$

So for each j

$$n_j = n_1 \frac{\sigma_j}{\sigma_1} \frac{N_j}{N_1} \quad (13.22)$$

Thus

$$n_1 = \frac{n}{1 + \frac{\sigma_2}{\sigma_1} \frac{N_2}{N_1} + \frac{\sigma_3}{\sigma_1} \frac{N_3}{N_1} + \dots + \frac{\sigma_m}{\sigma_1} \frac{N_m}{N_1}} \quad (13.23)$$

And, for any other n_j ,

$$n_j = n_1 \frac{\sigma_j}{\sigma_1} \frac{N_j}{N_1}. \quad (13.24)$$

Let us return to the election poll. Suppose in state j , there are N_j voters. Suppose for each voter, the probability of each voter voting for Obama is p_j . Then

$$\sigma_j = p_j(1 - p_j). \quad (13.25)$$

And, remembering,

$$n_1 = \frac{n}{1 + \frac{p_2(1-p_2)}{p_1(1-p_1)} \frac{N_2}{N_1} + \frac{p_3(1-p_3)}{p_1(1-p_1)} \frac{N_3}{N_1} + \dots + \frac{p_m(1-p_m)}{p_1(1-p_1)} \frac{N_m}{N_1}} \quad (13.26)$$

If we have some prior information as to the various probabilities—say \hat{p}_j —we simply substitute these values in (13.25) and (13.26).

Thus,

$$n_j = n_1 \frac{\sigma_j}{\sigma_1} \frac{N_j}{N_1}. \quad (13.27)$$

We see that as the standard deviation in urn j increases, the proportion of the sample used in urn j increases. And, as the number of units in the urn increases, so does the sample size in the j th urn increase.

What if our prior guesses for the proportions favoring Obama are off the mark? Will our estimate for the total number of voters favoring a candidate be flawed? The answer is that it will not. Our estimates for each state will be unbiased. As long as we are using correct values for the sizes of the voter populations on a state-by-state basis and our sampling is done randomly, our estimate for the total number of voters favoring a candidate will be unbiased. By using incorrect values for the \hat{p}_j in the formula for the sample sizes, state by state, we will have sacrificed some efficiency in the estimation process, but our answer will still be unbiased.

Next, let us apply stratification to an inventory problem. We have goods in a number of oil tool warehouses. We need to come up with a reasonable figure for the total worth of the goods on hand. In some of the warehouses, we may have things as simple as nuts and bolts. In others, we will have drilling bits and extenders worth many thousands of dollars. How shall we decide where to invest our sampling of n items. A general assumption, frequently a fair approximation, is that the coefficient of variation is constant across each warehouse. That is to say, for all the warehouses

$$\frac{\sigma_1}{\mu_1} \approx \frac{\sigma_2}{\mu_2} \approx \dots \approx \frac{\sigma_m}{\mu_m}. \quad (13.28)$$

And

$$n_j = n_1 \frac{\mu_j}{\mu_1} \frac{N_j}{N_1}. \quad (13.29)$$

where,

$$n_1 \approx \frac{n}{1 + \frac{\mu_2 N_2}{\sigma_1 N_1} + \frac{\mu_3 N_3}{\mu_1 N_1} + \dots + \dots + \frac{\mu_m N_m}{\mu_1 N_1}}. \quad (13.30)$$

13.2.1 A Warehouse Inventory

Let us suppose we have a warehouse with three different categories of items, as in Table 13.2.

Table 13.2. Warehouse Inventory.

j	N_j	\hat{p}_j (average value in dollars)
1	100	10,000
2	5000	500
3	25,000	5

Then from (13.30), we have

$$n_1 = \frac{100}{1 + \frac{500}{10,000} \frac{5,000}{100} + \frac{5}{10,000} \frac{25,000}{100}} = 28 \quad (13.31)$$

And, from (13.29),

$$n_2 = \frac{500}{10,000} \frac{5,000}{100} \times 28 = 68 \quad (13.32)$$

$$n_3 = \frac{5}{10,000} \frac{25,000}{100} \times 28 = 4 \quad (13.33)$$

13.3 The Saga of Happy Valley

It is common these days for these cities to attempt to annex affluent independent suburban communities. The reasons for doing so are various, but a major reason is to add the citizens of these communities to the real estate tax rolls of the city. Of course, once the suburb has been captured by the city, there may be some subsequent costs involved. For example, it may be necessary to add sewer lines and water mains, provide professional (as opposed to volunteer) fire protection, etc. One of the first things the tax assessor will want to do is to pick houses randomly from the suburb in an effort to determine the average value. In this way, the city authorities can make an estimate about the real estate tax revenue which incorporation of the suburb will add to the city's coffers. In Table 13.3 we display the results of such a sample. The Mayor Taxum of Monstropolis is deciding whether to begin proceedings to annex the upper middle income suburb of Happy Valley. The house prices may look relatively low, but these data are from the late 1980s.

Table 13.3. Home Values.

House Number	Estimated Value
1	\$107,500
2	\$329,000
3	\$274,500
4	\$174,000
5	\$87,500
6	\$495,500
7	\$295,000
8	\$310,000
9	\$290,500
10	\$478,000

The average of the sample is computed via

$$\bar{X} = \frac{1}{10} \sum_{i=1}^{i=10} X_i \quad (13.34)$$

where X_i is the estimated value of the i th house. The rationale for using \bar{X} as an estimate for the average value of a house in this suburban community might be

The ten houses we have examined randomly selected represent the entire suburban community in microcosm. Within this community, each house has probability 1/10 of being selected. If we had to pick one number to represent the idealized “representative house” from the suburb, \bar{X} seems a likely candidate. So, (13.35) represents the “world in a drop of water” estimate of the value of a house in the community.

The expression the “world in a drop of water” goes back to ancient notions which exist in most of the world’s major religions about some small part of the whole being, in a sense, equivalent in kind to the whole. For example, we find in the *Koran*:

5:32 That was why We laid it down for the Israelites that whoever killed a human being, except as a punishment for murder or other wicked crimes, should be looked upon as though he had killed all mankind; and that whoever saved a human life should be regarded as though he had saved all mankind.

From a data analytical perspective, we might seem to be going back to John Graunt’s treatment of the records of death in sixteenth-century London. Recall that analysts prior to Graunt had insisted on looking at each death individually, and so they could not grasp the big picture. But Graunt was able to “see the forest” by aggregating deaths into groups based on age at death. We are talking about grasping the details of the forest by looking at a small subset of the trees in the forest. There might well be 20,000 houses in the suburban community. We are not, at this time, aggregating the houses according to some stratification rule, such as “126 houses worth \$50,000 to \$75,000, 520 worth \$75,001 to \$90,000, etc.” Frequently, simple economics prevents us from grabbing the entire relevant data set. So, rather, we are saying that the ten houses we randomly selected have price characteristics representative of the full set of 20,000 houses in the suburban community.

The use of the *sample mean* \bar{X} as a representative value for a much larger set has gone on for as long as one can imagine. It is the basis of an area of data analysis called *sampling theory*. But, ancient though the use of \bar{X} is, the nuances of sampling theory are by no means trivial to grasp.

13.3.1 The Problem with Standards

Some years ago, a large manufacturer of trucks was deciding whether to accept a contract for 125,000 trucks from the U.S. Army. Everything looked reasonable in the contract except for one stipulation. The trucks were to be produced over a five year period at the rate of 25,000 per year. For each year of production, an inspector from the Army would pick one of the trucks and test it over a 100 mile stretch of road. The fuel consumption of the truck would be compared with the target value in the manufacturer's specifications. If the consumption rate was less than that of the specification, then the manufacturer would pay no penalty. However, were the consumption greater than the specification, then the manufacturer would pay a penalty given by subtracting the actual miles per gallon from the specified miles per gallon and dividing this into the expected life of the truck (100,000 miles) times \$1.50 per gallon.

Although the engineers at the plant were not statisticians, they felt that

1. The specification was unusual;
2. It could make the deal risky for the manufacturer.

Now the engineers were confident that the fleet average consumption of the 25,000 trucks would be better than the specified miles per gallon. If the penalty were based on an orderly recording of the miles per gallon of each truck, then they would be willing to sign off on the deal. But to base judgments on 25,000 trucks on the performance of just one of them seemed to be nonintuitive. Here, one truck is not a big enough drop of water to represent the world. The Army was unwilling to go to the expense of a record taking involving all the trucks.

How big must the sample of trucks be to come up with something fair? This question does not go back into the mists of prehistory. It was first considered by Karl Frederick Gauss in his lectures on the motions of comets given at Königsberg in 1809.

Gauss knew perfectly well how to take several estimates of speed and combine them to give one improved good estimate of speed. This is just the sample mean \bar{X} . But he did not know how to estimate the quality of his estimate. He tried various strategies. First, he considered looking at the average miss of each observation from \bar{X} :

$$\text{Average Error} = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X}). \quad (13.35)$$

Suppose the sample size were two and the true value of the comet's velocity were 1000 kilometers per hour. Suppose we consider two possible samplings. In the first, the measurements are 500 km/hr and 1500 km/hr. In the second, the measurements are 999 km/hr and 1001 km/hr. In both cases, the Average Error is zero. But, clearly, we feel much better about \bar{X}

results based on the second set of measurements, for they are less *variable*. Gauss noted that it was unacceptable to use a measure where positive and negative errors could cancel each other.

So Gauss sought alternative measures. One that he considered, and discarded, was the *Mean Absolute Deviation* (MAD):

$$MAD == \frac{1}{n} \sum_{i=1}^n |X_i - \bar{X}|. \quad (13.36)$$

Actually, this is not a bad measure. Why did Gauss reject it? One reason is the relative complexity of dealing with absolute values. In these days, when we have computers, the *MAD* makes a good deal more sense than in Gauss's. If we try and look at the problem of finding \hat{a} the value of a that minimizes

$$MAD == \frac{1}{n} \sum_{i=1}^n |X_i - a|, \quad (13.37)$$

the answer is not the sample mean \bar{X} , that wonderful estimate which goes back to prehistory. Moreover, the calculus developed separately by Leibnitz and Newton finds dealing with absolute values messy.

Gauss decided to use squared deviation of the observations from the sample mean as a practical measure of the confidence one might have in the quality of the estimate of how well "the drop represents the world." For example, consider the two cases above given for the estimation of the velocity of a comet. In the case where we have the two sets of observations above, our two estimates are

$$\begin{aligned} s_1^2 &= \frac{1}{2} [(500 - 1000)^2 + (1500 - 1000)^2] = 250,000 \\ s_2^2 &= \frac{1}{2} [(999 - 1000)^2 + (1001 - 1000)^2] = 1, \end{aligned}$$

respectively. For the first data set, most people would feel they had very little confidence that \bar{X} is close to the true velocity. For the second data set, we have reason to suppose that the true velocity probably is very close to 1,000 km/hr. For s^2 small, we probably feel that the drop of water is big enough to describe well the world. The statistic

$$s^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2 \quad (13.38)$$

is called the *sample variance*. The square root of s^2 is called the *sample standard deviation* s . Note that s and \bar{X} have the same units of measurement.

Of course, we need to quantify what we mean by the variance being small or large. Small or large relative to what? The usual way to determine this is

to divide s by \bar{X} . This quantity is called the *sample coefficient of variation*:

$$CV = \frac{s}{|\bar{X}|}. \quad (13.39)$$

So, for two cases above, we have

$$CV_1 = \frac{s_1}{\bar{X}} = \sqrt{\frac{250,000}{1,000}} = 15.81$$

$$CV_2 = \frac{s_2}{\bar{X}} = \sqrt{\frac{1}{1,000}} = .0316$$

Now, although it is not essential to follow convention, and sometimes it is even bad to do so, we are going, in the future, to follow the convention, in obtaining the sample variance of dividing through by $n - 1$ as opposed to n . Obviously, for n large, it will make little difference whether we divide by one or the other. But for n small, one advantage to dividing through by $n - 1$ rather than by n is that it gives us a larger value for s^2 . Thus, it makes us a bit more cautious in claiming that we have a "drop of water large enough to describe the world."

13.4 To Annex or Not

Returning to our example about annexation, the mayor of Monstropolis wants to be very sure before proceeding that the average house value in Happy Valley is over \$250,000 before proceeding with annexation. There are numerous disadvantages to incorporation of Happy Valley, not the least of which is that it may bring in tens of thousands of angry Republican voters to vote against him in the next election. If he is sure the average tax value for the houses in Happy Valley is at least \$250,000, he reckons it is worth the risk and costs to proceed with incorporation.

Performing the calculations, the mayor's statistician, Mr. Damnlys, quickly finds that the sample mean for house values is \$284,100, comfortably above the \$250,000 threshold. He tells the mayor that, indeed, it really looks as though Happy Valley houses average more than \$250,000. But Damnlys has not taken into account the variability of the data. We compute the sample standard deviation to be equal to \$136,295. And the sample coefficient of variation is a rather large .48. Probably, the mayor would like to be 95% sure that the Happy Valley average was at least \$250,000. Can we assure him of that?

13.4.1 The Bootstrap: The World in a Drop of Water

There are many ways we can try to answer Mayor Taxum's question. We are going to use the *bootstrap* here and quite a lot generally because

1. The bootstrap requires few assumptions;
2. The bootstrap intuitively gives us a practical feel as to whether “our drop of water” is large enough for the purpose at hand;
3. It is very computer intensive, replacing human effort by computer simulation effort.

Is the bootstrap “prehistoric” or of recent vintage? Here the answer is ambiguous. Practically speaking, it was little used before Julian Simon [6] used it in business courses in the 1960s (in those days, fast mainframe computers, such as the IBM 7040, were becoming rather common). Although Professor Simon used his resampling paradigm on a relatively main frame, it is generally available for use on most platforms at no cost. There is, for example, the freeware site **Statistics101.net**. Inasmuch as it is written in JAVA, Statistics101 will work on most platforms. A freeware manual is also included.

At the level of mathematical statistics, the “bootstrap” did not truly take off until Bradley Efron [4] and his colleagues at Stanford started massively demonstrating its use in the early 1980s. But it is also true that *bootstrap-like* algorithms go back at least to Block [1] in 1960 and Dwass [3] in 1957 when the first usable and generally available digital computer, the IBM 650, was on the scene.

When you consider it, a person who, in the 1920s, say, came up with a procedure which was practically impossible without a computer would have been very much like a person who today would assume the existence of an anti-gravity device as an essential part of a scheme for cheap transportation. It just did not happen. The author blushes to admit that he himself used a bootstrap-like scheme for economic forecasting purposes in an extensive consultation in 1971. The author blushes because one really cannot appropriately use bootstrapping for forecasting purposes. The bootstrap essentially works for one great purpose: assessing variation and its surrogates. As a practical matter, we will credit the bootstrap partly to Simon and particularly to Efron. In this course, we will be using a rather clever and compact computer package called *Resampling Stats*, which was developed by Simon and his associates.

13.4.2 Resampling and Histograms

Consider the house price data in Table 13.3. There are ten house prices in that “drop of water,” and we would like to know whether it is legitimate to tell Mayor Taxum that he may safely assume that the average house price in Happy Valley is at least \$250,000.

1. Consider first of all that our drop of water has become our mini-universe. We will treat the ten prices as the only values one can

possibly get. Each of the ten house prices will have probability one-tenth of being drawn.

2. Pick at random a house price from the list. It turns out, we pick house price number 3, namely \$274,500.
3. Pick at random, from the full list of ten, another house price. Again, we get number 3, namely \$274, 500.
4. Continue to resample from the full list of ten house prices until we have a sample of size ten. It turns out that our sample is \$107,500, \$329,000, \$274,500, \$274,500, \$174,000, \$87,500, \$495,500, \$495,500, \$310,000, \$290,000.
5. Compute \bar{X}_1 for this sample. It turns out to be \$283,890.
6. Carry out this operation 10,000 times, saving the entire list of $\bar{X}_1, \bar{X}_2, \dots, \bar{X}_{10000}$.
7. Sort these values and plot them in a histogram as shown in Figure 13.1.
8. See what percentage of the \bar{X}_i values lie to the left of \$250,000.
9. Graph the results

We will now take a Grauntian step, i.e., a step of aggregation. We will count all the \bar{X}_i which occur in the intervals (in increments of \$10,000) from \$135,000 through \$455,000. Note that, for example, the interval with center \$140,000 counts all \bar{X}_i values from \$135,000 through \$145,000. We note that we reexpress counts as percentages in the next column; simply divide each cell by 10,000 and multiply the resulting fraction by 100%. In the fourth column we accumulate the percentages. We note that 24% of the resampled \bar{X}_i values are less than \$250,000. We must report to Mayor Taxum that, assuming our sample of size ten is representative of the total house population of Happy Valley, there seems to be a 24% chance that the average value of houses there is less than \$250,000. tation of the sample *cumulative distribution function*, i.e., column four of Table 13.4.

$$CDF(x) = \hat{F}(x) = \frac{\text{number observations } \leq x}{\text{total number observations}} . \quad (13.40)$$

Table 13.4. Occurrence of 10,000 Bootstrapped X_i .

Interval Center	Number of X_i	Percent	Cumulative Percent
140000	2	0	0
150000	1	0	0
160000	7	0	0
170000	14	0	0
180000	34	0	1
190000	69	1	1
200000	110	1	2
210000	187	2	4
220000	289	3	7
230000	427	4	11
240000	576	6	17
250000	701	7	24
260000	873	9	33
270000	856	9	42
280000	1044	10	52
290000	900	9	61
300000	936	9	70
310000	700	7	77
320000	690	7	84
330000	472	5	89
340000	426	4	93
350000	275	3	96
360000	171	2	98
370000	86	1	99
380000	78	1	99
390000	44	0	100
400000	14	0	100
410000	13	0	100
420000	2	0	100
430000	1	0	100
440000	1	0	100
450000	1	0	100

What will Mayor Taxum decide to do? One can only conjecture. If he has other suburban areas which are more clearly ripe for the picking, he may leave Happy Valley alone, for the time being. Or he may decide to take a larger sample of estimated house values from Happy Valley. Clearly, taking things to the extreme, if he can obtain estimates for all the houses in Happy Valley, then he has eliminated any doubt as to the true value of the average of house values for Happy Valley. Generally speaking, however, we are driven by cost to take a sample much smaller than the entire relevant population.

Note that in Figure 13.1, we have essentially displayed a graphical representation of the third column in Table 13.4. Such a representation is called a *histogram*. Figure 13.1 gives a representation of the resampled house price averages in cumulative distribution format.

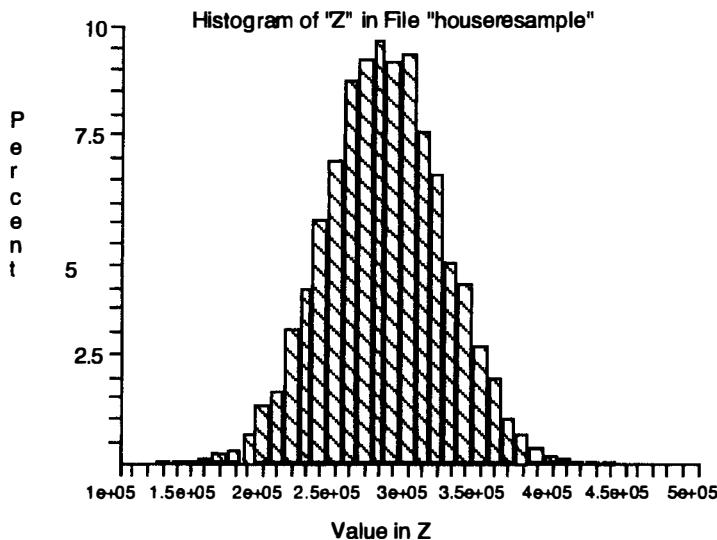


Figure 13.1. Resampled sample means of house prices.

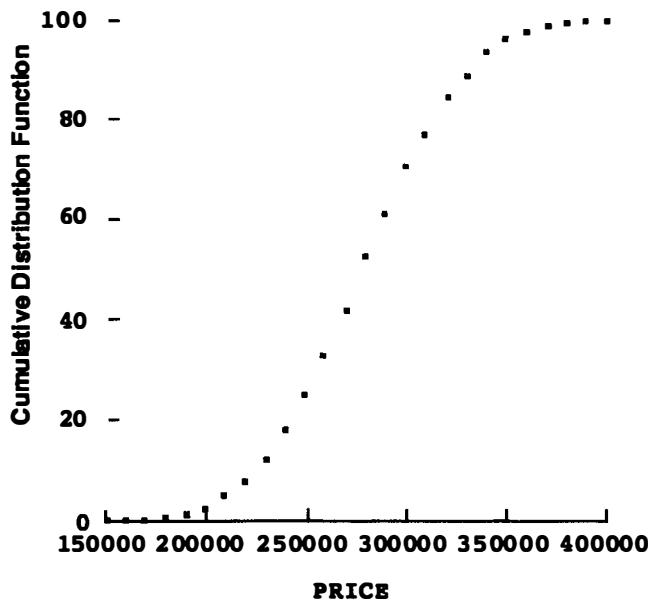


Figure 13.2. Cumulative distribution function of house prices.

'houseresample

'A sample of 10 houses in Happy Valley is taken, and the average valuation is
'283,890. How sure can we be that the true average house valuation is at least
'250,00?

'The sample values house valuations are listed in A.
'We use the "bootstrap" technique of drawing many bootstrap
'resamples with replacement from the original sample, and observing
'how the resample means are distributed.

```
Maxsize Default 50000
COPY (107500    329000    274500    174000    87500 495500
295000    310000    290500    478000)A
REPEAT 10000      'Do 10000 trials or simulations
SAMPLE 10 AB      'Draw 10 lifetimes from A, randomly and with replacement.
SUM B C
DIVIDE C 10 D      'Find the average lifetime of the 10
SCORE D Z      'Keep score
END
HISTOGRAM percent Z      'Graph the experiment results.
```

Figure 13.3. Resampling Stats program Happy Valley Houses.

There would appear to be around a 20% chance that the \$250,000 average is overly optimistic. Mayor Taxum's assistant, Mr. Damnlys, did not know anything about resampling. He advised the Mayor to go ahead and annex with confidence.

13.5 Using Sampling to Estimate Total Population Size

One relatively easy case would be the counting of penguins on an ice shelf in the Antarctic. Photographs from overflights would probably be fairly accurate for simply counting the penguins, particularly if these were taken at a time of day when the penguins were almost all at rest on the ice shelf.

Counting, say the number of robins in a county on a spring day would be much more difficult. One could very carefully compute the number of robins on a grid and then impute the density to be the same throughout the county. This would be a difficult and generally dubious assumption.

For endangered species—such as the Whooping Crane—there is a very careful and costly attempt to count nests around, say, Port Aransas, Texas, during nesting season. Fortunately, nest counts show that this species appears to be in a state of recovery.

Careful counting procedures have established the recovery of the American Bald Eagle. The California Condor's future is still in doubt, but even with major forest fire problems in the highlands of California, there appears to be evidence of recovery.

13.5.1 The United States Census

For the purpose of allocating congressional seats among the several states of the United States, the *United States Constitution* provides the method of *enumeration*. That means, an attempt should be made to count all U.S. citizens residing in the United States. The Indians who chose to live on tribal lands were treated as members of separate nations and were not to be counted. Any Amerinds who chose to live in the general taxed population were counted. The Amerind provision would seem to make clear, as to historical precedents, that the U.S. Census was to count American citizens and not nationals of other countries.

For the 2010 Census, the protocol has been changed to one of counting everybody, including, for example, members of a foreign sports team taking months long training in the USA. The headquarters of the United States Census has actually been moved to the East Wing of the White House, where its workings can be carefully monitored by the Obama Administration. It does not require too much cynicism to conjecture as to why this *de facto* change in the Constitution has been effected.

At this time, a majority of those persons who vote Democrat pay no income taxes at all. Nontaxpayers have a vested interest in voting for an entitlements oriented government. Aliens who reside in the United States make up well over 10,000,000 persons and tend to reside in areas where poorer Americans live. Therefore, the more illegal aliens one counts, the more Democrat congressmen and the more electoral votes for states that tend to vote Democrat (though not always: Texas will likely continue to vote Republican, and the counting of illegal aliens will increase, therefore, the weight of votes from a Republican state.)

Irregularities in voting in the United States are not uncommon. For example, in the Election of 1960, it may well have been the case that the political machinations of Mayor Daley in Chicago and Boss Parr in Duvall County, Texas, may have falsely given the election to Kennedy rather than Nixon. (In an earlier election, "missing votes" were presented by the Parr machine in which the voters voted in alphabetical order. In New Orleans there is the old joke that the tombs are built above ground to make it easier for the dead to walk to the polls and vote.)

13.5.2 Capture Recapture

Considering the wording of the Constitution, it really is not admissible for any procedure other than enumeration to be employed. That fact will not prevent creative wizardry going on in the East Wing for the counting of aliens residing illegally in the United States. One of the most popular suggestions being considered is that of "capture/recapture".

To start with, let us consider finding the total number of fish in a lake. We capture M fish out of a lake and tag them. The total (unknown) number

of fish in the lake is N . After a few days, we capture z fish from the lake, noting that x of these are tagged. Assuming no traumatization effect of the first capture (a big assumption) we can write

$$\frac{M}{N} \approx \frac{x}{z}. \quad (13.41)$$

Then we appear to have a natural estimate for N , namely

$$N \approx \frac{Mz}{x}. \quad (13.42)$$

Note, however, that if $x = 0$, we obtain an infinite estimate for N .

One popular alternative estimator is

$$\hat{N} \approx \frac{(M+1)(z+1)}{x+1}. \quad (13.43)$$

We need to find the expectation of $\frac{1}{x+1}$. Let us define $p = \frac{M}{N}$. Then, we have

$$E\left[\frac{1}{x+1}\right] = \sum_{x=0}^z \frac{z!}{(x+1)!(z-x)!} p^x (1-p)^{z-x}. \quad (13.44)$$

Going forward, we find after a few steps, letting $x = y - 1$,

$$\begin{aligned} E\left[\frac{1}{x+1}\right] &= \frac{1}{zp} \sum_{y=1}^{z+1} \frac{(z+1)!}{y!(z+1-y)!} p^y (1-p)^{z+1-y}. \\ &= \frac{1}{pz} [1 - (1-p)^{z+1}] \approx \frac{N}{Mz}. \end{aligned} \quad (13.45)$$

Returning to (13.44) we have that

$$E(\hat{N}) \approx N \frac{M+1}{M} \frac{z+1}{z} \approx N. \quad (13.46)$$

Further work shows that

$$Var(\hat{N}) \approx \frac{(z+1)(M+1)(z-x)(M-x)}{(x+1)^2(x+2)}. \quad (13.47)$$

Now, we are ready to go through a “practical” example. Suppose we go to an area where many illegal aliens are thought to live. We use sound trucks and flyers to inform the inhabitants that for the next week the area will exclude INS agents and any other law enforcement official seeking to capture illegal aliens. Then the Hispanic Census pollsters go into the neighborhood and start looking for persons to be counted. Perhaps they offer \$5 to anyone who will fill out a verbally administered census form. We give a metal token to the person interviewed. After a few days the Census workers return again

offering \$5 to anybody who will fill out the verbally administered census. If the person produces the metal token, we award him or her another \$5. The results of the poll are 500 persons found on the first census. On the second poll 1,000 persons are found. Of these 25 have the token.

Our estimate for the illegal aliens is given by

$$\hat{N} = \frac{(M+1)(z+1)}{x+1} = 19,289 \quad (13.48)$$

$$Var(\hat{N}) = \frac{(1001)(501)(1000-25)(500-25)}{(26)^2(27)} = 12,725,052 \quad (13.49)$$

How reliable would such an estimate be? Not very. Many will hide. Typically, each person paid several thousands to get here. Even the captured fish would be sensitive to a second capture. And people are much more clever than fish. Anything except enumeration is unconstitutional and, almost certainly, statistically invalid.

Problems

13.1 In a minor league, surprise drug testing was administered to 200 players, categorized by the number of homeruns hit last season.

Table 13.5. Home runs and Steroids.

Number Home runs	Number	Number in Sample	Testing Positive
More than 30	40	30	18
15 \leq 30	200	75	15
Number \leq 14	210	95	2

(a) What are the estimates for the overall percentage of drug use in the league?

(b) What is the estimate for overall variance of the percentage?

(c) Using what you have seen from this league, using minimum total variance, find the number to be sampled in each of the three homerun groupings from another league if the total sample size is 200.

13.2. In an attempt to estimate the total number of large mouth bass in a large lake, 523 are captured and tagged. Then after a week, a sample of size 724 is captured. Of these 72 are found to be tagged. Estimate the number of large mouth bass in the lake and also the variance of your estimate.

13.3. A chain of jewelry stores is attempting to estimate the value of diamond rings for *ad valorem* tax purposes. The chain wishes to minimize

the variance of the estimate. Give your best guess as to the number of rings to be sampled from each store if the total sample size is 100.

Table 13.6. Diamond Rings.

Store Number	Guess of Average Ring Value	Number Rings in Store
I	\$1000	250
II	\$2000	300
III	\$10,000	200

13.4. The Bayou City Foods chain of grocery stores is well known for its sales. Some have conjectured that the chain fails to enter the sales prices in its computer and overcharges its customers. This would be a version of “bait and switch,” and, if true, is illegal. The legal firm of Sueem, Sueem, and Settle has decided to conduct a sample survey of 35 randomly selected branches of a basket of \$50 to see whether a class action of behalf of customers of Bayou City Foods might be indicated. Use a bootstrapping simulation to determine whether there is a *prima facie* case for “bait and switch, using the resulting prices the sample customers actually paid \$50.00 \$43.14 \$50.00 \$48.75 \$50.00 \$43.71 \$50.00 \$50.00 \$56.13 \$50.00 \$50.00 \$50.00 \$48.35 \$56.75 \$50.00 \$49.10 \$50.00 \$57.85 \$61.12 \$50.00 \$50.00 \$48.75 \$52.35 \$58.06 \$50.00 \$50.00 \$49.44 \$56.12 \$50.00 \$48.74 \$58.98 \$50.00 \$51.23 \$50.00 \$51.12.

References

- [1] Block, J. (1960). “On the number of significant findings to be expected by chance,” *Psychometrika*, **25**, 158–168.
- [2] Cochran, W. G; Mosteller, F. and Tukey, J. W (1954). “Statistical problems of the Kinsey report on sexual behavior in the human male.” *Journal of the American Statistical Association*.
- [3] Dwass, M. (1957). “Modified randomization sets for nonparametric hypotheses,” *Annals of Mathematical Statistics*, **29**, 181–187.
- [4] Efron, B. (1979). “Bootstrap methods: another look at the jackknife,” *Annals of Statistics*, **8**, 1–26.
- [5] Kinsey, A. C. et al. (1948/1998). *Sexual Behavior in the Human Male*. Philadelphia: W.B. Saunders; Bloomington, IN: Indiana U. Press.

[6] Simon, J.L. (1990). *Resampling Stats*. Arlington, Va.: Resampling Stats, Inc.

[7] Walker, L. and Monahan, J. (1998). "Sampling damages," *IOWA Law Review*, v. 83, March 1998, pp. 545.

Chapter 14

The Stock Market: Strategies Based on Data versus Strategies Based on Ideology

14.1 Introduction

Many young professionals have the false notion that their futures will be secure if they simply excel in their professions. Generally speaking, unless they invest from their salaries in a fairly regular fashion, they will arrive at retirement with very little to show for it. One, generally speaking, does not get wealthy by socking away salary money in Treasury Bills or certificates of deposit. Inflation and taxes will lead them over the cliff. One should set aside a portion of one's salary monthly for investment in real estate, a new on the side venture, sound but high paying bonds, common stocks, etc. In this chapter, we shall restrict ourselves to buying publicly traded common stocks.

How shall one invest in the market of common stocks? There are well over 10,000 publicly traded stocks. Not surprisingly, there are many opinions. Some of these are quite bizarre. One tenured professor of economics from a major university appeared on a major television talk show some years ago. He was invited, since he had correctly forecasted a temporary crash in the market. As the program proceeded, to the apparent horror of the host of the program, the professor revealed that his insights were based

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

on the insights of his guru, who was some sort of Hindu astrologer. In this chapter, we will not include such bizarre forecasters, although the efficient market academics have sometimes some of the flavor of the unnamed guru.

The most venerable school is that of the "value investors" who base their strategies by seeking underpriced stocks. Perhaps the most famous of these is Warren Buffett, whose Berkshire-Hathaway portfolio has a 40 year track record of greater than 20% gain per year. Value investors tend to follow the strategy of Graham and Dodd, who carefully examined the internal statistics of companies to attempt to find those which appeared to be undervalued. This approach is not based on ideology or the configuration of the stars, but on actual data. Naturally, the huge number of stocks cries out for some sort of orderly computer analysis, but that is not so easy when one is examining the quality of management, working conditions, worker morale, market opportunities, etc. In large measure, Buffett disdains the use of computer analysis. But it is not necessary that one scan through thousands of balance sheets. One need not seek for a sharp optimum. Indeed, seeking for a sharp optimum is frequently a fragile approach. There are plenty of good stocks around. It is not necessary to scan all of them. Of course, the young professional could simply trust in the insights of somebody such as Buffett with a long and impressive track record. Buffett's company, Berkshire-Hathaway has much to recommend it. However, if one invests in Berkshire-Hathaway, he or she should be aware that they are investing in a portfolio of stocks which typically has a value, in terms of assets held, which may be less than 75% of the cost of the Berkshire-Hathaway stock. Part of the cost of the stock is the quite reasonable high value placed on the wisdom of "the genius from Omaha." A decision to invest part of one's assets in Berkshire-Hathaway stock, might be very wise. However, no mortal lives forever, and it is a fair question one must ask, "What happens to the value of a share of Berkshire-Hathaway when Buffett leaves the scene?" Successors of geniuses in investment firms usually fall short of their predecessors. Therefore, that the prudent young professional should take interest in managing his or her own retirement investments.

Of course, if one wishes to invest on "autopilot" there are ways to do so. John Bogle has effectively argued [4] that the value of investment counsellors is, in general, not worth their fees. Many years ago, he founded the Vanguard S&P 500 fund (among others) which maintains a portfolio balanced according to the market cap values of each of the members of the Standard and Poor selected basket of top 500 stocks. Thus the weight of investment in the i^{th} stock would be

$$w_i = \frac{V_i}{\sum V_j}, \quad (14.1)$$

where V_i is the total market value of all the stocks in company i . Interestingly, Bogle's strategy is actually very close to the "total market index fund" suggested by Nobel laureate William Sharpe.

This brings us to the “Efficient Market Hypothesis” (EMH). According to this theory, stock purchasers in the aggregate, at any given time, impute to each of the more than 10,000 publicly traded stocks its accurate value. It assumes that the current price information, although obviously a one dimensional vector, correctly sums up such factors as the past trading prices of the stock, the past trading prices of all the other stocks, the market situation, the political environment, the relative quality of management of the companies, etc. To many data analysts, the EMH appears too simplistic. In a sense, it smacks of the long defunct Marxist idea that value of an item consists of the labor which went into its creation.

Nevertheless, the EMH enjoys overwhelming dominance in most schools of business in the United States. Although the author is in strong disagreement with the EMH, he will devote some time to it, for readers need to be aware of its jargon and its consequences some of which have been quite destructive.

A third paradigm is that of the technical analyst (chartist, momentum trader, etc.). Like the EMH advocate the technical analyst will generally not pay close attention to the operational details of a company. Rather, the technical analyst will try and extrapolate from the historical records of stocks. For example, the technical analyst will generally prefer a stock which has gone up for ten straight trading days to one which has declined for ten straight trading days. Thanks to the advent of the high storage high speed computer, it is possible to build “expert systems” which can train themselves in the utilization of massive data sets for the purposes of making high returns for the user using constraints on the extrapolated downside risk.

14.2 Markowitz’s Efficient Frontier: Portfolio Design as Constrained Optimization

We will now consider strategies for reducing the risk to an investor as a well defined constrained optimization problem. The argument follows roughly that of Markowitz [7–8]. Let us suppose that the *proportional gain* of a security is given the symbol X_i . Here we are not necessarily assuming X_i is the growth rate of the stock. X_i may include as well, for example, the dividends accruing to the security. And we will treat X_i as a Gaussian (normal) random variable rather than as a constant. We will then assume that the average (expected value) of X_i is μ_i , and its *variance* is σ_i^2 . Then we shall assume that the *covariance* of X_i and X_j is given by σ_{ij} (alternatively, that the *correlation* is given by ρ_{ij}). That is, we assume that:

$$\begin{aligned} E(X_i) &= \mu_i \\ E(X_i - \mu_i)^2 &= \sigma_i^2 = \sigma_{ii} \end{aligned}$$

$$\begin{aligned} E[(X_i - \mu_i)(X_j - \mu_j)] &= \sigma_{ij} \\ \rho_{ij} &= \frac{\sigma_{ij}}{\sigma_i \sigma_j}. \end{aligned}$$

Our portfolio will be formed by a linear combination of n stocks where the fraction α_i of our portfolio will consist of shares of stock i . Clearly, it would be a fine thing to have a high value of

$$\mu_{ave} = \sum_{i=1}^n \alpha_i \mu_i. \quad (14.2)$$

On the other hand, we would like for the portfolio to be as close to a sure thing as possible, i.e., we would like to minimize the volatility of the portfolio

$$\begin{aligned} S &= E[\sum_{i=1}^n \alpha_i X_i - \sum_{i=1}^n \alpha_i \mu_i]^2 \\ &= \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j \sigma_{ij}. \end{aligned} \quad (14.3)$$

Clearly, there is a problem, for minimizing S would logically drive us to something like Treasury Bills, which strategy is not historically very good for maximizing μ_{ave} . It might be postulated that what should be done is to ask an investor what μ_{ave}^* he requires and then design a portfolio obtaining a mix of the n stocks so as to minimize S for that target value. This is not a very natural posing of the problem from the standpoint of the investor (picking a μ_{ave}^* is not natural for most). However, from the standpoint of the mathematician, it is a formulation easily solved. To see that this is so, consider the Lagrange multiplier formulation which seeks to minimize:

$$Z = \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j \sigma_{ij} + \lambda_1 (\sum_{i=1}^n \alpha_i \mu_i - \mu_{ave}^*) + \lambda_2 (\sum_{i=1}^n \alpha_i - 1) \quad (14.4)$$

Differentiating partially with respect to each of the α_i , λ_1 , and λ_2 and setting the derivatives equal to zero gives us the $n + 2$ equations (linear in $\alpha_1, \alpha_2, \dots, \alpha_n, \lambda_1, \lambda_2$):

$$\begin{aligned} \frac{\partial Z}{\partial \alpha_i} &= 2 \sum_{i=1}^n \alpha_i \sigma_i^2 + 2 \sum_{j>i} \alpha_j \sigma_{ij} + \lambda_1 \mu_i + \lambda_2 = 0 \\ \frac{\partial Z}{\partial \lambda_1} &= \sum_i^n \alpha_i \mu_i - \mu_{ave}^* = 0 \\ \frac{\partial Z}{\partial \lambda_2} &= \sum_{i=1}^n \alpha_i - 1 = 0. \end{aligned}$$

This is an easy formulation to solve in the current era of cheap high speed computing. Naturally, as formulated here, it is possible that some of the α_i may go negative, though it is easy to impose the additional constraint that all α_i be nonnegative. Furthermore, it will happen that the α_i will yield fractional shares of stocks. Generally rounding will give us a satisfactory approximation to the solution, though we can easily impose the restriction that all shares be bought as integer lots. All these little details can be dealt with easily. However, the formulation is not particularly relevant in practice, for few individuals will find it natural to come up with a hard number for μ_{ave}^* .

Well then, we could also pose the problem where we maximize the gain of the portfolio subject to some acceptable level of the volatility. But this is also not (for most people) a natural measure of portfolio riskiness. An investor probably would like to know his/her probabilities of achieving varying levels of value as time progresses. This is not an easy task. Indeed, the assumption of the equivalence of risk and variance is a dramatic oversimplification.

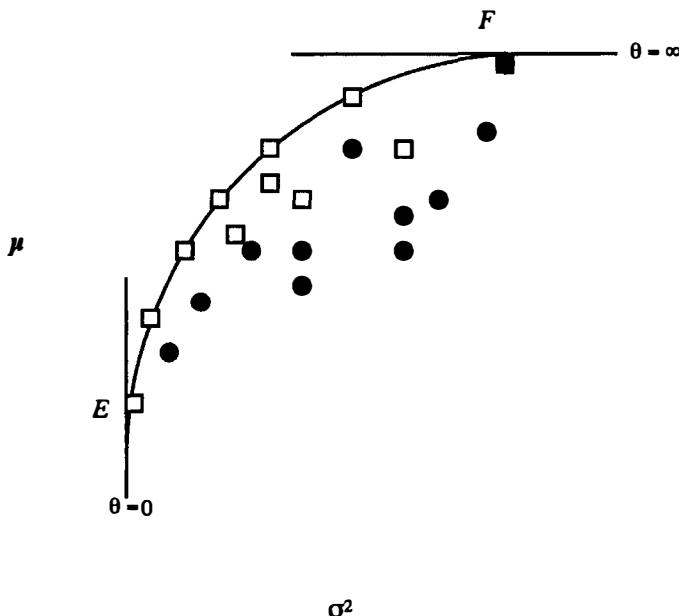


Figure 14.1. Markowitz's efficient frontier.

The set of all portfolios with maximum expected gain at a given level of volatility (or minimum volatility at a given level of expected gain) was de-

fined by Markowitz as the *efficient frontier*. His basic method, which is fairly similar conceptually to the other techniques discussed in this section, can perhaps be understood by reference to Figure 14.1. Here the dots represent security parameters and the boxes represent portfolio parameters. Markowitz set about to minimize a function of the type $\sigma^2 - \theta\mu$. By initializing the procedure at $\theta = \infty$, the highest return security (F) is obtained. Note that, because diversification cannot increase return, the highest return portfolio will be composed entirely of the highest return security. From this point, Markowitz employed a quadratic programming algorithm to trace the efficient frontier by allowing θ to decrease to 0 (at which point E , the minimum variance portfolio is obtained). In actuality, the iterative procedures only determine "corner" portfolios, which are those points at which a security enters or leaves the efficient portfolio. The efficient frontier between two corner portfolios is a linear combination of the corner portfolios. Aside from the objective function, these techniques also generally involve constraints, such as the requirement that the weights assigned to securities be nonnegative and/or sum to one. Now, it would appear that, according to the EMH, nobody would want to buy the black circle stocks in Figure 14.1. But people do. The assumption of some sort of instantaneous group wisdom is the kind of *deus ex machina* not borne out in the real world. Actually, if one looks at the relationship between growth and volatility (standard deviation) for large cap stocks over the period from 1926 through 2000, one finds the correlation to be, not positive, but rather $-.317$ as demonstrated in Figure 14.2

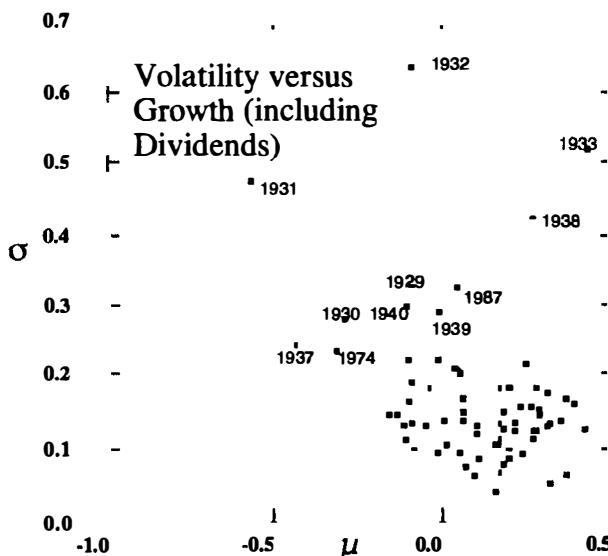


Figure 14.2. Historical relation between μ and σ .

14.3 Sharpe's Super Efficient Frontier: The Capital Market Line (CML)

Building upon the work of Markowitz, William Sharpe, formed the philosophical underpinnings of Index Fund Investment.

If we may assume that investors behave in a manner consistent with the EMH, then certain statements may be made about the nature of capital markets as a whole. Before a complete statement of capital market theory may be advanced, however, certain additional assumptions must be presented:

1. The μ and σ of a portfolio adequately describe it for the purpose of investor decision making [$U = f(\sigma, \mu)$].
2. Investors can borrow and lend as much as they want at the riskless rate of interest.
3. All investors have the same expectations regarding the future, the same portfolios available to them, and the same time horizon.
4. Taxes, transactions costs, inflation, and changes in interest rates may be ignored.

Under the assumptions above, all investors will have identical opportunity sets, borrowing and lending rates ($r_L = r_B$) and, thus, identical optimal borrowing-lending portfolios, say X (see Figure 14.3). Because all investors will be seeking to acquire the same portfolio (X), and will then borrow or lend to move along the *Capital Market Line* in Figure 14.3, it must follow for equilibrium to be achieved that all existing securities be contained in the total market portfolio (X). In other words, all securities must be owned by somebody, and any security not initially contained in X would drop in price until it did qualify. Therefore, the portfolio held by each individual would be identical to all others and a microcosm of the market, with each security holding bearing the same proportion to the total portfolio as that security's total market value would bear to the total market value of all securities. In no other way could equilibrium be achieved in the capital market under the assumptions stated above.

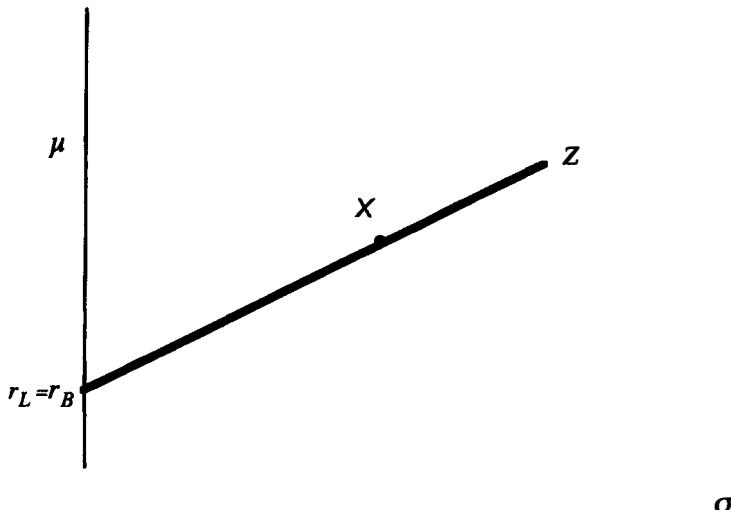


Figure 14.3. The capital market line (CML).

The borrowing-lending line for the market as whole is called the *Capital Market Line*. The securities portfolio (X) employed is the total universe of available securities (called the *market portfolio*) by the reasoning given above. The CML is linear and it represents the combination of a risky portfolio and a riskless security. One use made of the CML is that its slope provides the so-called *market price of risk*, or, that amount of increased return required by market conditions to justify the acceptance of an increment to risk, that is

$$\text{slope} = \frac{\mu(X) - r}{\sigma(X)}.$$

The simple difference $\mu(X) - r$ is called the *equity premium*, or the expected return differential for investing in risky equities rather than riskless debt.

This very elegant result of Sharpe indicates that one simply cannot do better than invest along the Sharpe Superefficient Frontier (CML). Unfortunately, a backlook at 50,000 randomly selected portfolios from the 1,000 largest market cap stocks over a period of 40 years shows that over half lie above the CML. How it has been that EMH enthusiasts apparently failed to crunch the numbers is a matter of conjecture. Nor is this result surprising, since the Standard and Poor Index fund over this period has averaged a return of somewhat in excess of 10% while Buffett's Berkshire-Hathaway has delivered well over 20%.

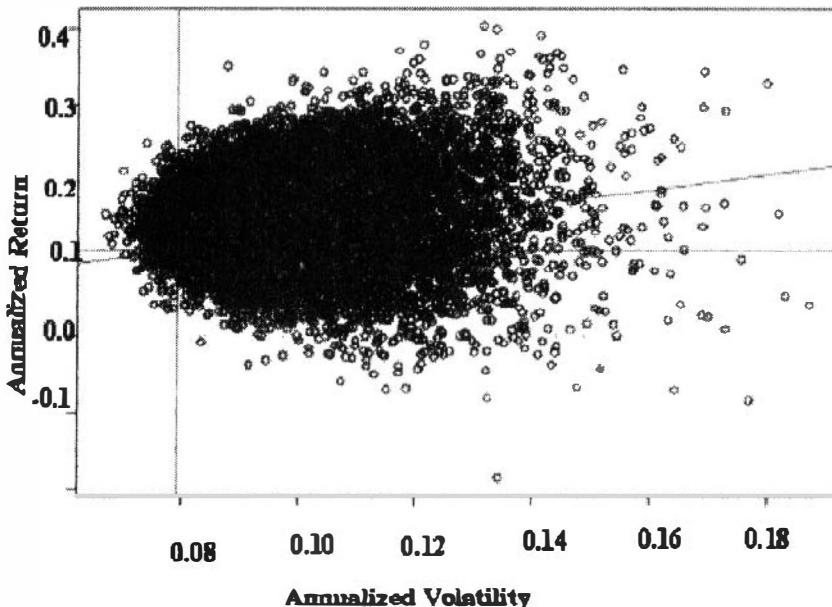


Figure 14.4. Randomly selected portfolios in 1993 beating the super efficient frontier portfolios.

14.4 The Security Market Line

One major question raised by CML analysis involves the means by which individual securities would be priced if such a system were in equilibrium. Throughout this chapter, we generally do not assume that markets are in equilibrium. Therefore, when we combine securities into a portfolio that has the average return of the component securities but less than the average risk, we simply ascribe this *gain from diversification* to our own shrewdness. In the type of efficient market assumed by the EMH enthusiasts, everyone will be doing the same thing, and the prices of securities will adjust to eliminate the windfall gains from diversification.

Sharpe [9,10] has suggested a logical way by which such security pricing might take place. If everyone were to adopt a portfolio theory approach to security analysis, then the risk of a given security might be viewed not as its risk in isolation but rather as the change in the total risk of the portfolio caused by adding this security. Furthermore, because capital market theory assumes everyone to hold a perfectly diversified (that is, the market) portfolio, the addition to total portfolio risk caused by adding a particular

security to the portfolio is that portion of the individual security's risk that cannot be eliminated through diversification with all other securities in the market.¹

Because the concept of individual security pricing is rather elusive, let us restate it. Sharpe argued that the price (and thus return) of a given security should not be determined in relation to its total risk, because the security will be combined with other securities in a portfolio and some of the individual risk will be eliminated by diversification (unless all the securities have correlation equal to 1). Therefore, the return of the security should only contain a risk premium to the extent of the risk that will actually be borne (that is, that portion of the total risk which cannot be eliminated by diversification—which is variously called *nondiversifiable risk* or *systematic risk*).

If this logic is accepted, it is then possible to generate a Security Market Line as shown in Figure 14.5, where the return on individual securities is related to their covariance with the market. If capital markets are in equilibrium and all the other assumptions of this chapter hold, then the parameters of each security should lie on the SML. Furthermore, because the risk of a portfolio is the weighted sum of the nondiversifiable risk of its component securities, all portfolios should also fall on the SML in equilibrium. (It should be noted that Sharpe's theory indicates that all portfolios will fall on the SML, and as a general rule, no individual securities or portfolios, should lie on or above the CML.) Of course, the inelegant reality should trump the elegant theory. Data analysis shows that stocks do not lie on the SML and it is not true that all stocks and portfolios like below the CML. If one can do better, on the average, than investing in the total market portfolio by random portfolio selection, then it should not be surprising that Buffett's value investing was so successful. Nor should it be surprising that we can beat index portfolios with astounding regularity by using technical coordinated momentum analysis. One might have hoped that the EMH enthusiasts would have done a bit of plotting of real world data to see whether it be true that all securities lie on the SML. Unfortunately, as is their custom, they found it unnecessary to see whether the real world conforms to their suppositions.

¹If the standard deviation of the market as whole is σ_M and the standard deviation of security i is σ_i , and the correlation of security i with the market is ρ_{iM} , then the non-diversifiable portion of the individual security's risk is the covariance of returns between the security and the market as a whole, i.e., $C_{iM} = \sigma_i \sigma_M \rho_{iM}$.

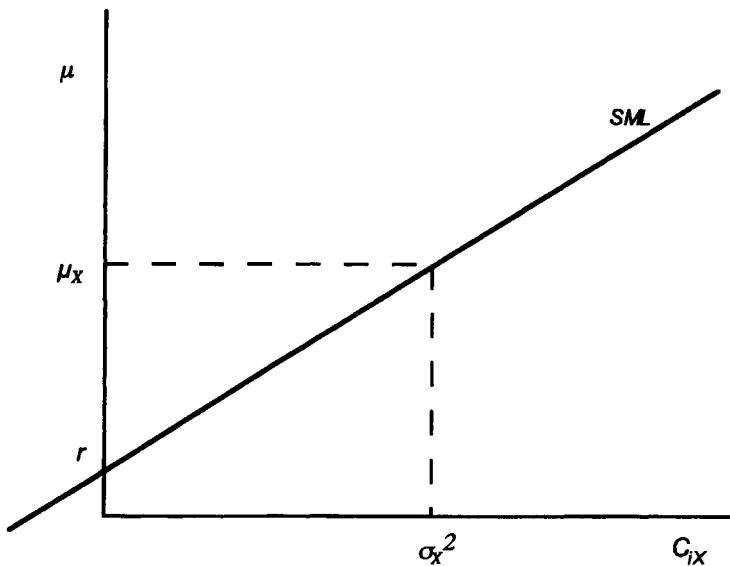


Figure 14.5. The security market line (SML).

14.5 The Sharpe Diagonal Model

Although Sharpe's theory does not lead to anything close to optimality, it is useful in that it assumes that the return on a security may be related to an index (such as the DJIA, S&P 500, Wilshire 5000, or whatever) as follows:

$$\text{Return}_i = a_i + b_i \text{Return}_I + c_i \quad (14.5)$$

$$\mu_i = a_i + b_i \mu_I + c_i$$

where:

a_i and b_i are constants,

μ_I is the return (including dividends) on the index,

c_i is an error term with $\mu_{c_i} = 0$ and $\sigma_{c_i} = \text{a constant}$.

It is further assumed that c_i is not correlated with μ_I , with itself over time, nor with any other security's c (the last implying that securities are only correlated through their common relationship to the index). Therefore, μ_i can be estimated as $(a_i + b_i \mu_I)$. The parameters a_i and b_i can either be estimated, computed by regression analysis, or both. Furthermore, σ_{c_i} can be viewed as the variation in μ_i not caused by variation in μ_I . The values a_i and b_i are referred to as Sharpe's alpha and beta, respectively. When people away from detailed information about stock values other than the increase in one of the major indices, such as the DOW or the S&P 500, hear that these indices have gone up by 2% day-to-day, they may well heave a

sigh of relief, for the stock values in the market are generally correlated with the values of the major indices.

The return of the portfolio becomes:

$$\begin{aligned}\mu &= \sum_{i=1}^n \alpha_i (a_i + b_i \mu_I + c_i) \\ &= \sum_{i=1}^n \alpha_i (a_i + c_i) + \left(\sum_{i=1}^n \alpha_i b_i \right) \mu_I\end{aligned}\quad (14.6)$$

where the first term is viewed as an investment in the essential nature of the securities, and the second term is an investment in the index. The risk of the portfolio is:

$$\sigma = \sqrt{\sum_{i=1}^n (\alpha_i \sigma_{c_i})^2 + \left(\sum_{i=1}^n \alpha_i b_i \right)^2 \sigma_I^2} \quad (14.7)$$

where, again, the first term under the radical may be viewed as the risk of the portfolio attributable to the particular characteristics of the individual securities, and the second term as the risk attributable to the index.

Thus, the Sharpe model simplifies the input problem by making it directly amenable to simple regression analysis. In addition, by assuming that securities are only related through the index, the nonzero elements in the covariance matrix are reduced to those on the diagonal, thus easing the computational burden.

To illustrate the Sharpe approach, assume that the index currently stands at 1000 and, with reinvestment of dividends, it is expected to be at 1100 at the end of the year. Given the following data, suppose we wished to determine portfolio μ and σ for $\alpha_1 = .2$, $\alpha_2 = .5$ and $\alpha_3 = .3$.

$$\begin{aligned}\sigma_I &= 0.10 \\ \mu_1 &= 0.06 + 0.1\mu_I; \sigma_{c_1} = 0.03 \\ \mu_2 &= -0.03 + 2\mu_I; \sigma_{c_2} = 0.20 \\ \mu_3 &= 0.00 + \mu_I; \sigma_{c_3} = 0.10.\end{aligned}$$

Employing the above we obtain:

$$\begin{aligned}\mu &= (.2)(.06) + (.5)(-.03) + (.3)(.00) + [(0.2)(.1) + (0.5)(.2) + (0.3)(1)](.10) \\ &= .012 - .015 + (1.32)(.10) = .129 \text{ or } 12.9\%.\end{aligned}$$

Employing (14.7),

$$\begin{aligned}\sigma &= \sqrt{[(0.2)(0.03)]^2 + [(0.5)(-0.2)]^2 + [(0.3)(0.0)]^2 + [(0.2)(0.1) + (0.5)(0.2) + (0.3)(1)]^2} \\ &= \sqrt{(0.006)^2 + (.1)^2 + (.03)^2 + (1.32)^2(.1)^2} \\ &= \sqrt{.000036 + .01 + .0009 + .017424} = \sqrt{.02836} = .168 \text{ or } 16.8\%.\end{aligned}$$

It is also possible to discuss the SML in terms of Sharpe's index model.

$$\mu_i = a_i + b_i \mu_I + c_i. \quad (14.8)$$

The b_i term (called Sharpe's *beta coefficient*), given μ_I , is equal to:

$$(b_i | \mu_I) = \frac{C_{iI}}{\sigma_I^2} = C_{iI} \frac{1}{\sigma_I^2} \quad (14.9)$$

which, if the index is a valid depiction of the market:

$$(b_i | \mu_I) = \frac{C_{iX}}{\sigma_X^2} = C_{iX} \frac{1}{\sigma_X^2}. \quad (14.10)$$

Under these assumptions, the abscissa of a point on the SML expressed in terms of b_i is merely $1/\sigma_X^2$ times that of the same point expressed in terms of C_{iX} and the two are directly comparable. Viewed another way, the risk premium an individual security would exhibit in equilibrium is:

$$\mu_i - r = \frac{\mu_X - r}{\sigma_X^2} C_{iX} = (\mu_I - r) b_i. \quad (14.11)$$

A major advantage of transferring the discussion into beta terminology is that the regression coefficient can be used directly to estimate the systematic risk of the asset. Unfortunately, the beta concept also possesses serious pitfalls. In the first place, its very simplicity and popularity cause it to be used by many who fail to understand its limitations. Because the concept is subject to all the assumptions of both linear regression and the efficient capital market hypothesis, statistical problems and economic imperfections may undermine its usefulness. Many investors are unaware of these limitations and have blithely assumed that one need only fill a portfolio with securities possessing large betas to get high returns. At best, the beta is a risk-measure surrogate and *not* an indicator of future returns. The idea that the assumption of large amounts of risk will generate large returns only approaches being correct over the long run in reasonably efficient markets in equilibrium. Even then it ignores utility considerations. A further difficulty with the beta concept follows from empirical findings that betas for small portfolios (and, of course, individual securities) over short periods can be highly unstable over long holding periods where, of course, beta approaches one by definition anyway. It would thus appear that one of the few valid applications of the beta concept would be as a risk-return measure for large portfolios. An example of how betas can be used in this regard is presented in the next section.

14.6 Portfolio Evaluation and the Capital Asset Pricing Model (CAPM)

Several measures directly related to capital market theory have been developed for the purpose of portfolio evaluation. The latter is essentially a

retrospective view of how well a particular portfolio or portfolio manager did over a specified period in the past. Most of the published research in this area has dealt with mutual funds, seemingly because they are controversial, economically important in certain financial markets, and possessed of long life with readily available data. Much of the early work in this area (including the advertisements of the funds themselves) was of a simple time series nature, showing how well an investor could have done over a given period in the past if he/she had invested then or else comparing these results to what the investor could have earned in other funds, or the market as a whole. The more recent work considers both return and its variability, contending that mutual funds that invest in riskier securities should exhibit higher returns. One result of this work, considered subsequently, has been the finding that investors do as well or better on average by selecting securities at random as they could with the average mutual fund. Another implication, of more relevance here, is the growing feeling that the managers of any kind of portfolio should be rated not on the return they earn alone, but rather on the return they earn adjusted for the risk to which they subject the portfolio and their management fees.

Before proceeding, however, a caveat is in order about the nature of *ex post* risk and return measures. As in any problem in measurement, one must delineate (1) why a measurement is being made, (2) what is to be measured, (3) which measurement technique is appropriate, and (4) the import of the results of the measurement. If one is not careful *ex post* return measurements can easily result in the "if only I had" syndrome, which is a waste of time and effort as far as making an investment in the present is concerned. For such measures to be of use, one must assume that the ability of a manager or fund to earn greater-than-average returns in the past is some indication of ability to do so in the future. As the empirical work cited below indicates, there is little evidence to support this contention. As far as risk is concerned, there is some doubt about what the concept of *ex post* risk means. Most of the writers in this area are careful to stress the term "return variability" instead of risk per se. Because the outcomes of all past events are currently known with certainty, the use of return variability as a measure of risk in this instance involves a different notion of risk than we have been using. Again, to make operational investment decisions, it would seem necessary to assume that past risk-return behavior of managers or portfolios either could or would be maintained in the future.

Deferring judgment for the moment on the above reservations, let us consider the proposed evaluation measures. Sharpe [10] has proposed the use of a reward-to-variability ratio related to the slope of the capital market line:

$$\text{Sharpe's Measure for } i\text{th portfolio} = \frac{\mu_i - r}{\sigma_i}. \quad (14.12)$$

In effect, Sharpe is computing the slope of the borrowing-lending line going through the given portfolio and arguing that a greater slope is more

desirable.

A second measure on the SML is that of Treynor [14]:

$$\text{Treynor's Measure for the } i\text{th portfolio or security} = \frac{\mu_i - r}{b_i} \quad (14.13)$$

and the line in (beta, return) space = $r + (\mu_i - r)/(b_i)$ = *characteristic line* of security or portfolio i . Treynor's methodology is fairly similar to that of Sharpe, except that by using the SML instead of the CML, the Treynor measure is capable of evaluating individual security holdings as well as portfolios. A disadvantage is that the accuracy the rankings depends in part upon the assumption (implicit in the use of the SML) that the fund evaluated would have been held in an otherwise perfectly diversified portfolio.

A third measure, also based on the SML but different from Treynor's, is that of Jensen [14]:

$$\text{Jensen's Measure} = (\mu_i - r) - b_i(\mu_X - r). \quad (14.14)$$

This measure is expressed in units of return above or below the riskless rate of a line drawn through the parameters of the security or portfolio parallel to the SML. This measure does allow comparisons of a portfolio to the market and is also amenable to estimation by regression; because of its treatment of differential risk, however, direct comparisons between funds or portfolios generally cannot be made. Furthermore, it has been suggested that all three of the above measures are biased against high risk portfolios by failing to recognize the inequality of borrowing and lending rates and the resulting nonlinearity of the SML and CML.

Use of geometric means as an evaluation tool should not be overlooked as well. Over a given period of time, the geometric mean portfolio return could be compared to that of other portfolios or some market index. There are several advantages to such a measure. Assuming that the interval considered is "sufficiently" long (and if it is not, one may doubt the validity of any evaluation technique), then undue risk taking should manifest itself in numerous low period returns and, thus, a reduced geometric mean (or terminal wealth, which is an equivalent concept in this context). If such is not the case, then the equivalence of historical variability and risk becomes increasingly dubious. The geometric mean also facilitates the use of very short investment periods (because funds value their holdings several times a day, thousands of observations per year could be obtained) and provides a cumulative effect if desired (by simply including each new set of observations without discarding any of the old).

In its simplest form, the capital asset pricing model (CAPM) is the more common name for the SML. Over the years, however, efforts have been made to extend the CAPM to multiple periods, other investment media, foreign markets, and even human wealth. Unfortunately, however, it became increasingly apparent that very little of the cross section of securities' returns was reliably explained by beta.

14.7 Views of Risk

Einstein has been quoted as saying that it was more difficult to understand compound interest than it was to understand special relativity. While this author cannot agree with Einstein's statement, I would certainly agree that the consequences of the manipulation of markets by bankers and politicians are more intricate and frequently more brutal than the operation of organized crime. Firstly there really is no honest way to eliminate risk without reducing profit margins to starvation levels. If one insists on investing in T-bills, then taxes and inflation will leave one a pauper at retirement time.

There are ways honestly to participate in a growing economy (if only politicians will allow it to grow). The economy of the United States rises on a wave of continuing technological progress, whether it be in agriculture, or health care, or the manufacturing of washing machines. If technological progress be sufficiently stressed by taxes and regulation, then the ship of state, including all possible investments (excepting perhaps those in submarines and lifeboats and survivalist provisions), will sink. The United States, with only 5% of the world's population, employs 70% of the world's lawyers. The proportion of lawyers in Congress exceeds even one third. Having a country run by one of its least productive segments might be thought to be other than wise. And, as to the 5% figure, it could be pointed out that the proportion of engineering baccalaureates in the United States hovers around 5%. We do produce vast numbers of sociologists, psychologists, historians, etc., who can perform as social workers or perhaps continue on to acquire law degrees or, as President George W. Bush, go for the MBA. Both President G.W. Bush and President Barack Obama expressed the desire that all American students could go to college. The number of trade high schools, where one might learn a skilled trade has dropped to near zero in this country.

That being said, there has always been a "muddling through" spirit in the United States. This author writes this chapter assuming that things will not be much more chaotic than they have been in the last 75 years.

14.7.1 Diversification

If we have ten stocks, each with the same growth rate and each with the same volatility, dividing our investment among the ten stocks rather than putting all our investment in any one of them is almost a "free lunch" (assuming their returns are not perfectly correlated with each other). Of course, the lunch is not entirely free. Such diversification should save us from losing everything in an Enron, but it might kill our hopes of becoming a Microsoft millionaire (as occurred to many Microsoft clerical personnel who had retirement plans invested heavily in the stock of their employer, which is the other side of the coin from the experience of Enron employees). Diversification of this sort has been used for a long time (in the nineteenth

century many farmers planted corn as well as wheat in the event that hail storms zapped the more profitable wheat).

Unfortunately, a bear market will tend to cause most stocks in the portfolio to drop. Just as an extended drought will zap both corn and wheat, a bear market will hurt stocks generally. (An old politically incorrect adage of Wall Street is “When the paddy wagon comes, good girls are arrested as well as the bad.”) What other variable can we use for “diversification”? The answer is **time**. The historical fact is that investors over longer periods of time, have the advantage of the fact that in roughly 70% of the years, the index of large cap U.S. stocks rises rather than falls. And there is the further encouraging news that in over 40% of the years, the index rises by over 20%. In 30% of the years, the market rises by over 25%. And in 25% of the years, the index has risen by over 30%. Over the roughly 75 year period such records have been kept, the United States has lived through the Great Depression, the Second World War, the Cold War, Korea, Vietnam, assorted massive sociological changes, shifts toward and away from free markets, and assorted epidemics. These can all be viewed as the political/economic/sociological analogs of major “droughts.” It is true that we have yet to experience Martian invasion, attacks by genetically engineered viruses or suitcase nuclear devices, or the costs of mounting the Sixth Crusade. We hope such events do not occur, but events of comparable angst have occurred to other countries of the West over the past 75 years. Poland was occupied by Russia and Germany in September of 1939, and the Russian occupation only ended (sort of) in June of 1989. It is hard to imagine a course of action (other than attempting to take oneself and one’s money out of Poland and moving to, say, the United States) which would have saved an investor in the Warsaw Stock Exchange. And it is hard today to imagine a safe harbor for oneself or one’s property in the event that the United States falls. Past performance is not an infallible guide for looking over the risk profile and we do not claim it to be. Moreover, readers have surely by this point learned that we have not produced any ethical schemes for getting rich quick. We hope, however, to give clues by which an investor can hope to become wealthy at a moderate rate of speed. Let us consider a federally insured certificate of deposit with interest rate r (interest reinvested) with a time horizon of, say, five years. Calling the value of the CD at time t , $X(t)$, we can compute the change in the value of the CD over an increment of time Δt via

$$\Delta(X(t)) = rX(t)\Delta(t). \quad (14.15)$$

We recognize, in the limit as Δt goes to zero, one of the most venerable of simple differential equations:

$$\frac{dX(t)}{X(t)} = rdt \quad (14.16)$$

with solution

$$X(t) = X(0) \exp(rt). \quad (14.17)$$

This is the equation of compound interest of which Einstein is reputed to have said, "There is no magic in special relativity. Compound interest, now that is magic." Actually, there is no magic in compound interest either, but there is a great deal of dishonest flim-flam in the way that it is used

14.8 Stock Progression as Geometric Brownian Motion

14.8.1 Ito's Lemma

Following Hull [6], let us suppose we have a continuously differentiable function of two variables $G(x, t)$. Then, taking a Taylor's expansion through terms of the second order, we have

$$\begin{aligned} \Delta G \approx & \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial t} \Delta t \\ & + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} (\Delta x)^2 + \frac{1}{2} \frac{\partial^2 G}{\partial t^2} (\Delta t)^2 + \frac{\partial^2 G}{\partial x \partial t} \Delta x \Delta t. \end{aligned} \quad (14.18)$$

Next let us consider the *general Ito process*

$$dx = a(x, t)dt + b(x, t)dz \quad (14.19)$$

with discrete version

$$\Delta x = a(x, t)\Delta t + b(x, t)\epsilon\sqrt{\Delta t}, \quad (14.20)$$

where dz denotes a Wiener process, and a and b are deterministic functions of x and t . We note that

$$(\Delta x)^2 = b^2\epsilon^2\Delta t + \text{terms of higher order in } \Delta t. \quad (14.21)$$

Now

$$\text{Var}(\epsilon) = E(\epsilon^2) - [E(\epsilon)]^2 = 1.$$

So, since by assumption $E(\epsilon) = 0$,

$$E(\epsilon^2) = 1.$$

Furthermore, since ϵ is $\mathcal{N}(0, 1)$, after a little algebra, we have that $\text{Var}(\epsilon^2) = 2$, and $\text{Var}(\Delta t \epsilon^2) = 2(\Delta t)^2$. Thus, if Δt is very small, through terms of order $(\Delta t)^2$, we have that it is equal to its expected value, namely,

$$(\Delta x)^2 = b^2\Delta t. \quad (14.22)$$

Substituting (14.20) and (14.22) into (14.18), we have

Lemma 14.1 (Ito)

$$\Delta G = \left(\frac{\partial G}{\partial x} a(x, t) + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right) \Delta t + \frac{\partial G}{\partial x} b \epsilon \sqrt{\Delta t} \quad (14.23)$$

or

$$dG = \left(\frac{\partial G}{\partial x} a(x, t) + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right) dt + \frac{\partial G}{\partial x} b dz. \quad (14.24)$$

14.8.2 A Geometric Brownian Model for Stocks

Let us look at stock growth as noisy compound interest

$$\frac{\Delta S}{S} = \mu \Delta t + \sigma \epsilon \sqrt{\Delta t}. \quad (14.25)$$

Again, σ (the *volatility*) is a measure of the variability of the process as time increases. Here we will formally take ϵ to be a normal variate with mean zero and variance 1. In the limit, as Δt goes to zero, such a process is uniquely defined and is commonly referred to as a geometric Brownian process.

$$d(\ln S) = \mu dt + \sigma dz. \quad (14.26)$$

Alternatively, we have

$$dS = \mu S dt + \sigma S dz. \quad (14.27)$$

Now, in Ito's lemma we define $G = \ln S$. Then we have

$$\frac{\partial G}{\partial S} = \frac{1}{S}; \quad \frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2}; \quad \frac{\partial G}{\partial t} = 0.$$

Thus G follows a Wiener process:

$$dG = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dz. \quad (14.28)$$

This tells us simply that if the price of the stock at present is given by $S(0)$, then the value t units in the future will be given by

$$\begin{aligned} S(t) &= S(0) \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) t + \sigma \epsilon \sqrt{t} \right] \\ &= S(0) \exp \left[\mathcal{N} \left(\left(\mu - \frac{\sigma^2}{2} \right) t, t\sigma^2 \right) \right] \\ &= \exp \left[\mathcal{N} \left(\log(S(0)) + \left(\mu - \frac{\sigma^2}{2} \right) t, t\sigma^2 \right) \right], \end{aligned} \quad (14.29)$$

where $\mathcal{N}(\log(S(0)) + (\mu - \sigma^2/2)t, t\sigma^2)$ is a normal random variable with mean $\log(S(0)) + (\mu - \sigma^2/2)t$ and variance $t\sigma^2$. Thus, $S(t)$ is a normal

variable exponentiated (i.e., it follows the *lognormal* distribution). The expectation of $S(t)$ is given by $S(0) \exp[\mu t]$. In the current context, the assumption of an underlying geometric Brownian process (and hence that $S(t)$ follow a lognormal distribution) is somewhat natural. Let us suppose we consider the prices of a stock at times t_1 , $t_1 + t_2$, and $t_1 + t_2 + t_3$. Then if we assume $S(t_1 + t_2)/S(t_1)$ to be independent of starting time t_1 , and if we assume $S(t_1 + t_2)/S(t_1)$ to be independent of $S(t_1 + t_2 + t_3)/S(t_1 + t_2)$, and if we assume the variance of the stock price is finite for finite time, and if we assume that the price of the stock cannot drop to zero, then, it can be shown that $S(t)$ must follow geometric Brownian motion and have the lognormal distribution indicated.

When a stock price $S(t)$ obeys the law given in the three equivalent forms we say that $S(t)$ is *lognormal* with *growth rate* μ and *volatility* σ . We note that the expected value of $S(t)$ is given by

$$\begin{aligned}
 E(S(t)) &= S(0) \int_{-\infty}^{\infty} \exp((\mu - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}Z) \frac{1}{\sqrt{2\pi}} e^{-Z^2/2} dZ \\
 &= S(0) \exp((\mu - \frac{1}{2}\sigma^2)t) \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{\sigma\sqrt{t}Z} e^{-Z^2/2} dZ \\
 &= S(0) \exp((\mu - \frac{1}{2}\sigma^2)t) \int_{-\infty}^{\infty} \exp[-\frac{1}{2}(Z - \sigma\sqrt{t})^2] dZ e^{\frac{1}{2}\sigma^2 t} \\
 &= S(0) \exp[(\mu - \frac{1}{2}\sigma^2)t + (\frac{1}{2}\sigma^2)t] \\
 &= S(0) e^{\mu t}.
 \end{aligned} \tag{14.30}$$

It is a straightforward matter to show that

$$\begin{aligned}
 \text{Var}(S(t)) &= E[S(t) - E(S(t))]^2 \\
 &= S(0)^2 \exp[2\mu t] [e^{\sigma^2 t} - 1]
 \end{aligned}$$

14.9 Estimating μ and σ

From (14.30), we have, for all t and Δt

$$r(t + \Delta t, t) = \frac{S(t + \Delta t)}{S(t)} = \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) \Delta t + Z \sigma \sqrt{\Delta t} \right]. \tag{14.31}$$

Defining $R(t + \Delta t, t) = \log(r(t + \Delta t, t))$, we have

$$R(t + \Delta t, t) = \left(\mu - \frac{\sigma^2}{2} \right) \Delta t + \epsilon \sigma \sqrt{\Delta t}.$$

Then

$$E[R(t + \Delta t, t)] = \left(\mu - \frac{\sigma^2}{2} \right) \Delta t. \tag{14.32}$$

Suppose we have a stock that stands at 100 at week zero. In 26 subsequent weeks we note the performance of the stock as shown in Table 9.1. Here $\Delta t = 1/52$. Let

$$\bar{R} = \frac{1}{26} \sum_{i=1}^{26} R(i) = .002931.$$

Table 14.1. 26 Weeks of Stock Performance

Week= i	Stock(i)	$r(i) = \text{Stock}(i)/\text{Stock}(i-1)$	$R(i) = \log(r(i))$
1	99.83942	0.99839	-0.001611
2	97.66142	0.97818	-0.02206
3	97.54407	0.99880	-0.00120
4	96.24717	0.98670	-0.01338
5	98.65675	1.02503	0.02473
6	102.30830	1.03701	0.03634
7	103.82212	1.01480	0.01469
8	103.91875	1.00093	0.00093
9	105.11467	1.01151	0.01144
10	104.95000	0.99843	-0.00157
11	105.56152	1.00583	0.00581
12	105.44247	0.99887	-0.00113
13	104.21446	0.98835	-0.01171
14	103.58197	0.99393	-0.00609
15	102.70383	0.99152	-0.00851
16	102.94174	1.00232	0.00231
17	105.32943	1.02320	0.02293
18	105.90627	1.00548	0.00546
19	103.63793	0.97858	-0.02165
20	102.96025	0.99346	-0.00656
21	103.39027	1.00418	0.00417
22	107.18351	1.03669	0.03603
23	106.02782	0.98922	-0.01084
24	106.63995	1.00577	0.00576
25	105.13506	0.98589	-0.01421
26	107.92604	1.02655	0.02620

By the strong law of large numbers, the sample mean \bar{R} converges almost surely to its expectation $(\mu - \sigma^2/2)\Delta t$. Next, we note that

$$[R(t + \Delta t, t) - E(R(t + \Delta t, t))]^2 = \epsilon^2 \sigma^2 \Delta t, \quad (14.33)$$

so

$$\text{Var}[R(t + \Delta t, t)] = E[R(t + \Delta t, t) - \left(\mu - \frac{\sigma^2}{2}\right) \Delta t]^2 = \sigma^2 \Delta t. \quad (14.34)$$

For a large number of weeks, this variance is closely approximated by the sample variance

$$s_R^2 = \frac{1}{26-1} \sum_{i=1}^{26} (R(i) - \bar{R})^2 = .000258.$$

Then $\hat{\sigma}^2 = .000258/\Delta t = .000258 \times 52 = .013416$, giving as our volatility estimate $\hat{\sigma} = .1158$. Finally, our estimate for the growth rate is given by

$$\hat{\mu} = \bar{R} \times 52 + \frac{\hat{\sigma}^2}{2} = .1524 + .0067 = .1591.$$

14.10 The Time Indexed Distribution of Portfolio Value

Let us construct, using simulation, the distribution of 1000 possible outcomes of an investment of \$10,000 in a stock with $\mu = 0.10$ and $\sigma = 0.10$ after ten years. We show the results in Figure 14.6. This is not, of course, a *histogram* in the usual sense of the term. A histogram is a relative count register of the number of historical observations which fall into the intervals of value observed in the past. Here, we have taken parameter values and used them in a model to obtain simulations. So, each simulation gives an observation of simulated value. (Of course, for this simple case, one can obtain the density as a Gaussian probability integral. Very quickly, we will be moving to a situation where such a “closed-form solution” is no longer practical.)

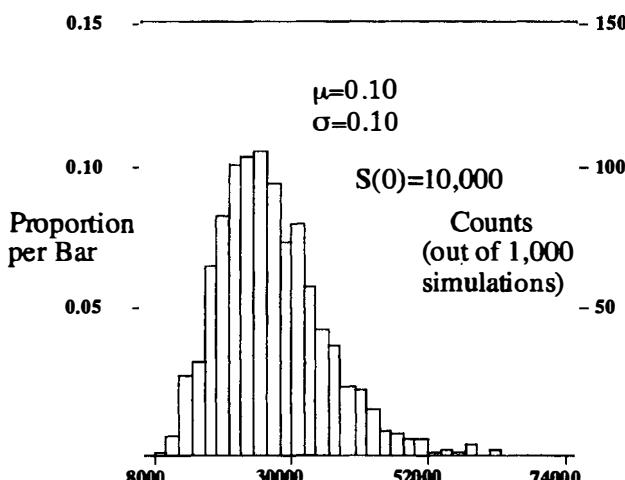


Figure 14.6. Simulated investment values at ten years.

Next, if we can find 20 stocks, each with $\mu = .10$ and $\sigma = .10$, then *assuming they are stochastically independent of each other*, we might take the \$10,000 and invest \$500 in each of the stocks. The distribution of value at ten years (using 200 possible outcomes) is shown in Figure 14.7. The sample means for both the one-stock investment and the diversified 20-stock mutual fund are \$27,088 and \$26,715, respectively. But the standard deviation for the one-stock investment (\$8,343) is roughly $\sqrt{20}$ times that for the mutual fund investment (\$1,875). A portfolio that has such such stochastic independence would be a truly *diversified* one. Generally speaking, one should expect some dependency between the tracks of the stocks.

Now let us recall the general equation for geometric Brownian motion

$$\frac{\Delta S_i}{S_i} = \mu \Delta t + \sigma \epsilon \sqrt{\Delta t}. \quad (14.35)$$

Let us modify (14.35) to allow for a mechanism for dependence:

$$\frac{\Delta S_i}{S_i} = \mu \Delta t + \sigma \epsilon_i \sqrt{\Delta t}. \quad (14.36)$$

We shall take η_0 to be a Gaussian random variable with mean zero and variance 1. Similarly, the 20 η_i will also be independent Gaussian with mean zero and variance 1. Then we shall let

$$\epsilon_i = c(a\eta_0 + (1 - a)\eta_i). \quad (14.37)$$

We wish to select c and a so that a is between zero and 1 and so that $\text{Var}(\epsilon_i) = 1$ and any two ϵ_i and ϵ_j have positive correlation r . After a little algebra, we see that this is achieved when

$$a = \frac{\rho - \sqrt{\rho(1 - \rho)}}{2\rho - 1} \quad (14.38)$$

and

$$c^2 = \frac{1}{a^2 + (1 - a)^2}. \quad (14.39)$$

At the singular value of $\rho = 0.5$, we use $a = 0.5$.

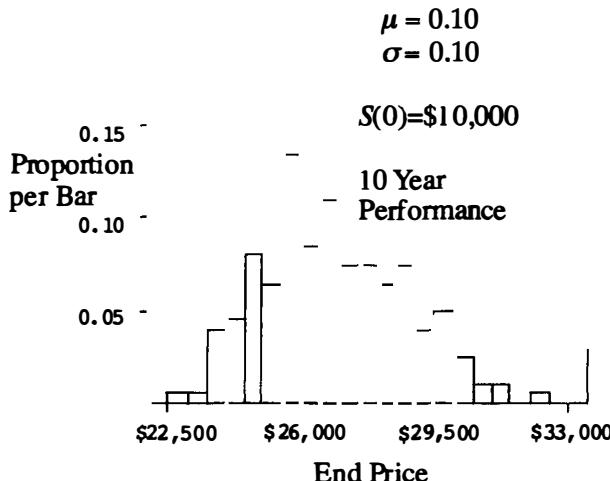


Figure 14.7. A simulation of an idealized mutual fund value at ten years.

Let us examine the situation with an initial stake of \$500 per stock with $\mu = \sigma = .10$ and $\rho = .8$ as shown in Figure 14.8. We employ 500 simulations. We note that the standard deviation of the portfolio has grown to 7,747. This roughly follows the rule that the standard deviation of a portfolio where stocks have the same variance and have correlation ρ , should be $\sqrt{1 + (n - 1)\rho}$ times that of an uncorrelated portfolio.

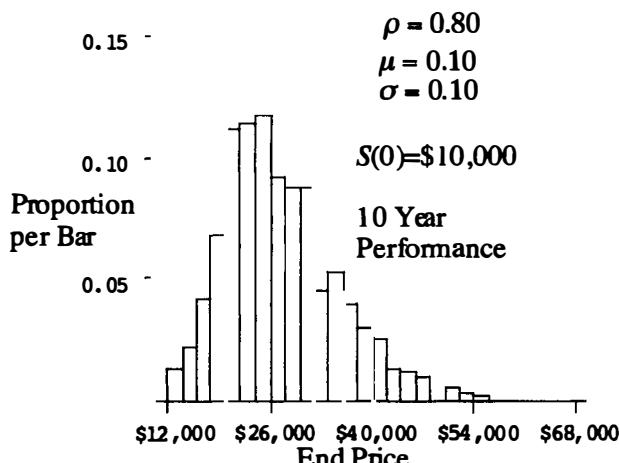


Figure 14.8. Simulation of mutual fund with correlated stock prices.

14.11 Negatively Correlated Portfolios

Is there anything more likely to reduce the variance of a mutual fund portfolio than the assumption that the stocks move in a stochastically independent fashion? We recall that if we have two random variables X_1 and X_2 , each with unit variance and the same unknown mean μ , the variance of the sample mean is given by

$$\text{Var}((X_1 + X_2)/2) = \frac{1}{4}[2 + 2\rho]. \quad (14.40)$$

Here the variance can be reduced to zero if $\rho = -1$. Let us consider a situation where we have two stocks each of value \$5000 at time zero which grow according to

$$\begin{aligned} \frac{\Delta S_1}{S_1} &= \mu\Delta t + \sigma\epsilon\sqrt{\Delta t} \\ \frac{\Delta S_2}{S_2} &= \mu\Delta t - \sigma\epsilon\sqrt{\Delta t}, \end{aligned} \quad (14.41)$$

where ϵ is a Gaussian variate with mean zero and unit variance. Then the resulting portfolio (based on 500 simulations) is exhibited in Figure 14.9.

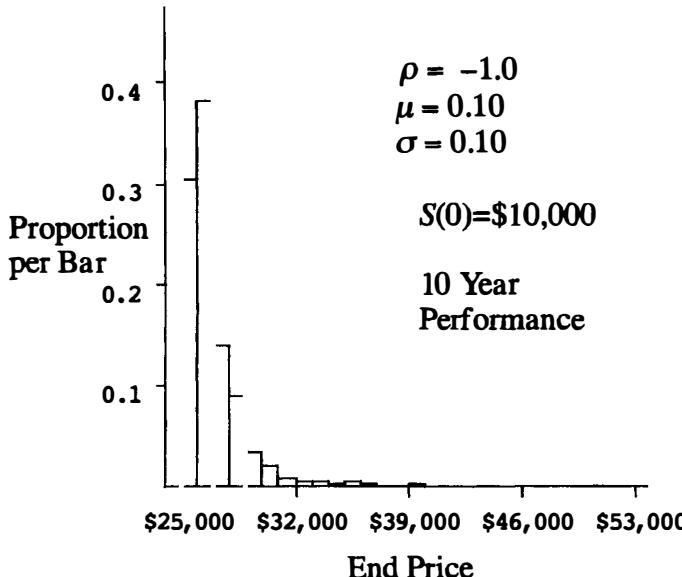


Figure 14.9. Simulation of two-stock portfolio with $\rho = -1$.

We note that the standard deviation of this two-stock portfolio is 1701, even less than that observed for the 20-stock portfolio with the assumption of independence of stocks. Now, the assumption that we can actually find stocks with negative correlation to the tune of -1 is unrealistic. Probably,

we can find two stocks with rather large negative correlation, however. This is easily simulated via

$$\begin{aligned}\frac{\Delta S_1}{S_1} &= \mu \Delta t + (a\epsilon_0 + (1-a)\epsilon_1)c\sqrt{\Delta t}\sigma & (14.42) \\ \frac{\Delta S_2}{S_2} &= \mu \Delta t - (a\epsilon_0 + (1-a)\epsilon_2)c\sqrt{\Delta t}\sigma,\end{aligned}$$

where

$$a = \frac{\rho + \sqrt{-\rho(1-\rho)}}{2\rho + 1},$$

$c = 1/\sqrt{a^2 + (1-a)^2}$ and ϵ_0 , ϵ_1 , and ϵ_2 are normally and independently distributed with mean zero and variance 1. Let us consider, in Figure 14.10, the situation where $\rho = -0.5$.

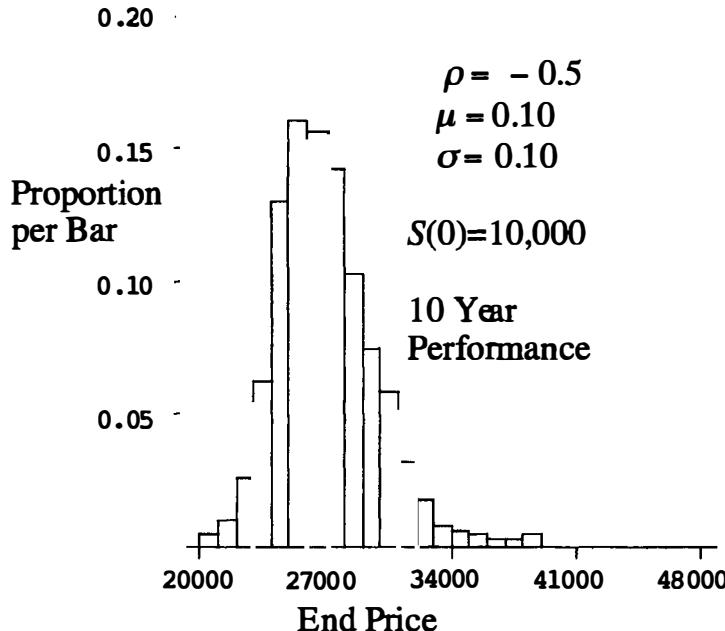


Figure 14.10. Simulation of two-stock portfolio with $\rho = -0.5$.

We note that the standard deviation here has grown to \$2,719. When it comes to utilizing negative correlation as a device for the reduction of the variance of a portfolio, a number of strategies can be considered. We know, for example, that if one wishes to minimize the variance of a sample mean, we can pose the problem as a constrained optimization problem to find the optimal correlation matrix, where we impose the constraint that the covariance matrix be positive definite. Our problem here is rather different, of course.

We could try something simple, namely take our two-stock negatively correlated portfolio and repeat it 10 times, (i.e., see to it that the stocks in each of the ten subportfolio have zero correlation with the stocks in the other portfolios). Here, each of the 20 stocks has an initial investment of \$500. In Figure 14.11, we show the profile of 500 simulations of such a portfolio. The standard deviation of the pooled fund is only \$1,487.

How realistic is it to find uncorrelated subportfolios with stocks in each portfolio negatively correlated? Not very. We note that we can increase the sizes of the subportfolios if we wish, only remembering that we cannot pick an arbitrary correlation matrix—it must be positive definite. If we have a subportfolio of k stocks, then if the stocks are all to be equally negatively correlated, the maximum absolute value of the correlation is given by $1/(k - 1)$.

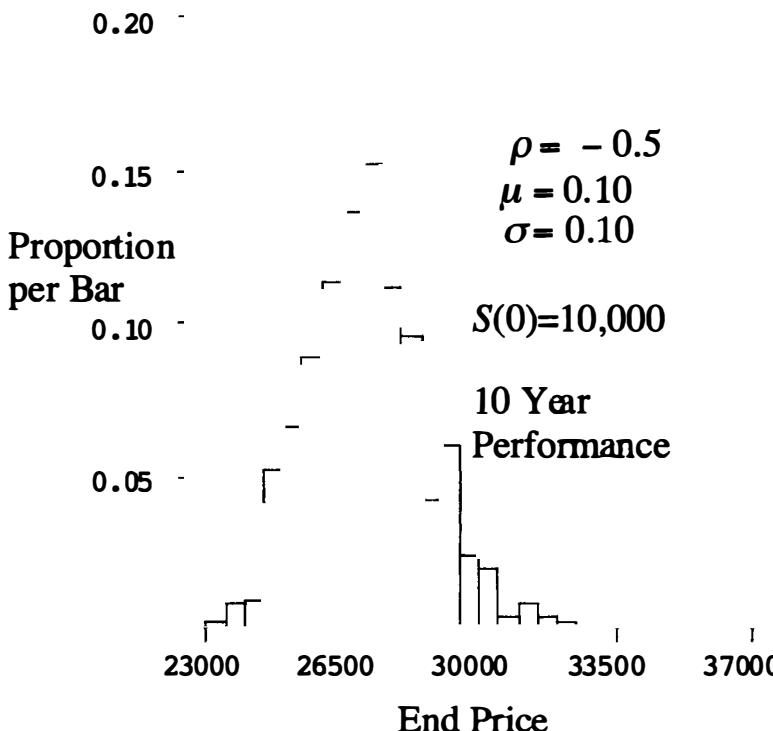


Figure 14.11. Simulation of ten independent two-stock portfolios with $\rho = -.5$.

Let us consider another type of randomness in the stock market. Superimposed over Brownian geometric motion of stocks there are periodic bear market downturns in the market overall. It is unusual for bull markets to exhibit sharp sudden rises. But 10% corrections (i.e., rapid declines) are

quite common, historically averaging a monthly probability of as much as 0.08. Really major downturns, say 20%, happen rather less frequently, say with a monthly probability of 0.015.

14.12 Bear Jumps

In Figure 14.12 we see the binned results of 500 simulations with the jumps modeled as above, $\sigma=.10$, $\rho = 0$, and $\mu = 0.235$. The mean here is \$ 27,080, very similar to that of the situation with independent stocks, with $\mu=.10$ and $\sigma = .1$. However, we note that the standard deviation is a hefty \$11,269.

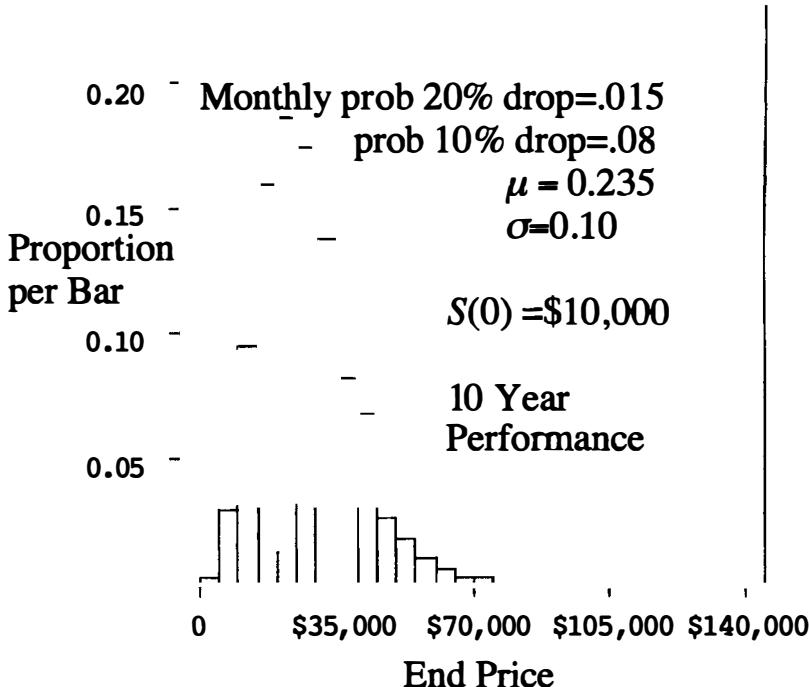


Figure 14.12. Simulation of portfolio of 20 independent stocks with bear jumps.

We note that these general (across the market) downward jumps take away some of the motivation for finding stocks which have local negative correlation in their movements. (For example, had our portfolio had a .8 correlation between the stochastic increments, the standard deviation would only have increased from 11,269 to 13,522.) It is rather easy to advise an investor not to put all his/her investments in Apple Computers. It is also clear that the investor ought not invest in a portfolio consisting of Apple,

Dell, Compaq, Oracle and HP, for these stocks tend to move together. However, when a big bear market hits, our investor is very likely to be damaged even if he/she includes department stores, utilities, drug companies and restaurant chains. Still, a diversified investor, in the bear market which started in the last half of 2000, generally fared better than the investor who only invested in a few high tech stocks. Now, we have arrived at a situation where nearly 25% of the time, our portfolio performs worse than a riskless security invested at a 6% return. If we increase the volatility σ to .5, then nearly 40% of the simulated portfolios do worse than the fixed 6% security.

Let us return to looking at the situation where a \$100 million university endowment, consisting of 20 stocks (\$5 million invested in each stock) with stochastically independent geometric Brownian steps, with $\mu = 0.235$, $\sigma = 0.1$ and with monthly probabilities 0.08 of a 10% drop in all the stocks and a 0.015 probability of a 20% drop in all the stocks. We shall "spend" the portfolio at the rate of 5% per year. Let us see in Figure 14.18 what the situation might be after 10 years. With probability 0.08, after 10 years, the endowment will have shrunk to less than \$50 million. With probability 0.22, it will have shrunk to less than \$75 million. With probability 0.34, it will have shrunk to less than the original \$100 million. Given that universities tend to spend up to their current cash flow, it is easy to see how some of those which tried such a strategy in the 1960s went through very hard times in the 1970s and 1980s.

I have discussed how a portfolio of stocks has, historically, been a superior investment over the long run for retirement purposes. The endowment problem just discussed applies to those who have followed such a policy as they near retirement (especially if the timing is mandatory). If the date falls soon after a bear jump, the options are ugly: Cash out into a fixed annuity and consume much less than one had anticipated or consume from the equity portfolio and hope for a big, quick recovery before the portfolio is literally "eaten up." Retirement flexibility can moderate this risk (and is why a lot of people who hoped to retire in 2000–2002 did not). Failing this, at least a partial shift out of equities as one approaches mandatory retirement is indicated. However, at the present time, when the Federal Reserve has continued a policy for some years of artificially keeping interest rates near zero, it could be argued that the burden of the United States government fighting, without justification, a huge war against the Iraqis while demanding savings and loan associations allow \$600,000 mortgages to be given to persons with incomes below \$50,000, has created a situation in which there is really no safe ground other than precious commodities, and that a conscious decision has been made to make retirees pay the cost of ridiculous government policies.

It is very likely the case that broad sector downward jumps ought to be included as part of a realistic model of the stock market. The geometric Brownian part may well account for the bull part of the market. But

downward movements may very well require a jump component. By simply noting variations on geometric Brownian drift in stock prices, analysts may be missing an essential part of the reality of the stock market, namely large broad market declines which occur very suddenly.

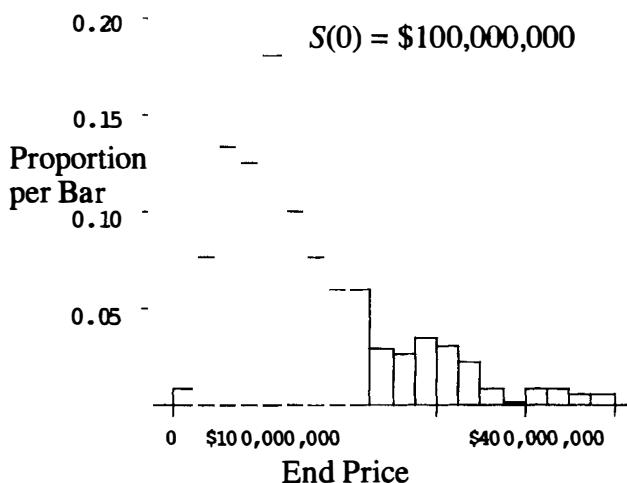


Figure 14.13. Simulation of \$100 million endowment with 5% payout.

14.13 “Everyman’s” MaxMedian Rule for Portfolio Management

If index funds, such as Vanguard’s S&P 500 are popular (and with some justification they are), this is partly due to the fact that over several decades the market cap weighted portfolio of stocks in the S&P 500 of John Bogle (which is slightly different from a total market fund) has small operating fees currently, less than 0.1% compared to hedge fund management rates typically around 30 times that of Vanguard. And, with dividends thrown in, it produces around a 10% return. Many people, perhaps damaged by MBA toting Poobahs in \$5,000 suits living in fancy offices, invest in managed funds. The results have not been promising, overall, although those dealing with people like Peter Lynch and Warren Buffet have done generally well. John Bogle probably did not build his Vanguard funds because of any great faith in fatwahs coming down from the EMH crowd at the University of Chicago. Rather, he was arguing that investors were paying too much for the “wisdom” of the Poobahs. There is little question that John Bogle has benefited greatly the middle class investor community.

That being said, we have shown earlier in this chapter that market cap weighted funds do no better (actually worse) than those selected by random choice. It might, then, be argued that there are nonrandom strategies which

the individual investor could use to his/her advantage. For example if one had invested in the stocks *with equal weight* in the S&P 100 over the last 40 years rather than by weighting according to market cap, she would have experienced a significantly higher annual growth (our backtest revealed as much as a 5% per year difference in favor of the equal weighted portfolio). Moreover, the downside losses in bad years would have been less than with a market cap weighted fund. Actually, there are now several equal weight funds. Moreover, their operating fees are rather modest. It would be nice if we could come up with a strategy which kept only 20 stocks in the portfolio. If one is into managing ones own portfolio, it would appear that Baggett and Thompson [1] did about as well as the equal weight of the S&P 100 using a portfolio size of only 20 stocks. I am harking back to the old morality play of "Everyman" where the poor average citizen moving through life is largely abandoned by friends and advisors except for *Knowledge* who assures him "Everyman, I will accompany you and be your guide."

The science of statistics exists, in large measure, as a vehicle for testing and, when appropriate, modifying models in other sciences. We have a great deal of evidence to the effect that the academic school of finance dominant in the United States, that of the efficient market hypothesis, has not stressed its models sufficiently with market data. This fact has led to market strategies with poor, frequently disastrous results. If one looks at a growth versus volatility (standard deviation) plot for stocks and portfolios of stocks, then the efficient frontier [7–8] of Markowitz gives a convex curve below and to the right of which no investor would wish to venture. Beyond that, Sharpe notes [9–10] that if one plots a point on the left for the zero volatility Treasury Bill, and then goes rightward and upward to a point representing the total market of all publicly traded stocks weighted by market cap, then one has the optimal Capital Market Line (super efficient frontier) above which no stocks or portfolios of stocks should be found. Were these results, which are a basis of efficient market hypothesis practice, correct, then one need look no further for investment in the stock market. One simply picks the volatility level with which one can live and makes an investment on the CML. That Warren Buffett's Berkshire-Hathaway has, for many years, produced returns roughly double those of an S&P 500 Index Fund seems not to shake the faith of believers in the EMH.

It turns out that it is not so difficult to form portfolios with performance above the CML. Wojciechowski and Thompson [18] have shown (looking back 40 years) how randomly selected portfolios selected from the 1,000 largest market cap stocks lie above the CML over half the time (see Figure 14.4). This empirical fact demonstrates that the search for optimal portfolios is by no means a closed issue, as would be supposed by advocates of the EMH. Further, we have observed ([13]–[18]) that the assumptions of EMH option theory are contradicted by market data. Market prices are different from those suggested by Black and Scholes [3] analysis. The dis-

crepancy cannot be avoided by such sharp practices as implied volatility, since implied volatilities for the same stock at the same execution time tend to be different for different strike prices. The building of a vast structure for the trading of derivatives based on specious assumptions has led to a concatenation of market catastrophes, starting with the relatively mild (\$3.3 billion) Long Term Capital Management fiasco in 1998, then proceeding to the failure of ENRON in 2001 (\$62 billion). Looking for villains on whom to blame these disasters, few seemed to consider the flawed EMH theories which had more to do with the problem than the fraudulent off-shore corporations of some of Enron's directors. The collapse of a number of high bandwidth firms seems to have been caused by an unexpected rise in interest rates by the Federal Reserve, attempting to curb the irrational exuberance caused, in large measure, by its earlier managed bail-out of LTCM. The same Federal Reserve then dropped interest rates to such a level that millions of retirees living on fixed incomes experienced significant hardship. Proceeding to the present day, we have observed the subprime mortgage fiasco including the insurance policies written on subprime mortgages, all based on EMH premises. We observe a societal cost which is in the trillions of dollars.

It is not my purpose directly to effect system change. Rather, I present a simple nonproprietary rule which appears significantly to enhance the return to an individual investor. (I have developed and patented a better proprietary algorithm of great computational complexity which will not be dealt with in this book). However the MaxMedian Rule [1], given below, is easy to use and appears to beat the Index, on the average, by up to an annual multiplier of 1.05, an amount which is additionally enhanced by the power of compound interest. Note that that $(1.15/1.00)^{45} = 7.4$, a handy bonus to one who retires after 45 years. A purpose of the MaxMedian Rule was to provide individual investors a tool which they could use and modify without the necessity of massive computing. Others in my classes have developed their own paradigms, such as the MaxMean Rule. In order to use such rules, one need only purchase for a very modest one time fee the Yahoo base *hquotes* program from hquotes.com. (The author owns no portion of the *hquotes* company.)

The MaxMedian Rule

1. For the 500 stocks in the S&P 500 look back at the daily returns $S(j,t)$ for the preceding year the day to day ratios $r(j,t) = S(j,t)/S(j,t-1)$.
2. Sort these for the year's trading days.
3. Discard all r values equal to one. in the 500 medians of the ratios.
4. Invest equally in the 20 stocks with the largest medians.

5. Hold for one year and then liquidate.

In Figure 14.14 we examine the results of putting one present value dollar into play in three different investments: 5% yielding T-bill, S&P 500 Index Fund, MaxMedian Rule. First, we shall do the investment simply without avoiding the intermediate taxing structure. The assumptions are that interest income is taxed at 35%; capital gains and dividends are taxed at 15%; and inflation is 2%. As we see, the T-Bill invested dollar is barely holding its one dollar value over time. The consequences of such an investment strategy are disastrous as a vehicle for retirement. On the other hand, after 40 years, the S&P 500 Index Fund dollar has grown to 11 present value dollars. The MaxMedian Rule dollar has grown to 55 present value dollars. Our investigations indicate that the MaxMedian Rule performs about as well as an equal weighted S&P 100 portfolio, though the latter has somewhat less downside in bad years. Of course, it is difficult for the individual investor to buy into an SP equal weight S&P 100 index fund. The advantage of the equal weight S&P 500 index fund is only 2% greater than that of the market cap weight S&P 500. Even so, when one looks at the compounded advantage over 40 years, it appears to be roughly a factor of two. It is interesting to note that the bogus Ponzi scheme of Bernie Madoff claimed returns which appear to be legally attainable either by the MaxMedian Rule or the equal weight S&P 100 rule. This leads the author to the conclusion that most of the moguls of finance and the Federal Reserve Bank have very limited data analytical skills or even motivation to look at market data.

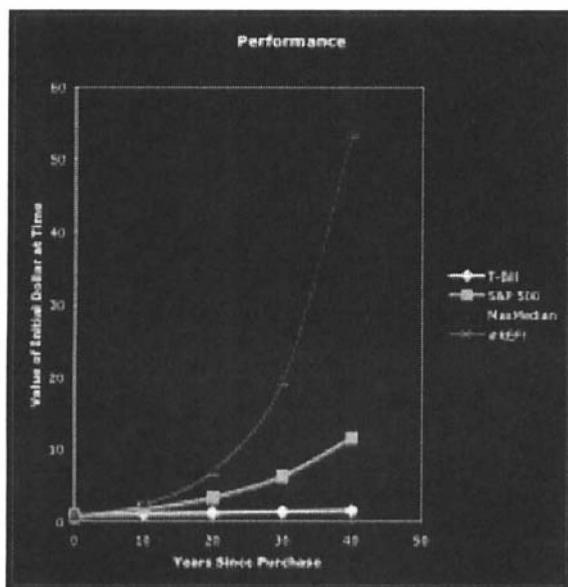


Figure 14.14. A comparison of three investment strategies.

14.13.1 Investing in a 401k

Money invested in a 401-k plan avoids all taxes until the money is withdrawn, at which time it is taxed at current level of tax on ordinary income. In Table 14.2, we demonstrate the results of adding an annual inflation adjusted \$5,000 addition to a 401k for 40 years, using different assumptions of annual inflation rates. All values are in current value dollars.

Table 14.2. 40– Year End Results of Three 401k Strategies.

Inflation	2%	3%	5%	8%
T-bill	447,229	292,238	190,552	110,197
S&P Index	1,228,978	924,158	560,356	254,777
MaxMedian	4,660,901	3,385,738	1,806,669	735,977

We recall that when these dollars are withdrawn, taxes must be paid. So, in computing the annual cost of living, one should figure in the tax burden. Let us suppose the cost of living including taxes for a family of two is \$70,000 beyond Social Security retirement checks. (Of course, the federal government may well decide to eliminate all or part of a family's Social Security payments.) We realize that the 401k portion which has not been withdrawn will continue to grow (though the additions from salary will have ceased upon retirement). Even for the unrealistically low inflation rate of 2% the situation is not encouraging for the investor in T-bills. Both the S&P Index holder and the Max Median holder will be in reasonable shape. For the inflation rate of 5%, the T-bill holder is in real trouble. The situation for the Index Fund holder is also risky. The holder in the MaxMedian rule portfolio appears to be in reasonable shape. Now, by historical standards, 5% inflation is high for the United States. On the other hand, we observe that the decline of the dollar against the Euro during the Bush Administration was as high as 8% per year.

Hence, realistically, 8% could be a possibility to the inflation rate for the future in the United States. In such a case, of the four strategies considered only the return available from the MaxMedian rule leaves the family in reasonable shape. Currently, even the Euro is inflation stressed due to the social welfare excesses of some of the Eurozone members. From a societal standpoint, it is not necessary that an individual investor achieve spectacular returns. What is required is effectiveness, robustness, transparency, and simplicity of use so that the returns will be commensurate with the normal goals of families: education of children, comfortable retirement, etc. Furthermore, it is within the power of the federal government to bring the economy to such a pass where even the prudent cannot make do. The history of Western societies shows that high rates of inflation cannot be sustained without some sort of revolution, such as that which occurred at the end of the Weimar Republic. The lack of awareness of basic intuitive economics among the American people is depressing. Unscrupulous bankers encourage indebtedness on the unwary, taking their profits at the front end

and leaving society as a whole to pick up the bill. Naturally, as a scientist, I would hope that the empirical rules such as the MaxMedian approach of Baggett and Thompson will lead to fundamental insights about the market and the economy more generally. Caveat: The MaxMedian rule is freeware not quality assured or extensively tested. If you use it, remember what you paid for it. The goal of the MaxMedian rule is to enable the individual investor to develop his or her own portfolios without the assistance of generally overpriced and underachieving investment fund managers. Note that I am not badmouthing the equal weighted index funds, which can be an attractive autopilot strategy and do indeed have modest management fees. The investor gets to use all sorts of readily available information in public libraries, e.g., *Investors Business Daily*. Indeed, many private investors will subscribe to *IBD* as well as to other periodicals. Obviously, even if a stock is recommended by the MaxMedian rule (or any rule) and there is valuable knowledge, such as that the company represented by the stock is under significant legal attack for patent infringement, oil spills, etc., exclusion of the stock from the portfolio might be indicated. The bargain brokerage *Fidelity* provides abundant free information for its clients and generally charges less than 8 dollars per trade.

Obviously, one might choose rather a MaxMean rule or a Max 60 Percentile rule or an equal weight Index rule. There are many which might be tested by a forty year backtest. My goal is not to push the MaxMedian rule or the MaxMean rule or the equal weight S&P 100 rule or any rule, but rather allow the intelligent investor to invest without paying vast sums to MBAs in \$5000 suits. If, at the end of the day, the investor chooses to invest in market cap based index funds, that is suboptimal but not ridiculous. What is ridiculous is not to work hard to understand as much as practicable about investment. This chapter is a very good start. It has to be observed that at this time in history, investment in US Treasury Bills or bank cds would appear to be close to suicidal. Both the Federal Reserve and the investment banks are doing the American middle class no good service. 0.2% return on Treasury Bills is akin to theft, and what some of the investment banks do is akin to robbery.

The author has no magic riskless formula for getting rich. (The author created and owns the patent on a highly computer intensive artificial intelligence algorithm for buying stocks, which backtested over a 40 year period at roughly the rate of return of Buffett's Berkshire-Hathaway. We shall not discuss that paradigm in this book.) Rather, I shall offer some opinions about alternatives to things such as buying T-bills. Investing in market cap index funds is certainly suboptimal. However, it is robustness and transparency rather than optimality which should be the goal of the prudent investor. It should be remembered that most investment funds (including, alas TIAA-CREF) do charge the investor a fair amount of his/her basic investment whatever be the results. The EMH is untrue and does

not justify investment in a market cap weighted index fund. However, the fact is that, with the exception of such gurus as Warren Buffett and Peter Lynch, the wisdom of the professional market forecaster seldom justifies the premium of the guru's charge. There are very special momentum based programs (on one of which the author holds a patent), in which the investor might do well. However, if one simply manages one's own account, using MaxMean or MaxMean within an IRA, it would seem to be better than trusting in gurus who have failed again and again. Berkshire-Hathaway has proved to be over the years a vehicle which produces better than 20% return. For any strategy that the investor is considering, backtesting for, say, 40 years, is a very good idea. That is not easy to achieve with equal weight funds, since they have not been around very long. Baggett and Thompson had to go back using raw S&P 100 data to assess the potential of an S&P 100 equal weight fund, since, to my knowledge, no such fund is currently available. If Bernie Madoff (who had the resources to do it), had set up such a fund, he might well have been able to give his investors the 15% return he promised but did not deliver.

Note that the United States government, with its war of choice in the Middle East, its forcing of commercial banks to grant mortgage loans to persons unlikely to be able to repay them, and its willingness to allow commercial banks to engage in speculative derivative sales, is the driving force behind the market collapse of the late Bush Administration and the Obama Administration. Just the war cost part of the current crisis due to what Nobel Laureate Stiglitz has described as something beyond a three trillion dollar war in the Middle East has damaged both Berkshire-Hathaway's and other investment strategies. To survive in the current market situation, one must be agile indeed. Stiglitz keeps upping his estimates of the cost of America's war in the Middle East. Anecdotally, I have seen estimates as high as six trillion dollars. If we realize that the cost of running the entire US Federal government is around three trillion dollars per year, then we can see what a large effect Bush's war of choice has had on our country's aggregate debt. This fact alone would indicate that a future damaging inflation is all but certain. To some extent, investing in the stock market could be viewed as a hedge against inflation. In the author's opinion, it was Bush's war of choice in the Middle East which caused the recent and continuing recession with real employment rates approaching some of the years of the Great Depression. In the next section, we will examine another cause of denigration and instability in the economy, the use of derivatives.

14.14 Derivatives

In the period frequently referred to as "The Roaring Twenties," there were investment counselors who advised their clients to buy stocks "on margin." Margin investing is most easily explained by an example. Suppose an in-

vestor has \$10,000 to invest in a stock. The current price of the stock is \$100. Thus, the investor could simply buy 100 shares of stock. If the stock's value climbed to \$110, then the investor would have made \$1000. His account in the stock would be worth \$11,000. But if the stock declined in value to \$90, then the investor would have lost \$1000. His account would be still be worth \$9000.

The stock broker advises the investor to use a margin loan with a factor of ten times. That would enable the investor to buy 1000 shares of stock. If the stock price rose to \$110, then the investor would have made a tidy profit of \$10,000. This was the good news the investment counselors told potential investors.

Less common was the warning that if the stock declined in value to \$90, then the margin loan would capture all the shares of stock purchased by the investor. His losses would have brought his account to zero value. Notice, also, that as the investment broker started selling recouped shares, this puts downward pressure on the value of the stock.

On October 28, 1929 there was a 13% fall in the market followed by a 12% fall on October 29. On December 10, there was a run on the Bronx branch of the (privately owned) Bank of the United States. Thus began America's Great Depression. Many things were tried to ameliorate the situation with various stimuli, but the Great Depression was still continuing when America entered the Second World War in 1941. To ameliorate the effect of bank failures on depositors, the Federal Deposit Insurance Corporation was founded. It has had a positive effect, but the amount of funds available to it is really rather small. President Roosevelt changed investment laws to reduce margin loans from ten fold to two fold. At any rate, it would appear that prudent steps were taken by President Roosevelt, which might have eliminated the start of the Great Depression in the first place.

14.14.1 Black-Scholes and the Search for Risk Neutrality

It is a fair statement that those traders who profited greatly from selling margin portfolios to their customers were harmful to the nation. It is also true that President Roosevelt did implement some strategies which minimized this practice. Unfortunately, the Nobel winning Black-Scholes equation has been, in large measure, responsible for the continuing recession which started with Bush's War and continues with Obama's War.

A *call option* is the right (but not the obligation) to buy a specified asset S at a *strike price* X T days in the future. If the seller of the option owns the stock (or commodity) connected with the option, then the option is said to be *covered*. If the seller of the option does not own the stock, then the option is *uncovered*.

Before deriving the Black-Scholes pricing formula for the price of a *Eu-*

ropes an *call option* (defined below), let us look at two older pricing formulas C_A and C_B giving the arguments for each. What should be the fair value for an option to purchase a stock at an exercise price X starting with today's stock price $S(0)$ and an expiration time of T ? If the rate of growth of the stock is μ , and the volatility is σ , then, assuming the lognormal distribution for $S(t)$, the stock value at the time T should be

$$S(0) \exp \left(\mathcal{N} \left(\left(\mu - \frac{\sigma^2}{2} \right) T, T\sigma^2 \right) \right),$$

where $\mathcal{N}(a, b)$ is a Gaussian random variable with expectation a and variance b . If we borrow money at a fixed riskless rate r to purchase the option, then the value of the option could be argued² to be equal to C_A in

$$\text{Method A } C_A = \exp(-rT) E[\max(0, S(T) - X)],$$

where E denotes expectation.

On the other hand, it could also be argued that the person buying the option out of his assets is incurring an opportunity cost by using money to buy the option which might as well have been used for purchasing the stock so that the value of the option should be given by

$$\text{Method B } C_B = \exp(-\mu T) E[\max(0, S(T) - X)].$$

After a little calculus [11], it can be shown that we have for **Method A**,

$$C_A = e^{-rT} \{ e^{\mu T} S(0) \Phi \left(\frac{\log(S(0)/X) + [\mu + (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) - X \Phi \left(\frac{\log(S(0)/X) + [\mu - (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) \}.$$

For **Method B**, we have

$$C_B = e^{-\mu T} \{ e^{\mu T} S(0) \Phi \left(\frac{\log(S(0)/X) + [\mu + (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) - X \Phi \left(\frac{\log(S(0)/X) + [\mu - (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) \}.$$

The formula for the pricing of an option which gives the purchaser the right to buy a stock for X dollars at future time T , as formulated above, in the case of C_A depends on knowledge of three parameters, μ , σ , and r . Ostensibly, r should be easy to determine. It is frequently argued to be

²The use of the expectation criterion as the measure of value is questionable. We recall, for example, the St. Petersburg Paradox discussed in Chapter 13 shows how expectation of gain can give ridiculous results. The question of appropriate criteria is extremely important, and we have addressed it extensively throughout this book. For the moment, we stick with expectation, for that is, unfortunately, the standard view.

simply the riskless interest rate (i.e., that of a Treasury bill). A bit later, we will question whether a universal r value is reasonable. Knowledge of μ and σ would appear to be more uncertain. They can be estimated quite handily from past stock records, but we are not talking about the μ and σ values of the past. Rather we want to use those which are appropriate from now until the time T (when the option is exercised or is allowed to expire).

14.14.2 A Game Independent of the Odds

Let us consider the game indicated in Figure 14.15

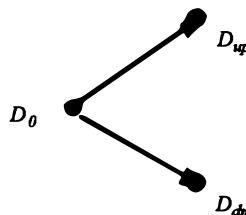


Figure 14.15. Game unlinked to probabilities.

There are two prizes D_{up} and D_{dn} . What is the fair price D_0 to pay the casino to play this game? As we have seen, if the prizes are reasonably small relative to the wealth of the player, the fair value might be taken to be the expected value:

$$D_0 = p_{up}D_{up} + p_{dn}D_{dn}. \quad (14.43)$$

We can be sure, in the real world, that the casino operator will charge a bit more than D_0 , say $D_1 > D_0$ to play this game, for a zero rate of return on his investment (casino, employees, utilities, security, etc.) would soon put him out of business. It is not reasonable, one might argue, for the player to become involved with the wager, since in paying D_1 to play, he will, over the long run, lose money. Yet casinos do a thriving business. In general, the casino organization will be happy, in the case of many of its games, to operate on a reasonably small margin $(D_1 - D_0)/D_1$ of profit, for, over time, it is almost certain to return, on the game, a profit margin per wager equal to the profit margin times the total bets wagered on the game.

In casino games of chance, we know precisely the probabilities associated with each game. But what if one is engaging in a game where one does not know the probabilities associated with winning and losing? Suppose one is taking bets, for example, on a horse race. There are horse racing experts who will set preliminary odds of the sort, Purple Martian has a ten percent chance of winning the race. So, preliminary odds might be that Purple Martian, for a one dollar ticket, will pay ten dollars if it wins the race. But as the bets come in, the track may note that too many people are betting on Purple Martian at the bet as stated. That means, if Purple

Martian should win, the track would actually lose money. So, a woman who bet on Purple Martian would know that the ticket she bought might pay less than the rate stated at the time of the purchase. The track has the right to readjust payoffs right up until the time of the start of the race to make sure it has locked in a profit, irrespective of which horse actually wins the race. Similarly, if the track finds that too few people are betting on Purple Martian, it will typically improve the payoff. The whole idea of the track, then, is to readjust payoffs for all the horses, continuously in time, to guarantee a profit for the track. Now, in the world of bookmakers, competition from several “agencies” will tend to drive the payoffs from all to be similar, for why buy a ticket which pays \$10 if another bookmaker is paying \$15? If communications are as rapid as they have been for some decades, there will be, essentially, a national market for betting on Purple Martian to win. The only thing which can restrict (partially) the creation of a nearly efficient market in the victory of Purple Martian is interference by some governmental or quasi-governmental agency (e.g., organized crime) to restrict trade in Purple Martian tickets. Notice, then, that the world of horse racing bookmakers provides us with an example of a market in outcomes in which the probabilities are not very well known, and that the bookmakers actually set the rates based, not to their retrospective guesses as to the probabilities, but rather on volumes of purchases of tickets for specific pay-off rates.

Notice that the payoff rates are computed by the bookmakers with little reliance on probability estimates. From the standpoint of an individual buying Purple Martian tickets, however, the decision is generally based on the probability of a Purple Martian win as intuited by the individual buying tickets. Here, the bookmaker has a very different strategy from that of the woman making the wager. The bookmaker is attempting to hedge his bets so that he makes a return based, essentially, on commissions.³

Earlier in this chapter we have seen how, given assumptions about μ , r , and σ , to obtain the price of a *European call option* under those assumptions. Moreover, it is clear how one might argue that if we want to determine a “fair” price to promise to deliver a stock T units in the future, when that price is to be paid on the day of delivery, it can be determined by buying the stock at today’s known price $S(0)$ and then borrowing the money for that purchase at the going interest rate r . That would mean that the *futures price at time T* is $S(0) \exp(rT)$. That means that, assuming the broker could borrow the money $S(0) \exp(rT)$ for rate r , then he could have a perfect hedge and make a profit exactly equal to his commission for the transaction. He would be in somewhat the same position of a bookie at the

³The aforementioned description of horse race bookmaking follows closely the way investment bankers price new stock issues (IPOs). The initial price increases or decreases depending on the “book” (even the terminology is similar). Substantial initial interest (buy orders) will raise the IPO price. If there is weak interest, the IPO price will be set lower. The size of the issue may also be adjusted depending on interest in the deal.

races.

From the standpoint of the buyer, however, he or she probably has reason to believe that at time T the stock will be worth more than $S(0) \exp(rT)$. (There are other reasons, of course, for the buyer to purchase the future instead of simply buying the stock today. For example, it might well be that other investments will bring into the futures purchaser's hands a sum of money at time T at least as great as $S(0) \exp(rT)$, a sum which the purchaser does not have in hand today. This rather obvious hedging price whereby the vendor can sell a futures contract goes back so far in history that it can probably be called a "folk theorem." In the case of a European call option, one might have inferred that the way to price a riskless (to the vendor) option would be simply to use the formulae associated with Method B, except replacing the growth rate μ with the bond interest rate r . This method turns out to be correct, given the assumptions.⁴ However, it was not proved to be the case until 1973 [3] at some profit to the provers. In 1998, it earned the two surviving members of the research project (Scholes and Merton) the Nobel Prize in Economics.

In their provocative book *Financial Calculus*, Baxter and Rennie [2, p. 7] state that for the purpose of pricing derivatives, "seductive though the strong law is, it is also completely useless." They show how serious they are about this view when they consider a wager which is based on the progress of a stock [2, p. 15]. At present, the stock is worth \$1. In the next tick of time, it will either move to \$2.00 or to \$0.50. A wager is offered which will pay \$1.00 if the stock goes up and \$0.00 if the stock goes down. The authors form a portfolio consisting of $2/3$ of a unit of stock and a borrowing of $1/3$ of a \$1.00 riskless bond. The cost of this portfolio is \$0.33 at time zero. After the tick, it will either be worth $2/3 \times \$2.00 - 1/3 \times \$1.00 = \$1.00$ or $2/3 \times \$0.50 - 1/3 \times \$1.00 = \$0.00$. From this, they infer that "the portfolio's initial value of \$0.33 is also the bet's initial value."⁵

⁴This "simple" futures approach ignores the risk of "nonexercise" to the seller of the option. At any given time, an investor who believes that a stock is going to make serious upward progress over the next six months may purchase an option to buy that stock six months in the future at a price set today. Such an option is called a *European call option*. The price at which he may buy the stock is called the *exercise price* or *strike price* and the time at which the option can be exercised is called the *expiration date*. The option need not be exercised and will not be unless the value of the stock at the expiration date is at least as great as the *exercise price*.

⁵Some might argue as follows: In fact, several things are occurring in the example which are often confused and combined. Observe that there are two immediate future states, up and down. The portfolio exactly replicates the payoff of the bet in each state and involves a net investment of \$0.33. Hence, unless the bet also sold for \$0.33, an arbitrage profit could be earned by buying the cheaper and selling the dearer of the two. In that no risk would be involved in such a transaction, it would be invariant to (and devoid of information about) risk preferences. Even if the market for the stock were in equilibrium (which would imply \$0.33 was an equilibrium price for the bet), including general agreement about the state outcomes, probabilities and risk preferences create a jointness which cannot be untangled from the information provided. For example, if the market is employing risk neutral pricing, then $\text{prob(up)}=1/3$, $\text{prob(down)}=2/3$, the

14.14.3 The Black–Scholes Hedge Using Binomial Trees

Suppose, however, we were to ask another question (“another” because options are not precisely analogized by payoffs in horse races). Namely, can we come up with a way in which a brokerage firm could match the value of a call option by continually readjusting a portfolio consisting of a mixture of shares of the stock on which the option is based and Treasury bills, in such a way that there would be no risk to the brokerage firm? If we neglect transaction costs to the firm, the answer is “Yes”. The brokerage firm, like a bookie at a race track, makes its money on service fees built into the cost of the commodity being sold. The brokerage firm does not care whether the stock goes up or down, provided that it can determine a “risk neutral” price. Given that the goal has changed from one of forecasting to one of accounting, we will not be surprised if the stochastic effects are modest in the solution of the problem of achieving “risk neutrality.”

Let us return to the situation in Figure 14.15 but with the addition of a stock (with outcomes in the next increment of time s_{up} and s_{dn}) to which the prizes D_{up} and D_{dn} are linked. We will refer now to the prizes D_{up} and D_{dn} as *derivatives*.

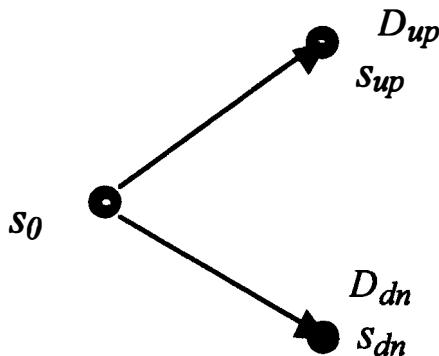


Figure 14.16. Simple binomial tree with options.

A woman approaches the dealer and states she would like to play a game

expected value of the stock, portfolio and bet are their prices and the expected return on each upon the revelation is 0. Suppose, however, the market is engaging in risk averse pricing, such that the probabilities are 50-50. The expected value of the stock is \$1.25 and of the portfolio and bet is \$.50. The instantaneous expected return on the first is 25% and on the latter two is 50%. Further, there is no way to earn an arbitrage profit between the portfolio and the bet. In particular, there is no way to arbitrage this result back to risk neutral pricing. If the only concern is the derivation of the hedged, no arbitrage price of the bet, a shortcut may be employed. In that preferences do not enter this computation, the simplest approach is to assume the stock price is risk-neutral, extract the implied probabilities (i.e., 1/3, 2/3), and price the bet (i.e., $1/3 \times \$1 + 2/3 \times 0 = \0.33). By using this approach, it is not even necessary to derive the hedge portfolio to price the bet. As discussed above, this procedure in no way implies that the price of the stock or the bet is *actually* risk neutral.

in which she receives D_{up} if the stock goes up in the next step and D_{dn} if it goes down. How much should the dealer charge her to play this game?

The dealer decides to emulate the derivative by forming a portfolio of the u units of the stock on which the derivative is based plus a position of v in a bond paying interest rate r . Using Figure 14.16, let us determine u and v .

$$\begin{aligned} us_{up} + ve^{r\Delta(t)} &= D_{up} \text{ when the stock goes up,} \\ vs_{dn} + ve^{r\Delta(t)} &= D_{dn} \text{ when the stock goes down.} \end{aligned} \quad (14.44)$$

Solving for u and v , we have

$$\begin{aligned} u &= \frac{D_{up} - D_{dn}}{s_{up} - s_{dn}}, \\ v &= e^{-r\Delta(t)}(D_{up} - us_{up}). \end{aligned}$$

At time zero, the value of the portfolio (and therefore of the emulated derivative) is given by

$$\begin{aligned} D_0 &= s_u + v \\ &= \exp(-r\Delta(t)) \left(\frac{s_0 e^{r\Delta(t)} - s_{dn}}{s_{up} - s_{dn}} D_{up} + \left(1 - \frac{s_0 e^{r\Delta(t)} - s_{dn}}{s_{up} - s_{dn}}\right) D_{dn} \right) \\ &= qD_{up} + (1 - q)D_{dn} \end{aligned} \quad (14.45)$$

where

$$q = \frac{s_0 e^{\Delta(t)} - s_{dn}}{s_{up} - s_{dn}}. \quad (14.46)$$

(Verify that $0 \leq q \leq 1$.)

1. Now, at time zero, the lady who wishes to play the game gives the dealer D_0 in cash.
2. The dealer buys u shares of the stock at a cost of us_0 .
3. So, at time zero, the dealer has u shares of stock, and cash equal to $D_0 - us_0 = v$.
4. If the stock moves up at time $\Delta(t)$, the dealer sells his u shares at price s_u . He then has cash equal to

$$us_{up} + ve^{r\Delta(t)} = \frac{D_{up} - D_{dn}}{s_{up} - s_{dn}} s_{up} + D_{up} - \frac{D_{up} - D_{dn}}{s_{up} - s_{dn}} s_{up} = D_{up}$$

5. If the stock moves down at time $\Delta(t)$, the dealer sells his u shares at price s_u . He then has cash equal to

$$us_{dn} + v = \frac{D_{up} - D_{dn}}{s_{up} - s_{dn}} s_{dn} + D_{up} - \frac{D_{up} - D_{dn}}{s_{up} - s_{dn}} s_{up} = D_{dn}$$

So, then, the dealer has achieved a strategy of buying a portfolio which gives him zero profit and zero loss whether the stock moves up or down. By imposing a discipline of buying the right amount of stocks and bonds at each step, he has achieved, apparently, a situation where he controls the "state of nature." We note that

$$q = \frac{s_0 e^{r\Delta(t)} - s_{dn}}{s_{up} - s_{dn}}$$

has the formal properties of a probability. That is, q is between zero and one. We have seen that the value of the dealer's portfolio at time zero is given by

$$D_0 = e^{-r\Delta(t)}[qD_{up} + (1 - q)D_{dn}].$$

For a price received from the lady D_0 , he agrees to pay the lady D_{up} at time $\Delta(t)$ if the stock moves to s_{up} and D_{dn} if it moves down to s_{dn} . At time $\Delta(t)$, the dealer has a portfolio of value D_{up} if the stock goes up and D_{dn} if it goes down. The dealer has, it would appear, created his own reality. Maintaining his strategy of updating the portfolio at each tick of time, he will have achieved the ability to have assets, absent his obligations, of zero. This q is, accordingly, frequently referred to as the *martingale* measure: it has "expectation" zero. Like a bookie, the dealer has managed to eliminate his risk and simply live on commissions.

Let us consider a modified version of equation (14.16), namely,

$$\begin{aligned} S(t + \Delta(t)) &= S(t) \exp[(\mu - \frac{1}{2}\sigma^2)\Delta(t)] \exp[\sigma\sqrt{\Delta(t)}] \text{ with prob } q \\ &= S(t) \exp[(\mu - \frac{1}{2}\sigma^2)\Delta(t)] \exp[-\sigma\sqrt{\Delta(t)}] \text{ with prob } 1 - q. \end{aligned} \tag{14.47}$$

Note that

$$\begin{aligned} q &= \frac{s_0 [e^{r\Delta(t)} - e^{\mu\Delta(t) - \sigma\sqrt{\Delta(t)}}]}{s_0 e^{\mu\Delta(t)} [e^{\sigma\sqrt{\Delta(t)}} - e^{-\sigma\sqrt{\Delta(t)}}]} \\ &= \frac{e^{\Delta(t)(r - \mu)} - e^{-\sigma\sqrt{\Delta(t)}}}{e^{\sigma\sqrt{\Delta(t)}} - e^{-\sigma\sqrt{\Delta(t)}}} \\ &= \frac{1 + \Delta(t)(r - \mu) + \Delta(t)^2(r - \mu)^2/2 - 1 + \sigma\sqrt{\Delta(t)} - \sigma^2\Delta(t)/2}{2\sigma\sqrt{\Delta(t)}} \\ &= 1/2 \left(1 - \sqrt{\Delta(t)} \frac{[\mu + \sigma^2/2 - r]}{\sigma} \right). \end{aligned}$$

Again, we look at

$$S(t) = S(0) \exp[(\mu - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}\frac{2X_n - n}{\sqrt{n}})], \quad (14.48)$$

But, with the martingale measure, we have

$$\begin{aligned} E_Q(2X_n - n) &= 2nq - n \\ &= n \left(1 - \sqrt{\Delta(t)} \frac{[\mu + \sigma^2/2 - r]}{\sigma} \right) - n \\ &= -n\sqrt{t/n} \frac{[\mu + \sigma^2/2 - r]}{\sigma}. \end{aligned} \quad (14.49)$$

$$\begin{aligned} \text{Var}(2X_n) &= 4nq(1-q) \quad (14.50) \\ &= n \left(1 - \sqrt{t/n} [\mu + \sigma^2/2 - r] / \sigma \right) \left(1 + \sqrt{t/n} [\mu + \sigma^2/2 - r] / \sigma \right) \\ &= n (1 - (t/n)((\mu + \sigma^2/2 - r)/\sigma)^2) \\ &\rightarrow n + \text{a constant as } n \text{ becomes large, i.e., as } \Delta(t) \rightarrow 0 \end{aligned}$$

Therefore,

$$\frac{2X_n - n}{\sqrt{n}} \rightarrow \mathcal{N} \left(-\sqrt{t}[\mu + \sigma^2/2 - r]/\sigma, 1 \right). \quad (14.51)$$

Let

$$Z = \frac{2X_n - n}{\sqrt{n}} + \frac{\sqrt{t}[\mu + \sigma^2/2 - r]}{\sigma}. \quad (14.52)$$

Z is normal with mean zero and unit variance. Then

$$\begin{aligned} S(t) &= S(0) \exp \left[(\mu - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}\frac{2X_n - n}{\sqrt{n}} \right] \\ &\rightarrow S(0) \exp \left[(\mu - \frac{1}{2}\sigma^2)t - \sigma\sqrt{t}\frac{\sqrt{t}[\mu + \sigma^2/2 - r]}{\sigma} + \sigma\sqrt{t}Z \right]. \end{aligned} \quad (14.53)$$

Then, under the risk-neutral Q measure,

$$\begin{aligned} E_Q(S(t)) &= S(0) \exp[(r - \frac{1}{2}\sigma^2)t] E_Q \left(\exp[\sigma\sqrt{t}Z] \right) \\ &= S(0)e^{rt}. \end{aligned} \quad (14.54)$$

Not surprisingly, considering the way the portfolio is constantly rebalanced, its return is the same as that of a riskless bond paying the rate r . In other words, although under the probability model, $E(S(t)) = S(0) \exp(\mu t)$, the risk neutral hedging model buys and sells the stock in such a way that the expected value of the stock under the hedging discipline is only $E(S(t)) = S(0) \exp(rt)$ regardless of the stock's growth rate μ . In essence, we have taken a risky stock and split it into riskless debt (which the dealer holds) and a risky call (which is sold to the lady). In a perfect world, this might work, but it does not in the real world. Here are some reasons why:

- Transaction costs are not really free. The closer the hedge gets to being riskless, the more frequently one must rebalance (and this results in material transaction costs).
- The realistic value of r will be significantly higher than that of a T–bill.
- Historical records show that the Black–Scholes formula [3], which we develop below, generally does not give the actual market price of a call option. To correct this imperfection in nature, it is customary for some traders to plug in whatever value is necessary for σ to give the market price for the option. We may recall in a chemistry or physics lab when we did not get the answer demanded by the science, there was some temptation to plug in whatever would conform to the established physical model. Such a procedure was called “dry labing” and generally regarded as cheating. Amongst believers in EMH, such a plug-in approach to a σ so obtained has a much more respectable name, *implied volatility*.
- Stock prices may jump (with substantial discontinuities), and this may defeat the hedging strategy. Stock price evolution is not just a smooth function of time.
- Investors will vary a great deal on their expectations as to the future price of the stock. Even in the aggregate, the investment bankers believe the investors to be more leming-like than they are.
- In the case of horse betting, there is an arbitrary mechanism which sets payoffs at the instant the race starts. The bookmaker is allowed to set payoffs with his profit margin locked into the payoff. Suppose that the bet is placed a week before the instant the race starts. What costs must the bookmaker incur in the intervening time period to rebalance the payoffs for the bets he has covered? The answer is “zero.” And this is the reason the analogy between bookmaking and selling options is flawed.

Perhaps the Black–Scholes Theorem can be described as a proof of a result wished by many to be true. Perhaps the best introduction to the result comes from asking the old question of how much an investor should pay at time T for a stock which today has price $S(0)$. The answer is deemed (by many) to be obvious. The seller of the stock future buys the stock today at a price $S(0)$ using money he borrows at interest rate r . If the agreed upon price is $S(0) \exp(rT)$, then when the buyer pays it at time T , the seller can pay back the loan he entered upon to buy the stock. So, the story goes, it really makes no difference what buyer or seller believes the growth rate of the stock is (that has already been incorporated magically into the current price of the stock). The broker of the future naturally will

add a commission to the cost. The only market aspect of the deal will be the possible competition between brokers to lower their commissions.

Let us turn to the “risk-free” purchase of a European call option. Recall that here we pay at time zero for the right to buy the stock at time T for strike price X .

$$\begin{aligned} C_{BS} &= e^{-rT} E[\text{Max}(S(T) - X, 0)] \\ &= e^{-rt} \frac{1}{\sqrt{2\pi\sigma^2 t}} \int_{\ln(\frac{X}{S(0)})}^{\infty} (S(0)e^z - X) \exp\left[-\frac{1}{2\sigma^2 t}(z - (r - \frac{\sigma^2}{2})t)^2\right] dz \end{aligned} \quad (14.55)$$

where we note that the formula is the same as that for Method B except that the growth rate is r rather than μ .

The risk neutral determination value for a call option was first given by Black and Scholes [3]. From Smith’s Lemma (see Problem 2 in this chapter), taking $\psi = 1$, $\lambda = \gamma = e^{-rT}$, and setting $\mu = r$,

$$\begin{aligned} C_{BS} &= e^{-rT} \{ e^{rT} S(0) \Phi\left(\frac{\log(S(0)/X) + [r + (\sigma^2/2)]T}{\sigma\sqrt{T}}\right) \\ &\quad - X \Phi\left(\frac{\log(S(0)/X) + [r - (\sigma^2/2)]T}{\sigma\sqrt{T}}\right) \}. \end{aligned} \quad (14.56)$$

The Black-Scholes model has had enormous impact on the trading of options. Consequently, it has itself changed the mechanism of the market. However, if one looks at the actual market price of an option at a given time, it is seldom the case that it is the same or even close to the Black-Scholes value if one uses historical measures of volatility. But since the true believers know the Black-Scholes value is correct, they take the actual market price of the option and calculate backwards to determine the *implied volatility*. That this *implied volatility* will be different for different time horizons is taken care of by noting that the appropriate average volatility will naturally be different for longer and shorter time epochs. That the *implied volatility* can also be different for the same stock using the same time horizon but different strike prices is the kind of rude remark that is best left unsaid in polite financial circles.

14.15 The Black-Scholes Derivation Using Differential Equations

We recall the brownian model for stock progression:

$$dS = \mu S dt + \sigma S dz. \quad (14.57)$$

Let f be a derivative security (i.e., one that is contingent on S). Then, from Ito’s lemma, we have:

$$df = \left(\frac{\partial f}{\partial S} \mu S + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right) dt + \frac{\partial f}{\partial S} \sigma S dz. \quad (14.58)$$

Multiplying (14.57) by $\partial f / \partial S$, and isolating $\partial f / \partial S \sigma S dz$ on the left side in both (14.57) and (14.58), we have:

$$\begin{aligned}\frac{\partial f}{\partial S} \sigma S dz &= \frac{\partial f}{\partial S} dS - \frac{\partial f}{\partial S} \mu S dt \\ \frac{\partial f}{\partial S} \sigma S dz &= df - \left(\frac{\partial f}{\partial S} \mu S + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right) dt.\end{aligned}$$

Setting the two right-hand sides equal (stochastic though they be), we have:

$$df - \frac{\partial f}{\partial S} dS = \left(\frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right) dt. \quad (14.59)$$

Let us consider a portfolio which consists of one unit of the derivative security and $-\partial f / \partial S$ units of the stock. The instantaneous value of the portfolio is then

$$\mathcal{P} = f - \frac{\partial f}{\partial S} S. \quad (14.60)$$

Over a short interval of time, the change in the value of the portfolio is given by

$$d\mathcal{P} = df - \frac{\partial f}{\partial S} dS = \left(\frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right) dt. \quad (14.61)$$

Now since (14.61) has no dz term, the portfolio is riskless during the time interval dt . We note that the portfolio consists both in buying an option and selling the stock. Since, over an infinitesimal time interval, the Black–Scholes portfolio is a riskless hedge, it could be argued that the portfolio should pay at the rate r of a risk free security, such as a Treasury short term bill. That means that

$$d\mathcal{P} = r\mathcal{P} dt = r \left(f - \frac{\partial f}{\partial S} S \right) dt = \left(\frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right) dt. \quad (14.62)$$

Finally, that gives us the Black–Scholes differential equation

$$rf = \left(rS \frac{\partial f}{\partial S} + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} (\sigma S)^2 \right). \quad (14.63)$$

It is rather amazing that the Black–Scholes formulation has eliminated both the Wiener term and the stock growth factor μ . Interestingly, however, the stock's *volatility* σ remains. Essentially, the Black–Scholes evaluation of a stock is simply driven by its volatility, with high volatility being prized. We note that μ has been replaced by the growth rate r of a riskless security. Over a short period of time, the portfolio will be riskless. (We recall how, in the Black–Scholes solution, we used a hedge where we bought options and sold stock simultaneously.) This risklessness will not be maintained at the level of noninfinitesimal time. However, if one readjusts the portfolio,

say, daily, then (making the huge assumption that sudden jumps cannot happen within a short period of time), it could be argued that assuming one knew the current values of r and σ , a profit could be obtained by purchasing options when the market value was below the Black-Scholes valuation and selling them when the market value was above that of the Black-Scholes valuation (assuming no transaction costs). (Such a fact, it could be argued, in which all traders acted on the Black-Scholes valuation, would drive the market. In reality, if the Treasury Bill rate is used for r and historical estimates are used for μ and σ , the actual value for which an option is traded is generally significantly different from the Black-Scholes value.)

Now, we recall that a *European call option* is an instrument which gives the owner the right to purchase a share of stock at the *exercise price* X , T time units from the date of purchase. Naturally, should the stock actually be priced less than X at time T , the bearer will not exercise the option to buy at price X . Although we get to exercise the option only at time T , we must pay for it today. Hence, we must discount the value of an option by the factor $\exp(-rt)$. Since we have seen that the Black-Scholes equation involves no noise term, it is tempting to conjecture that the fair evaluation of an option to purchase a share of stock at exercise price X is given by

$$\begin{aligned} C_{BS} &= e^{-rT} E[\text{Max}(S(T) - X, 0)] \\ &= e^{-rt} \frac{1}{\sqrt{2\pi\sigma^2 t}} \int_{\ln[\frac{X}{S(0)}]}^{\infty} (S(0)e^z - X) \exp\left[-\frac{1}{2\sigma^2 t}(z - (r - \frac{\sigma^2}{2})t)^2\right] dz \end{aligned} \quad (14.64)$$

where we note that the growth rate is r rather than μ .

14.16 Black-Scholes: Some Limiting Cases

Consider, in Tables 14.3 and 14.4, the Black-Scholes pricing model compared to Model A and Model B in the case where a stock has a rather high growth rate $\mu = 0.15$ with a fixed riskless interest rate of 5% and a variety of volatilities and strike prices X . (Of course, we are looking at a case where the option buyer's estimates of the growth rate μ and volatility σ were correct. From the standpoint of the buyer, who decides to buy the call option, standing at time zero, he probably *believes* his estimate for (μ, σ) is correct. One question we should be examine is the value of the option to the buyer, given his current state of information.) We shall assume the option is for an exercise time of six months in the future, and that the price of the stock at the present time is \$100. The purpose of this exercise is simply to look at Black-Scholes in comparison to two older pricing models in the very optimistic case where the person using the model knows μ and σ .

Table 14.3. Six-Month Options: $\sigma = 0.20, \mu = 0.15$.

X	102	104	106	108	110	112	114	116	118	120
C_{BS}	5.89	4.99	4.20	3.51	2.91	2.40	1.98	1.63	1.36	1.16
C_B	8.58	7.47	6.46	5.55	4.73	4.00	3.37	2.82	2.35	1.95
C_A	9.02	7.85	6.79	5.83	4.97	4.21	3.54	2.96	2.47	2.05

Table 14.4. Six-Month Options: $\sigma = 0.40, \mu = 0.15$.

X	102	104	106	108	110	112	114	116	118	120
C_{BS}	11.48	10.63	9.82	9.07	8.37	7.72	7.11	6.54	6.01	5.53
C_B	13.84	12.89	12.00	11.16	10.37	9.62	8.92	8.26	7.64	7.06
C_A	14.55	13.55	12.62	11.73	10.90	10.11	9.37	8.68	8.03	7.43

We note that as the volatility σ increases, the three strategies become more similar. To note the effect of increasing and decreasing σ , we show in Tables 14.5 and 14.6 results for very low σ (.001) and very high σ (2.00).

Table 14.5. Six-Month Options: $\sigma = .001 \mu = .15$.

X	102	104	106	108	110	112	114	116	118	120
C_{BS}	0.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
C_B	5.37	3.51	1.66	0.00	0.00	0.00	0.00	0.00	0.00	0.00
C_A	5.65	3.69	1.74	0.00	0.00	0.00	0.00	0.00	0.00	0.00

**Table 14.6. Six-Month Options:
 $\sigma = 2.00 \mu = .15$.**

EP	102	104	106	108	110	112	114	116	118	120
C_{BS}	52.17	51.71	51.25	50.80	50.36	49.93	49.50	49.08	48.67	48.26
C_B	53.37	52.91	52.45	52.00	51.56	51.13	50.70	50.28	49.87	49.47
C_A	56.11	55.62	55.14	54.67	54.21	53.75	53.30	52.88	52.43	52.00

Let us consider limiting behavior as the volatility goes first to infinity and then to zero. Suppose that a stock is currently selling for $S(0)$. We wish to buy an option T time units in the future with strike price X . As the volatility of the stock goes to infinity, then we note that both Black–Scholes (14.56) and Method B (14.3) tell us that the option is so valuable that its fair price is simply the current value of the stock, namely $S(0)$, irrespective of the value of μ .

On the other hand, let us suppose that the value of the volatility is zero. Then the Black–Scholes price is

$$\begin{aligned} C_{BS} &= S(0) - e^{-rT} X \text{ if } S(0)e^{rT} \geq X \\ &= 0 \text{ otherwise.} \end{aligned} \quad (14.65)$$

Next, let us consider the situation where the growth rate of the stock is actually negative ($\mu = -0.15$) in Tables 14.7 and 14.8, respectively.

Table 14.7. Six-Month Options: $\sigma = .001 \mu = -.15$.

X	102	104	106	108	110	112	114	116	118	120
C_{BS}	0.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
C_B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
C_A	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 14.8. Six-Month Options: $\sigma = 2.00 \mu = -.15$.

X	102	104	106	108	110	112	114	116	118	120
C_{BS}	52.17	51.71	51.25	50.80	50.36	49.93	49.50	49.08	48.67	48.26
C_B	49.77	49.30	48.84	48.39	47.95	47.51	47.08	46.66	46.25	45.84
C_A	45.03	44.61	44.19	43.78	43.38	42.99	42.60	42.22	41.85	41.48

We note that Black-Scholes values a call option at, say \$102, equally whether the growth rate of the stock is +.15 or -.15. Naturally, it is unfair to note that Method A and Method B are more accurate than Black-Scholes. That would be true if we really knew μ , but generally we have only noisy estimates for this parameter. Nevertheless, markets are made, in large measure, by differences in information (opinion).

Table 14.9. Computed Values of Six-Month Options (with Bear Jumps).

$$\sigma = .20, \mu = .15.$$

Ex. Pr.	102	104	106	108	110	112	114	116	118	120
C_{BS}	5.89	4.99	4.20	3.51	2.91	2.40	1.98	1.63	1.36	1.16
C_B	8.58	7.47	6.46	5.55	4.73	4.00	3.37	2.82	2.35	1.95
C_A	9.02	7.85	6.79	5.83	4.97	4.21	3.54	2.96	2.47	2.05
Sim.	5.97	5.13	4.38	3.72	3.13	2.62	2.18	1.80	1.48	1.21

Table 14.10. Six-Month Options (with Bear Jumps.) $\sigma = .40, \mu = .15$

EP	102	104	106	108	110	112	114	116	118	120
C_{BS}	11.48	10.63	9.82	9.07	8.37	7.72	7.11	6.54	6.01	5.53
C_B	13.84	12.89	12.00	11.16	10.37	9.62	8.92	8.26	7.64	7.06
C_A	14.55	13.55	12.62	11.73	10.90	10.11	9.37	8.68	8.03	7.43
Sim.	10.89	10.08	9.33	8.62	7.96	7.34	6.76	6.23	5.73	5.28

Now, based on Tables 14.9 and 14.10, the pricing of the Black-Scholes model appears inspired. Of course, we have simply added on the kind of unexpected downward turn which is not accounted for by the geometric Brownian walk unmodified. On the other hand, our imposition of bear jumps has depressed the expected growth rate of the stock to essentially 1%, and most of the value of the option is due to volatility.

The vendor of the option typically has no strong views about a particular stock. He or she is selling options in many stocks and is only interested that he or she retrieves his or her supposed opportunity cost rate η . Accordingly, the vendor of the option might use Black-Scholes with r replaced by η .

$$C_{\text{vendor}} = e^{-\eta T} \{ e^{\eta T} S(0) \Phi \left(\frac{\log(S(0)/X) + [\eta + (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) - X \Phi \left(\frac{\log(S(0)/X) + [\eta - (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) \}. \quad (14.66)$$

On the other hand, the buyer of the option will have fairly strong views about the stock and its upside potential. The buyer could use Black-Scholes replacing r by μ , where, typically, $\mu > \eta > r$. Thus,

$$C_{\text{buyer}} = e^{-\mu T} \{ e^{\mu T} S(0) \Phi \left(\frac{\log(S(0)/X) + [\mu + (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) - X \Phi \left(\frac{\log(S(0)/X) + [\mu - (\sigma^2/2)]T}{\sigma\sqrt{T}} \right) \}. \quad (14.67)$$

We note that typically, in the mind of the call option buyer, μ is rather large. Perhaps some investors will buy solely on the basis of a large stock volatility, but this is unusual. Option buying is frequently a leveraging device whereby an investor can realize a very large gain by buying call options rather than stocks. The seller of the option is probably expecting an $\eta < \mu$ value as the reasonable rate of return on his/her investments overall. It is observed [1] that the arithmetic mean annual return on U.S. common stocks (including dividends) from 1926 on is over 10%. Let us suppose we are dealing with an initial stock price of \$100 and that the vendor uses $\eta = 0.10$ and the buyer believes $\mu = 0.15$. In Tables 14.11 and 14.12, we show the values of C_{vendor} and C_{buyer} , respectively. This may appear confusing, for we have arrived at a price for the vendor and one for the buyer, and they are generally not the same. *Pareto efficiency* is the situation where all parties are better off by undertaking a transaction. Clearly, at least from their respective viewpoints, we do have Pareto efficiency (assuming that the commission is not so high as to swamp the anticipated profit to the buyer). The difference between the price the buyer is willing to pay and that for which the vendor is willing to sell must be positive, or there will be no trade.

Suppose that an investor believes the rate of growth of a stock is 0.15 overall, bear jumps included. Then, if we are to include the bear jumps, we need to increase the value of the Brownian growth to $0.15 + 0.14 = 0.29$. So, let us now compute the simulated buyer's price, with discount to present value rate being $\mu = 0.15$. We also compute the vendor's price using the Black-Scholes formula with riskless rate $\eta = 0.10$ (we will assume that the vendor will use the nominal volatility values of 0.20 and 0.40, as shown in Tables 14.11 and 14.12).

Table 14.11. Six-Month Options (with Bear Jumps): $\sigma = 0.20$, $\mu = 0.15$.

Ex Pr	102	104	106	108	110	112	114	116	118	120
C_{BS}	5.89	4.99	4.20	3.51	2.91	2.40	1.98	1.63	1.36	1.16
C_{vendor}	7.17	6.16	5.26	4.45	3.74	3.13	2.59	2.13	1.75	1.42
C_{buyer}	10.53	9.34	8.24	7.23	6.30	5.46	4.71	4.04	3.44	2.91

Table 14.12. Six-Month Options (with Bear Jumps): $\sigma = 0.40$, $\mu = 0.15$.

EP	102	104	106	108	110	112	114	116	118	120
C_{BS}	11.48	10.63	9.82	9.07	8.37	7.72	7.11	6.54	6.01	5.53
C_{ven}	12.63	11.73	10.88	10.08	9.34	8.63	7.98	7.36	6.79	6.25
C_{buy}	15.56	14.54	13.58	12.67	11.80	10.98	10.21	9.49	8.81	8.17

It is unlikely that a buyer will be able to acquire options at the orthodox Black-Scholes rate (i.e., the one using Treasury bill interest rates of .05). But suppose she can. Suppose she correctly guesses $\sigma = .2$. The Black-Scholes price for a strike price of \$108 six months in the future is \$3.5357. Suppose she gets lucky and the growth rate over the next six months is $\mu = .15$. However, this is the aggregate growth (including Poissonian bear jumps of size 10% once a year and of size 20% once every five years).

Simulation allows us to see what she can expect. We display the results in the simugram in Figure 14.17. The expected value of the option is \$7.23. However, she should realize that around 55% of the time she will have lost her purchase price of the option. There are many other things the prospective buyer might choose to try before making the decision as to whether or not the option should be bought. Most of these are rather easy to achieve with simulation. The point here is that she should view the purchase of the call option as something risky. Mathematics has not secured for her a free lunch.

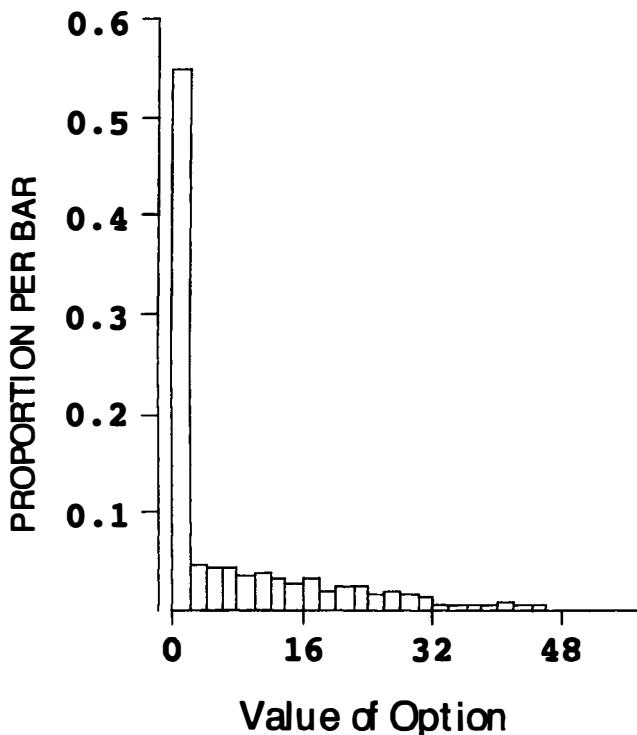


Figure 14.17. Simugram of option (present) values.

In actuality, a buyer of options has many to choose from. Let us suppose that a rational buyer has computed the value of options (using his own estimates for the growth rates and the volatilities). Suppose he has confined himself to stocks with roughly the same volatilities. Suppose further that he is interested in European call options maturing one year from today. If he has no emotional attachment to one stock or the other (big and frequently false assumption), he would then be well advised to purchase options in the stock which give him the largest positive difference between what his computations give him for the value of the option and the actual market price of that option. It will be unusual for the buyer to have the Black-

Scholes price as his personal valuation of the option, and he will seldom see the stock offered for the Black-Scholes price on the market.

So, our hypothetical buyer puts down his money for a one year option in the stock which seems, for a given risk measure, to give him the greatest expected (or median or twenty percentile) return. When he and others buy the option, if the volume of sales over a given time interval appears high, this will encourage vendors of the stock to take note that perhaps they can raise the option's price. Similarly, vendors of options which are not selling at expected volume levels, may decide to lower the prices of their options.

And, as all this is going on, we are aware that there exists a wide diversity of ways for each buyer to consider balancing expected gain against volatility. And there are many time horizons for which he can buy an option. And there are many other possible ways he can invest his capital: in real estate, bonds, wheat, etc. It is the concatenation of all these opportunities for buying and selling, viewed from the standpoints of each buyer and seller, which make up the market. The dynamism of the world in which the market exists is such that any notion of reaching equilibrium for most potential investments is generally unlikely so long as the market is allowed to "work its will." The possibilities for finding undervalued (from the standpoint of the buyer) stocks and derivatives to purchase exist day by day, hour by hour, minute by minute.

14.17 Conclusions

Finally, I cannot emphasize too strongly that the time indexed profile of the probability distribution of the proposed investment is much more reliable than simply looking at expected values. It is frequently the case that the purchase of an option with a high expected value of gain will, most of the time, be a losing proposition.

At the end of the day, an option based on a stock is simply a security to be purchased or not depending on its risk profile and the way an investor views that risk profile in comparison to those of other investments. It might seem absurd that one could come up with a formula which would give, at a given time, strike price and execution time, the value of an option on a stock which is following a random trajectory. Indeed, it is absurd. If transaction costs are truly zero, then the Black-Scholes evaluation of an option is correct *if one is basing evaluation on the expected value of the option*. And if cold fusion were a reality and anti-gravitational devices existed, that would be nice too.

The reality is that transaction costs are not free, and looking only at the expected value of an option (as opposed to its entire risk profile) will frequently lead to disaster. In many ways, option trading has become a useful surrogate for margin buying. Before the Crash of 1929, an investor could use his portfolio to leverage purchasing of stocks by a factor of ten

to one. As the market started downward, the broker would dump stocks in the portfolio to meet margin calls. This put downward pressure on the prices of the stocks being sold. This kind of feedback mechanism led to the Crash of 1929. These days, in the United States, margin leveraging is a more modest two to one.

The purchasing of options appears to be a relatively benign alternative to margin leveraging. If the buyer purchases an option for a strike price which the stock does not realize, then he loses the purchase price of the option without any direct negative pressure on the stock. Indeed, if the stock price is rising, then vendors of uncovered options will have to go into the market and purchase shares to cover the calls, thus putting an upward pressure on the stock price. Though some argue that the availability of options exerts a stabilizing effect on the market, one has to question this judgment. The ability of large companies to sell uncovered options can have disastrous consequences. Enron failed when it could not fulfill the call options in electricity it had sold in California.

The selling of covered call options may be very desirable for a fund manager who is trying to maintain a somewhat steady rate of return in markets good and bad. If the market moves into a phase of low or negative growth, the selling of the call options will bring in some income even though the purchasers will not exercise the options. If the market moves into a good growth phase, then the selling of the call options limits the upside profit from the stock to that of the strike price minus the original cost of the stock to the fund. Suppose that, by very clever balancing, the manager of such a fund managed to obtain a return of, say, 3% in bear markets and 10% in strong markets. That might be the basis for an attractive alternative to bonds. In other words, the vendor might be able to use call options as a way of trading high gains for low risk. In this author's opinion, however, the selling of uncovered call options should not be allowed.

From the standpoint of the buyer of call options, it is frequently the case that the purpose of the trade is to assume high risk in the hopes of substantial gains. There are situations where this can make a great deal of sense. But the author wonders how often the purchasers of call options bother to crank out and examine a risk profile such as that shown in Figure 14.17. The purchaser of a call option ought not believe that the equation of Black, Scholes and Merton will bring determinism into what is, in fact, the very risky business of using options for leveraging purposes.

If one wishes to put his or her investment strategy on autopilot, then finding a good equal weight S&P index fund (such as Vanguard's) is not a bad idea. Even a Fidelity or Vanguard market cap weighted fund may be alright. If one is willing to take the time, one can do better still. What is really unwise is investing in a fund with high management fees managed by individuals who do the investor little good service. And investing in Treasury Bills or derivative funds is probably very dangerous indeed.

The one sure thing about the stock market continues to be that it will

382 STRATEGIES BASED ON DATA VERSUS STRATEGIES BASED ON IDEOLOGY

fluctuate. And this fluctuation produces risk. However, the investor has two weapons at his/her disposal to reduce risk: portfolio diversification and time. The investor also has another weapon: reason and the knowledge that if a deal appears too good to be true, it probably is. Without careful backtesting, investing is beyond risky. The assumption that one can trust the wisdom of mutual fund managers is generally misplaced. Unfortunately, most companies with 401-k plans for its employees generally force them into mutual funds. Backtesting of mutual funds comparing their historical performance to the DOW or S&P 500 is a minimal time investment which can protect the employee. And equity in most mutual funds can be converted into others after a holding period of 90 days. The days of "buy and hold" and "trust thy employer who knoweth best" are gone and not likely to reappear. The area of personal investing is time consuming, but the time is well spent.

Problems

14.1. Find the option values for Method A and Method B for a stock with present value 100, $\mu = 0.10$, $r = 0.06$, $\sigma = 0.1$, strike price $X = 110$ and time horizon one year.

14.2. Verify Smith's Lemma [11].

If S is lognormal with growth rate μ and volatility σ and if

$$\begin{aligned} Q &= \lambda S - \gamma X \text{ if } S - \psi X \geq 0 \\ &= 0 \quad \text{if } S - \psi X < 0, \end{aligned}$$

then

$$\begin{aligned} E(Q) &= \int_{\psi X}^{\infty} (\lambda S - \gamma X) f(S) dS \\ &= e^{\mu T} \lambda S(0) \Phi \left(\frac{\log(S(0)/X) - \log(\psi) + [\mu + (\sigma^2/2)]T}{\sigma \sqrt{T}} \right) \\ &\quad - \gamma X \Phi \left(\frac{\log(S(0)/X) - \log(\psi) + [\mu - (\sigma^2/2)]T}{\sigma \sqrt{T}} \right), \end{aligned} \quad (14.68)$$

where λ , γ , and ψ are arbitrary parameters and Φ is the standard Gaussian cumulative distribution function.⁶

14.3. Consider the Black–Scholes differential equation (11.24) with the boundary condition that

$$\begin{aligned} f(S, T) &= S(T) - X \text{ if } S(T) - X > 0 \\ &= 0, \text{ otherwise.} \end{aligned}$$

Prove that

$$\begin{aligned} f(S, t) &= S \Phi \left(\frac{\log(S/X) + [r + (\sigma^2/2)](T-t)}{\sigma \sqrt{T-t}} \right) \\ &\quad - X e^{[-r(T-t)]} \Phi \left(\frac{\log(S/X) + [r - (\sigma^2/2)](T-t)}{\sigma \sqrt{T-t}} \right). \end{aligned}$$

14.4. A group of investors is considering the possibility of creating a European option-based mutual fund. As a first step in a feasibility study, they

6

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-z^2/2) dz$$

decide to compare investments of \$1 million in a portfolio of 20 stocks as opposed to a portfolio of 20 options. They want to obtain histograms of investment results after one year. Let us suppose that all stocks in the portfolios are bought at a cost of \$100 per share. Let us assume the usual model of stock growth,

$$S(t) = S(0) \exp(\mu t + \sigma \sqrt{t} \epsilon),$$

where ϵ is normally distributed with mean 0 and variance 1. Let us take two values of μ , namely, .10 and .15. Also, let us consider two values of σ , namely, .15 and .30. Consider several strike prices for the options: for example, the expected value of the stock at the end of one year, and various multiples thereof. Assume that the options are completely fungible. Thus, at the end of the year, if a stock is \$10 over the strike price, the option purchased for the Black–Scholes price is worth \$10 (i.e., one does not have to save capital to buy the stock; one can sell the option). For the “riskless” interest rate, use two values: .06 and .08. Clearly, then, we are considering a leveraged portfolio and seeing its performance in relationship to a traditional one. Carry out the study assuming that there is no correlation between the stocks.

14.5. Carry out the study in Problem 14.4 with the following modification. Use the μ value of .24 and the two σ values of .15 and .30. Then assume that there is an across-the-board bear jump mechanism whereby a sudden drop of 10% happens, on the average, once a year and a sudden drop of 20% happens on the average once every five years. The overall growth is still roughly .10. Use the Black–Scholes riskless price as before without adding in the effect of the Poisson jumps downward.

References

- [1] Baggett, L. S., and Thompson, J. R. (2007). "Every man's MaxMedian rule for portfolio management", in *Proceedings of the 13th Army Conference on Applied Statistics*.
- [2] Baxter, M. and Rennie, A. (1996). *Financial Calculus: An Introduction to Derivative Pricing*. New York: Cambridge University Press.
- [3] Black, F. and Scholes, M. (1973). "The pricing of options and corporate liabilities," *Journal of Political Economy*, 81, 637–659.
- [4] Bogle, J .C. (1999). *Common Sense and Mutual Funds: New Imperatives for the Intelligent Investor*. New York: John Wiley & Sons.
- [5] Findlay, M. Chapman, Williams, E. E. and Thompson, J. R. (2003). "Why we all held our breath when the market reopened," *Journal of Portfolio Management*, Spring, 91–100.
- [6] Hull, J.C., (1993). *Options, Futures, and Other Derivative Securities*. Englewood Cliffs, N.J.: Prentice Hall.
- [7] Markowitz, H.(1952). "Portfolio selection," *Journal of Finance*, March, 77–91.
- [8] Markowitz, H. (1959). *Portfolio Selection*, New York: John Wiley & Sons.
- [9] Sharpe, W. E. (1964). "Capital asset prices: A theory of market equilibrium under conditions of risk," *Journal of Finance*, September, 425–442.
- [10] Sharpe, W. E. (2000).*Portfolio Theory and Capital Markets*, New York: McGraw Hill, 102–103.
- [11] Smith, C. W. (1976). "Option pricing: a review," *Journal of Financial Economics*, 3, 3–51.
- [12] Thompson, J. R. (1999). *Simulation: A Modeler's Approach*. New York: John Wiley& Sons, 115–142.
- [13] Thompson, J. R. and Williams, E. E. (1999) "A post Keynesian analysis of the Black-Scholes option pricing model," *The Journal of Post Keynesian Economics*, Winter, 251–267.
- [14] Thompson, J. R., Williams, E. E., and Findlay, M. Chapman (2003). *Models for Investors in Real World Markets*. Hoboken, N.J.: John Wiley Sons.
- [15] Thompson, J. R. Baggett, L. S. and Wojciechowski, W. C. (2004). "Some things economists know that just aren't so," in *The Proceedings of*

the 10th Conference on Applied Statistics.

- [16] Thompson, J. R., Baggett, L. Scott, Wojciechowski, W. C., and Williams, E. E. (2006). “Nobels for nonsense.” *The Journal of Post Keynesian Economics*, Fall, pp. 3–18.
- [17] Williams, E. E. and Findlay, M. Chapman (1974). *Investment Analysis*. Englewood Cliffs, N.J.: Prentice-Hall.
- [18] Wojciechowski, W. C. and Thompson, J. R. (2006). “Market truths: theory versus empirical simulations.” *Journal of Statistical Computation and Simulation*, 76, no 5, 385–395.

Appendix A

A Brief Introduction to Probability and Statistics

A.1 Craps: An Intuitive Introduction to Probability

In this game, played with great gusto by millions, the player throws two six-sided dice. We shall assume that one of these dice is white and the other is black. If, in the first throw, the player throws a seven (W1B6, W2B5, W3B4, W4B3, W5B2, W6B1) or an eleven (W5B6, W6B5), he wins the game. We note that there are 36 possible results of the throw: (W1B1, W1B2, W1B3, W1B4, W1B5, W1B6; W2B1, W2B2, W2B3, W2B4, W2B5, W2B6; W3B1, W3B2, W3B3, W3B4, W3B5, W3B6; W4B1, W4B2, W4B3, W4B4, W4B5, W4B6; W5B1, W5B2, W5B3, W5B4, W5B5, W5B6; W6B1, W6B2, W6B3, W6B4, W6B5, W6B6).

This collection would be looked upon as the *sample space* S . Clearly

$$P(S) = 1.$$

We have to get one of the 36 results. Intuitively, each of these 36 results has the same chance of occurring: 1/36. These 36 elements represent the basic primitive events of the probability space.

We now give an example of a *random variable*. For a toss of the dice:

Let $X = \text{number of white pips} + \text{number of black pips}$.

A.1.1 Random Variables, Their Means and Variances

What is the probability of winning on the first throw? We need to find $P(7 \text{ or } 11)$. How shall do this? We look at the primitive elements which

⁰*Empirical Model Building: Data, Modeling, and Reality*, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

map under the random variable to 7 or 11. These are W1B6, W2B5, W3B4, W4B3, W5B2, W6B1, W5B6, W6B5. Now, we know the probability of each of these primitive events: each has probability 1/36. So, then the concept of a *random variable* is a mapping from the space of primitive events to some other space (here to the integers 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12) in such a way that the inverse map gets us to the primitive events on which the probability is naturally defined. A random variable has the property that the probability that the random variable is equal to a particular value can be computed from the space of original events:

$$\begin{aligned}
 P(7 \text{ or } 11 \text{ on first toss}) &= \\
 P(\text{W1B6, W2B5, W3B4, W4B3, W5B2, W6B1, W5B6, W6B5}) &= \\
 P(\text{W1B6}) + P(\text{W2B5}) + P(\text{W3B4}) + P(\text{W4B3}) + P(\text{W5B2}) + P(\text{W6B1}) + \\
 P(\text{W5B6}) + P(\text{W6B5}) &= \frac{8}{36} = 0.222222.
 \end{aligned}$$

By simply looking at the primitive events that map into 2,3,4,5,6,7,8,9,10, 11,12, we note that

$$\begin{aligned}
 P(2) = P(12) &= \frac{1}{36} \\
 P(3) = P(11) &= \frac{2}{36} \\
 P(4) = P(10) &= \frac{3}{36} \\
 P(5) = P(9) &= \frac{4}{36} \\
 P(6) = P(8) &= \frac{5}{36} \\
 P(7) &= \frac{6}{36}.
 \end{aligned}$$

The *expected value* of a random variable X is its average value. Here

$$\begin{aligned}
 \mu = E(X) &= \sum_{x=2}^{x=12} x P(X) = 2 \frac{1}{36} + 3 \frac{2}{36} + 4 \frac{3}{36} + 5 \frac{4}{36} + \dots \\
 &= 7.
 \end{aligned} \tag{A.1}$$

The *variance* of a random variable is given by

$$\sigma^2 = E(X - \mu)^2 = E(X^2) - 2\mu E(X) + \mu^2 = E(X^2) - \mu^2 = 54.833 - 7^2 = 5.833. \tag{A.2}$$

Now, the rules of the game of craps tell us that the player loses if he gets a 2, 3, or 12 on the first throw. What is the probability of this?

$$\text{Prob (2, 3, or 12 on first throw)} = P(\text{W1B1, W1B2, W2B1, W6B6})$$

$$\begin{aligned}
 &= \frac{1}{36} + \frac{1}{36} + \frac{1}{36} + \frac{1}{36} \\
 &= \frac{4}{36}.
 \end{aligned}$$

Now, the player may throw some number other than 2, 3, 12, 7 or 11. Suppose the number is 4. The probability of this is

$$P(4) = P(W1B3) + P(W2B2) + P(W3B1) = \frac{3}{36}.$$

The rule is that the player wins if he throws a second 4 before throwing a 7. We have seen already that the probability of a 7 is 6/36. So, the probability of getting a 4 before a 7 is

$$\frac{3}{3+6} = \frac{1}{3}.$$

The probability of the player winning the game by throwing a 4 on the first round and then, on subsequent throws, getting a 4 before rolling a 7 is:

$$\frac{3}{36} \times \frac{1}{3} = \frac{2}{72} = 0.027778.$$

Another way the player could win is to throw a 10 on the first round and then throw a second 10 before throwing a 7. Now the probability of throwing a 10 is

$$P(10) = P(W4B6, W5B5, W6B4) = \frac{3}{36},$$

the same as the probability of throwing a 4. Thus, the probability of winning by throwing first a 10 and then throwing another 10 before throwing a 7 is

$$\frac{3}{36} \times \frac{1}{3} = \frac{2}{72} = 0.027778.$$

We then note that the probability of getting a 5 on the first toss is

$$P(5) = P(W1B4, W2B3, W3B2, W4B1) = \frac{4}{36}.$$

So, the probability of winning by throwing first a 5 and then throwing another 5 before getting a 7 is

$$\frac{4}{36} \times \frac{4}{10} = \frac{4}{10} = 0.0444444.$$

By symmetry, we see that this is the same probability as throwing first a 9 and then getting a second 9 before throwing a 7.

Finally, the probability of throwing first a 6 is

$$P(5) = P(W1B5, W2B4, W3B3, W4B2, W5B1) = \frac{5}{36}.$$

Then, the probability of winning by throwing first a 6 and then a second 6 before throwing a 7 is

$$\frac{5}{36} \times \frac{5}{11} = 0.0631313.$$

By symmetry, this is the same as the probability of throwing first an 8 and then throwing another 8 before throwing a 7.

In summary, the probability of winning the game of craps is given by

$$\begin{aligned} P(\text{Winning}) &= P(7 \text{ or } 11 \text{ on first toss}) + \\ 2 \times [P(4 \text{ on first toss})P(4 \text{ before } 7) &+ P(5 \text{ on first toss})P(5 \text{ before } 7) \\ &+ P(6 \text{ on first toss})P(6 \text{ before } 7)] = 0.492928. \end{aligned}$$

Craps is interesting in that it can be used to capture at the level of intuition the key concepts of probability theory. We notice, for example, that all the probabilities in the game are actually generated from the 36 elementary events. Each of these has probability 1/36. But the game itself pays off on the basis of the sum total of the two dice (without regard to their color). The sum total of the two faces is a random variable. For any value of the random variable, we can find the elementary events which form the basis for the necessary computation.

Here, as we have mentioned, there are only eleven throws of interest to us (2,3,4,5,6,7,8,9,10,11,12). Suppose we hear that someone has won at a game of craps, just that he has won. We then want to compute the probability that the person won on the first round. This could only have been done if he had thrown a 7 or an 11. That probability we know is 8/36 = 0.222222. We have, however, the additional information that the player has won. Common sense might lead us to the following formula

$$\begin{aligned} P(\text{win}) P(\text{win on first round} | \text{win}) &= P(\text{win and win on first round}) \\ &= P(\text{win} \cap \text{win first round}) \quad (\text{A.3}) \end{aligned}$$

Here $P(\text{win on first round} | \text{win})$ is termed the "*conditional probability* that he won on the first toss given that he won at all." We can then solve for this conditional probability by using a bit of algebra:

$$P(\text{win on first round} | \text{win}) = \frac{P(\text{win and win on first round})}{P(\text{win})} \quad (\text{A.4})$$

Now, here, we note that the event that he won on the first round implies that he won at all. Hence the solution to (A.2) is given by

$$P(\text{win on first round} | \text{win}) = \frac{0.222222}{0.492928} = 0.450816. \quad (\text{A.5})$$

This discussion should be used as a template any time one needs reminding what a random variable is and how we compute the probability that a random variable has a particular set of values.

Let us get more abstract and ask how many sets could we construct from the 36 primitive elements (without replacement)? Let us construct one such set. We can put any number of primitive elements into the set. The first element W1B1 can either be in our set or out of it. That means two choices for inclusion of W1B1. And the same choice is available as to whether to include W1B2 or not. Aggregating across each of the primitive elements, we find that the number of sets is

$$2 \times 2 \times \dots \times 2 = 2^{36} = 6,871,947,674.$$

We can easily compute the probability P of each of these sets by looking at the primitive events included in each of them and adding up their probabilities (each equal to $1/36$). ¹

¹We can do more and take arbitrarily many unions of these 6,871,947,674 sets and intersections, and complements and so forth and so on *ad infinitum*. The class of sets so obtained is called the *sigma field* generated by the primitive events. Now, as noted, we can find the probability of each of these sets by simply looking to see which primitive events are included.

For example, the case where none of the elementary sets is included is called the *empty set* or \emptyset . We have to get one of the elementary events; something must happen; so the probability that nothing happens is zero. Thus,

$$P(\emptyset) = 0.$$

Now the entire sample space S of all the elementary events (36 of them here) must have probability one, for one of the elementary tooses must happen:

$$P(S) = 1.$$

Now, in the set of real numbers, we can consider the connected intervals as the basic building blocks of a sigma field. And then we can ad infinitum look at unions, intersections, and complements for the real numbers. The resulting sigma field is called the *Borel field*. In order for us to compute the probability of a set B of real numbers, we need to be able to assure ourselves that the inverse $X^{-1}(B)$ is a member of the sigma field in our primitive probability space (where we know what the probabilities are). In such a case,

$$P'(B) = P(X^{-1}(B)).$$

So, a random variable is a real variable such that the inverse image of the Borel sets is a member of the sigma field in the primitive probability space. Thus, the probability measure P and the random variable X will induce a probability measure P' on the Borel sets.

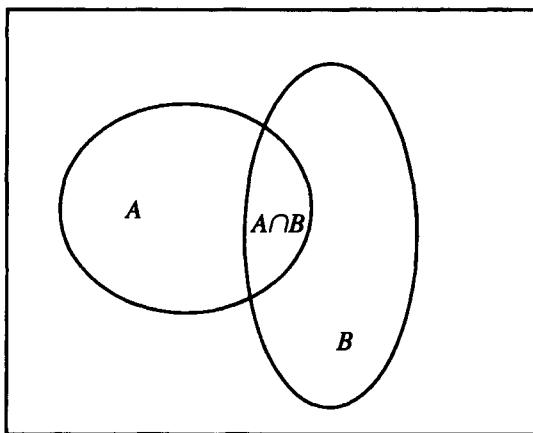


Figure A.1. Sets and their intersections.

As we see from Figure A.1,

$$P(A \cup B) = P(A) + P(B) - P(A \cap B). \quad (\text{A.6})$$

Suppose that we have two sets A and B which have the property that

$$P(A \text{ and } B) = P(A \cap B) = P(A) \times P(B). \quad (\text{A.7})$$

Then, we say that A and B are *stochastically independent* under the probability measure P .

If A and B are *disjoint* (have no points in common), then we can write

$$P(A \text{ and/or } B) = P(A \cup B) = P(A) + P(B). \quad (\text{A.8})$$

For any set A on which a probability measure is defined, we must have

$$0 \leq P(A) \leq 1. \quad (\text{A.9})$$

Let us return briefly to the game of craps. Suppose we have a second random variable, Y , which is equal to 1 if the sum of the two dice is odd and 2 if it is even. Let us view the two random variables in the light of the eleven possible outcomes of X and their probabilities

Table A.1. The Game of Craps.

Primitive Events	X	Y	XY	P(X, Y)
W1B1	2	2	4	1/36
W1B2, W2B1	3	1	3	2/36
W1B3, W2B2, W3B1	4	2	8	3/36
W1B4, W2B3, W3B2, W4B1	5	1	5	4/36
W1B5, W2B4, W3B3, W4B2, W5B1	6	2	12	5/36
W1B6, W2B5, W3B4, W4B3, W5B2, W6B1	7	1	7	6/36
W2B6, W3B5, W4B4, W5B3, W6B2	8	2	16	5/36
W3B6, W4B5, W5B4, W6B3	9	1	9	4/36
W4B6, W5B5, W6B4	10	2	20	3/36
W5B6, W6B5	11	1	11	2/36
W6B6	12	2	24	1/36

Firstly, we can easily compute the mean and variance of Y .

$$\begin{aligned} E(Y) &= 2 \times 1/36 + 3 \times 2/36 + \dots = 1.5 \\ E(Y^2) &= 2^2 \times 1/36 + 3^2 \times 2/36 + \dots = 2.5 \\ Var(Y) &= E(Y^2) - [E(Y)]^2 = 2.5 - 1.5^2 = .25. \end{aligned}$$

Generally, we may try to find a simple measure of the apparent interaction between two variables, X and Y . One such is the covariance of X and Y

$$Cov(X, Y) = E[(X - \mu_X)(Y - \mu_Y)]. \quad (\text{A.10})$$

A more popular measure of apparent interaction is the *correlation* between X and Y ,

$$\rho(X, Y) = \frac{Cov(X, Y)}{\sqrt{\sigma_X^2 \sigma_Y^2}}. \quad (\text{A.11})$$

Now, although covariances can take values from $-\infty$ to $+\infty$, the correlation can only take values between -1 and $+1$. To prove this fact, we note that:

$$0 \leq E[a(X - \mu_X) - (Y - \mu_Y)]^2 = a^2 \sigma_X^2 + \sigma_Y^2 - 2aCov(X, Y), \quad (\text{A.12})$$

where a is an arbitrary real constant which we elect to be $Cov(X, Y)/\sigma_X^2$. This gives us immediately a version of *Cauchy's Inequality*,

$$\rho^2 \leq 1. \quad (\text{A.13})$$

Now, the reader should verify that if Y is simply a positive multiple of X , then $\rho = 1$. If Y is a negative multiple of X , then $\rho = -1$. We are interested, in portfolio design, in looking at the correlation between two securities. To the extent that this correlation is close to 1, the diversification benefit of including both stocks in the portfolio is marginal.

A.2 Combinatorics Basics

Let us first compute the number of ways that we can arrange in a distinctive order k objects selected without replacement from n , $n \geq k$, distinct objects. We easily see that there are n ways of selecting the first object, $n - 1$ ways of selecting the second object, and so on until we select $k - 1$ objects and note that the k th object can be selected in $n - k + 1$ ways. The total number of ways is called the *permutation* of n objects taken k at a time, $P(n, k)$, and is seen to be given by

$$P(n, k) = n(n - 1)(n - 2) \cdots (n - k + 1) = \frac{n!}{(n - k)!}, \quad (\text{A.14})$$

where $m! = m(m - 1)(m - 2) \cdots 2 \times 1$, $0! = 1$. In particular, there are $n!$ ways that we can arrange n objects in a distinctive order. Next, let us compute in how many ways we can select k objects from n objects when we are not concerned with the distinctive order of selection. This number of ways is called the *combination* of n objects taken k at a time, and is denoted by $C(n, k)$. We can find it by noting that $P(n, k)$ could be first computed by finding $C(n, k)$ and then multiplying it by the number of ways k objects could be distinctly arranged (i.e., $k!$). So we have

$$P(n, k) = C(n, k)P(k, k) = C(n, k)k!$$

and thus

$$\binom{n}{k} = C(n, k) = \frac{n!}{(n - k)!k!}. \quad (\text{A.15})$$

For example, the game of stud poker consists in the drawing of 5 cards from a 52 card deck (4 suits, 13 denominations). The number of possible hands is given by

$$C(52, 5) = \frac{52!}{47!15!} = 2,598,960.$$

We are now in a position to compute some basic probabilities which are slightly harder to obtain than, say, those concerning tossing a die. Each of the 2,598,960 possible poker hands is equally likely. To compute the probability of a particular hand, we simply evaluate

$$P(\text{hand}) = \frac{\text{number of ways of getting the hand}}{\text{number of possible hands}}.$$

Suppose we wish to find the probability of getting an all-spade hand. There are $C(13, 5)$ ways of selecting 5 spades (without regard to their order) out of 13 spades. Hence,

$$\begin{aligned} P(\text{an all spade hand}) &= \frac{C(13, 5)}{C(52, 5)} \\ &= \frac{(9)(10)(11)(12)(13)}{(5!)(2,598,960)} = 0.0000495. \end{aligned}$$

Finding the probability of getting four cards of a kind (e.g., four aces, four kings) is a bit more complicated. There are $C(13, 1)$ ways of picking a denomination, $C(4, 4)$ ways of selecting all the four cards of the same denomination, and $C(48, 1)$ of selecting the remaining card. Thus,

$$\begin{aligned} P(\text{four of a kind}) &= \frac{C(13, 1)C(4, 4)C(48, 1)}{C(52, 5)} \\ &= \frac{(13)(1)(48)}{2, 598, 960} = 0.00024. \end{aligned}$$

Similarly, to find the probability of getting two pairs, we have

$$\begin{aligned} P(\text{two pairs}) &= \frac{C(13, 2)C(4, 2)C(4, 2)C(44, 1)}{C(52, 5)} \\ &= \frac{(78)(6)(6)(44)}{2, 598, 960} = 0.0475. \end{aligned}$$

A.3 Bayesian Statistics

A.3.1 Bayes's Theorem

Suppose that the sample space S can be written as the union of disjoint sets: $S = A_1 \cup A_2 \cup \dots \cup A_n$. Let the event H be a subset of S which has nonempty intersections with some of the A_i 's. Then

$$P(A_i|H) = \frac{P(H|A_i)P(A_i)}{P(H|A_1)P(A_1) + P(H|A_2)P(A_2) + \dots + P(H|A_n)P(A_n)}. \quad (\text{A.16})$$

To explain the conditional probability given by equation (A.16), consider a diagram of the sample space, S . Consider that the A_i 's represent n disjoint states of nature. The event H intersects some of the A_i 's.

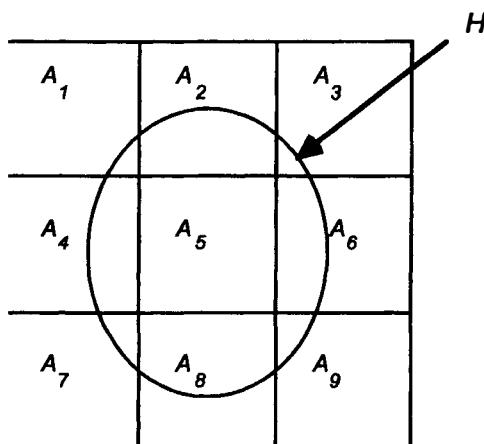


Figure A.2. Bayesian Venn diagram.

Then,

$$P(H|A_1) = \frac{P(H \cap A_1)}{P(A_1)} = \frac{P(A_1|H)P(H)}{P(A_1)}.$$

Solving for $P(A_1|H)$, we get

$$P(A_1|H) = \frac{P(H|A_1)P(A_1)}{P(H)},$$

and in general,

$$P(A_i|H) = \frac{P(H|A_i)P(A_i)}{P(H)}. \quad (\text{A.17})$$

Now,

$$\begin{aligned} P(H) &= P((H \cap A_1) \cup (H \cap A_2) \cup \dots \cup (H \cap A_n)) \\ &= \sum P(H \cap A_i), \text{ since the intersections } (H \cap A_i) \text{ are disjoint} \\ &= \sum P(H|A_i)P(A_i), \text{ where } i = 1, 2, \dots, n. \end{aligned}$$

Thus, with (A.17) and $P(H)$ given as above, we get (A.16).

The formula (A.16) finds the probability that the true state of nature is A_i given that H is observed. Notice that the probabilities $P(A_i)$ must be known to find $P(A_i|H)$. These probabilities are called *prior* probabilities because they represent information prior to experimental data. The $P(A_i|H)$ are then *posterior* probabilities. For each $i = 1, 2, \dots, n$, $P(A_i|H)$ is the probability that A_i was the state of nature in light of the occurrence of the event H .

A.3.2 A Diagnostic Example

Consider patients being tested for a particular disease. It is known from historical data that 5% of the patients tested have the disease, further, that 10% of the patients that have the disease test negative for the disease, and that 20% of the patients who do not have the disease test positive for the disease. Denote by D^+ the event that the patient has the disease, by D^- the event that the patient does not, and denote by T^+ the event the patient tests positive for the disease, and by T^- the event the patient tests negative.

If a patient tests positive for the disease, what is the probability that the patient actually has the disease? We seek the conditional probability, $P(D^+|T^+)$. Here, T^+ is the observed event, and D^+ may be the true state of nature that exists prior to the test. (We trust that the test does not cause the disease.) Using Bayes's theorem,

$$\begin{aligned}
 P(D^+|T^+) &= \frac{P(T^+|D^+)P(D^+)}{P(T^+)} & (A.18) \\
 &= \frac{P(T^+|D^+)P(D^+)}{P(T^+|D^+)P(D^+) + P(T^+|D^-)P(D^-)} \\
 &= \frac{0.9 \times 0.05}{0.9 \times 0.05 + 0.2 \times 0.95} = 0.1915.
 \end{aligned}$$

Thus, there is nearly a 20% chance given a positive test result that the patient has the disease. This probability is the posterior probability, and if the patient is tested again, we can use it as the new prior probability. If the patient tests positive once more, we use equation (A.18) with an updated version of $P(D^+)$, namely, 0.1915.

The posterior probability now is:

$$\begin{aligned}
 P(D^+|T^+) &= \frac{P(T^+|D^+)P(D^+)}{P(T^+)} \\
 &= \frac{P(T^+|D^+)P(D^+)}{P(T^+|D^+)P(D^+) + P(T^+|D^-)P(D^-)} \\
 &= \frac{0.9 \times 0.1915}{0.9 \times 0.1915 + 0.2 \times 0.8085} = 0.5159.
 \end{aligned}$$

Twice the patient tests positive for the disease and the posterior probability that the patient has the disease is now much higher. As we gather more and more information with further tests, our posterior probabilities will better and better describe the true state of nature.

In order to find the posterior probabilities as we have done, we needed to know the prior probabilities. A major concern in a Bayes application is the choice of priors, a choice which must be made sometimes with very little prior information. One suggestion made by Bayes is to assume that the n states of nature are equally likely (Bayes' Axiom). If we make this

assumption in the example above, that is, that $P(D^+) = P(D^-) = 0.5$, then

$$P(D^+|T^+) = \frac{P(T^+|D^+)P(D^+)}{P(T^+|D^+)P(D^+) + P(T^+|D^-)P(D^-)}$$

$P(D^+)$ and $P(D^-)$ cancel, giving

$$\begin{aligned} P(D^+|T^+) &= \frac{P(T^+|D^+)}{P(T^+|D^+) + P(T^+|D^-)} \\ &= \frac{0.9}{0.9 + 0.2} \\ &= 0.8182. \end{aligned}$$

This is much higher than the accurate probability, 0.1912. Depending upon the type of decisions an analyst has to make, a discrepancy of this magnitude may be very serious indeed. As more information is obtained, however, the effect of the initial choice of priors will become less severe.

A.4 The Binomial Distribution

Let us suppose we are selecting from a very large—effectively infinite—population of black and white balls. Suppose the probability a ball is black is p and that the probability a ball is white is $q = 1 - p$. Suppose that out of n draws, the first x are black and the next $n - x$ are white. Suppose that out of n draws we get x black balls and $n - x$ white balls. The probability is given by

$$\underbrace{pp \cdots p}_{x} \underbrace{(1-p)(1-p) \cdots (1-p)}_{n-x \text{ times}} = p^x(1-p)^{n-x}.$$

But suppose that we are not interested in the order in which the black balls appear, just their total number x out of n draws. Then we have the *binomial probability function*

$$P(X = x) = \binom{n}{x} p^x(1-p)^{n-x}, \quad x = 0, 1, 2, \dots, n. \quad (\text{A.19})$$

The binomial distribution may be viewed as the sum of n independent *Bernoulli variables*. A Bernoulli variable, Y , takes the value 1 with probability p and the value 0 with probability $q = 1 - p$. Thus, the expected value of $Y = p \times 1 + q \times 0 = p$. Similarly, the expected value of $Y^2 = p \times 1^2 = p$. Then the variance of $Y = E(Y^2) - [E(Y)]^2 = p - p^2 = p(1 - p)$.

Returning to the binomial variable x , we have

$$X = y_1 + y_2 + \dots + y_n. \quad (\text{A.20})$$

So, for the binomial distribution,

$$\begin{aligned}
 E(X) &= E(y_1) + E(y_2) + \dots + E(y_n) \\
 &= p + p + \dots + p \\
 &= np.
 \end{aligned} \tag{A.21}$$

$$E(X^2) = E(y_1^2 + y_2^2 + \dots + y_n^2 + n^2 - n \text{ terms like } y_i y_j \text{ where } i \neq j). \tag{A.22}$$

Thus,

$$E(X^2) = np + n(n-1)p^2. \tag{A.23}$$

And

$$Var(X) = E(X^2) - [E(X)]^2 = np + n(n-1)p^2 - (np)^2 = np(1-p). \tag{A.24}$$

We now bring in a concept which is used extensively throughout this book, that of the *cumulative probability distribution function*:

$$F(x) = P(X \leq x). \tag{A.25}$$

For the binomial distribution, we have

$$F(x) = \sum_{j=0}^{j \leq x} \binom{n}{j} p^j (1-p)^{n-j}, \quad j = 0, 1, 2, \dots, n. \tag{A.26}$$

We note that the binomial distribution is discrete. $F(x)$ is described by step functions. We show the (cumulative) distribution function of a binomial variate when $p = 0.7$ and $n = 3$ in Figure A.3.

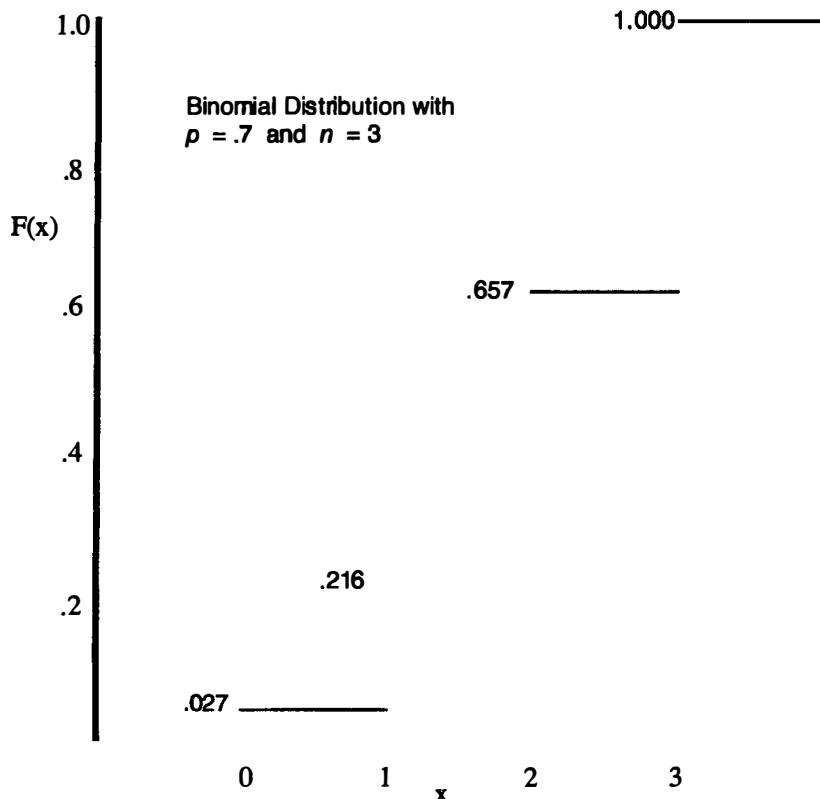


Figure A.3. CDF of the binomial distribution.

A.5 The Uniform Distribution

We now look at a random variable X which is characterized by its cdf

$$\begin{aligned}
 F(x) &= 0, \text{ if } x < 0 \\
 &= x, \text{ if } 0 \leq x \leq 1 \\
 &= 1, \text{ if } x > 1.
 \end{aligned} \tag{A.27}$$

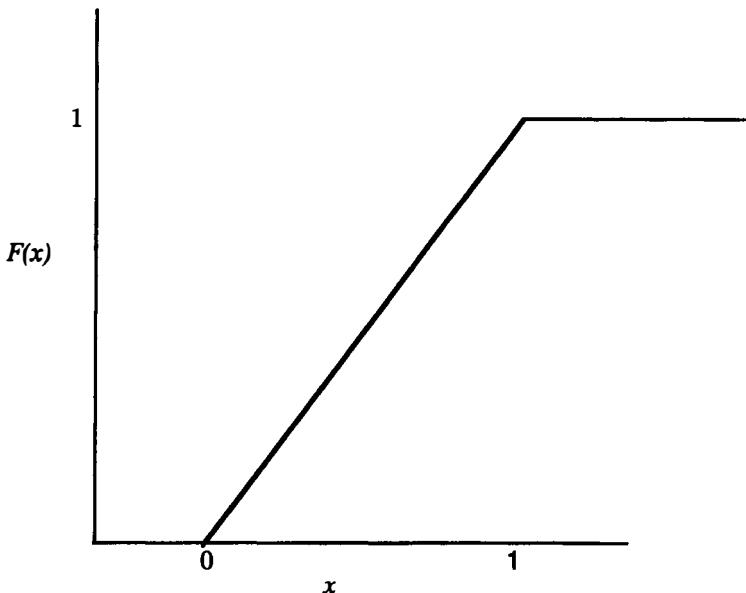


Figure A.4. CDF of the uniform distribution.

The uniform distribution has no jumps in the cdf. A uniform random variable is, therefore, an example of a *continuous random variable*. For a continuous random variable, we may desire to look at its derivative, $f(x) = dF(x)/dx$. $f(x)$ is called the *probability density function* of the random variable X . In the case of the uniform distribution, we note that

$$f(x) = \frac{dF(x)}{dx} = 1 \text{ for } 0 \leq x \leq 1; \\ = 0 \text{ otherwise.}$$

For simulation purposes, the uniform distribution is of particular importance. Suppose that we have another continuous random variable X with cdf $F(\cdot)$. We will now consider F as itself a random variable, $Y = F(x)$. Its cdf $G(y)$ is easily found by the following argument:

$$G(y) = P(Y \leq y) = P(F(x) \leq y) = P(x \leq F^{-1}(y)) = y. \quad (\text{A.28})$$

The consequence is that if we know the cdf $F(x)$, then we may obtain a simulated value from the distribution by finding u randomly distributed on the unit interval. Then

$$x = F^{-1}(u). \quad (\text{A.29})$$

Practically every numerical computer compiler has a uniform random number generator. So (A.29) will generally yield an easy way for us to generate a simulated observation of the random variable with cdf F .

A.6 Moment-Generating Functions

We now consider the joint density function of n random independent and identically distributed random variables of a continuous random variable X having density function $f(\cdot)$. Then a natural definition of the joint density of (x_1, x_2, \dots, x_n) is

$$\begin{aligned}
 f(x_1, x_2, \dots, x_n) &= \lim_{\epsilon_1, \epsilon_2, \dots, \rightarrow 0} \frac{P[x_1 < X_1 < x_1 + \epsilon_1]}{\epsilon_1} \frac{P[x_2 < X_2 < x_2 + \epsilon_2]}{\epsilon_2} \dots \\
 &= \lim_{\epsilon_1, \epsilon_2, \dots, \rightarrow 0} \frac{f(x_1)\epsilon_1}{\epsilon_1} \frac{f(x_2)\epsilon_2}{\epsilon_2} \dots \frac{f(x_n)\epsilon_n}{\epsilon_n} \\
 &= f(x_1)f(x_2) \dots f(x_n). \tag{A.30}
 \end{aligned}$$

The late Salomon Bochner once mentioned the rather modest result in (A.30) as being R. A. Fisher's greatest contribution to statistics. Note that it enables us to write the density of an n -dimensional random variable as the product of n one-dimensional densities.

Next, let X be a random variable with cumulative distribution function $F(x)$. The *moment-generating function* $M_X(t)$ via

$$M_X(t) = E(e^{tX}), \tag{A.31}$$

where t is an arbitrary real variable.

Assuming that differentiation with respect to t commutes with expectation operator E , we have

$$\begin{aligned}
 M'_X(t) &= E(Xe^{tX}) \\
 M''_X(t) &= E(X^2e^{tX}) \\
 &\vdots \\
 M_X^{(k)}(t) &= E(X^{(k)}e^{tX}).
 \end{aligned}$$

Setting t equal to zero, we see that

$$M_X^{(k)}(0) = E(X^k). \tag{A.32}$$

Thus, we see immediately the reason for the name moment-generating function (m.g.f.). Once we have obtained $M_X(t)$, we can compute moments of arbitrary order (assuming they exist) by successively differentiating the m.g.f. and setting the argument t equal to zero. As an example of this application, let us consider a random variable distributed according to the binomial distribution with parameters n and p . Then,

$$\begin{aligned}
 M_X(t) &= \sum_0^n e^{tx} \binom{n}{x} p^x (1-p)^{n-x} \\
 &= \sum_0^n \binom{n}{x} (pe^t)^x (1-p)^{n-x}.
 \end{aligned}$$

Now recalling the binomial identity

$$\sum_0^n \binom{n}{x} a^x b^{n-x} = (a+b)^n,$$

we have

$$M_X(t) = [pe^t + (1-p)]^n. \quad (\text{A.33})$$

Ne t, differentiating with respect to t, we have

$$M'_X(t) = npe^t [pe^t + (1-p)]^{n-1}. \quad (\text{A.34})$$

Then, setting t equal to zero, we have

$$E(X) = M'_X(0) = np. \quad (\text{A.35})$$

Differentiating again with respect to t and setting t equal to zero, we have

$$E(X^2) = M''_X(0) = np + n(n-1)p^2. \quad (\text{A.36})$$

To calculate the variance, it suffices to recall that for any r.v. X we have

$$\text{Var}(X) = E(X^2) - [E(X)]^2. \quad (\text{A.37})$$

Thus, for the binomial X ,

$$\text{Var}(X) = np(1-p). \quad (\text{A.38})$$

Of course we have already found the mean and variance of the binomial distribution via (A.20)–(A.24). Generally speaking the moment–generating function is an easier way to compute moments than the direct approach. However, we shall shortly see an even more important use of the moment generating function.

A.7 The Normal (Gaussian) Distribution

Consider the normal density function

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right), \quad -\infty < x < \infty. \quad (\text{A.39})$$

We would like to satisfy ourselves that we have a true density function. We note, first of all, that $f(x) > 0$ for all $-\infty < x < \infty$. Ne t, we need to show that $F(\infty) = 1$, i.e, that

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right) dx = 1. \quad (\text{A.40})$$

Let us make the transformation

$$z = \frac{x - \mu}{\sigma}. \quad (\text{A.41})$$

The left hand side of (A.40) becomes

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz = A. \quad (\text{A.42})$$

Clearly, A is non-negative. Hence it will suffice to show that $A^2 = 1$. Now

$$A^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \exp[-\frac{1}{2}(z^2 + w^2)] dz dw. \quad (\text{A.43})$$

Let us transform to polar coordinates, with

$$r^2 = z^2 + w^2; \tan(\theta) = w/z.$$

Thus,

$$A^2 = \frac{1}{2\pi} \int_0^{\infty} \int_0^{2\pi} e^{-r^2/2} r dr d\theta = \frac{1}{2\pi} 2\pi \int_0^{\infty} e^{-r^2/2} r dr = 1. \quad (\text{A.44})$$

The moment-generating function of a normal random variable can be found via

$$\begin{aligned} M_X(t) &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} e^{tx} \exp\left(-\frac{1}{2\sigma^2}(x - \mu)^2\right) dx \\ &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2}(x^2 - 2\mu x - 2\sigma^2 tx + \mu^2)\right) dx \\ &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2}(x^2 - 2x(\mu + t\sigma^2) + \mu^2)\right) dx \\ &= \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2\sigma^2}(x - \mu^*)^2\right) dx \exp\left(t\mu + \frac{t^2\sigma^2}{2}\right), \end{aligned}$$

where $\mu^* = \mu + t\sigma^2$. But recognizing that the integral is simply equal to $\sqrt{2\pi}\sigma$, we see that the m.g.f. of the normal distribution is given by

$$M_X(t) = \exp\left(t\mu + \frac{t^2\sigma^2}{2}\right). \quad (\text{A.45})$$

By evaluating the first two derivatives of $M_X(t)$ at 0, the reader may now easily verify that the mean and variance of the normal distribution are μ and σ^2 , respectively.²

²A related distribution, of particular interest to persons involved with market models, is the *lognormal distribution*. Suppose that we have a random variable X such that its logarithm is normally distributed with mean μ and variance σ^2 . Then we say that X has the lognormal distribution with density:

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp\left(-\frac{1}{2\sigma^2}(\ln x - \mu)^2\right) \text{ for } x > 0.$$

The possible mechanical advantages of the m.g.f. are clear. One integration (summation) operation plus k differentiations yield the first k moments of a random variable. However, the moment-generating aspect of the m.g.f. pales in importance to some of its properties relating to the summation of independent random variables. Let us suppose, for example, that we have n independently distributed r.v.'s X_1, X_2, \dots, X_n with m.g.f.'s M_1, M_2, \dots, M_n , respectively. Suppose that we wish to investigate the distribution of the r.v.

$$Y = c_1 X_1 + c_2 X_2 + \dots + c_n X_n,$$

where c_1, c_2, \dots, c_n are fixed constants. Let us consider using the moment-generating functions to achieve this task. We have

$$M_Y(t) = E[\exp \{t(c_1 X_1 + c_2 X_2 + \dots + c_n X_n)\}].$$

Using the independence of X_1, X_2, \dots, X_n we may write

$$\begin{aligned} M_Y(t) &= E[\exp t c_1 X_1] E[\exp t c_2 X_2] \cdots E[\exp t c_n X_n] \\ &= M_1(c_1 t) M_2(c_2 t) \cdots M_n(c_n t). \end{aligned} \quad (\text{A.46})$$

Given the density (or probability) function, we know what the m.g.f. will be. But it turns out that, under very general conditions, the same is true in the reverse direction; namely, if we know $M_X(t)$, we can compute a unique density (probability) function that corresponds to it. The practical implication is that if we find a random variable with an m.g.f. we recognize as corresponding to a particular density (probability) function, we know immediately that the random variable has the corresponding density (probability) function. Thus, in many cases, we are able to use (A.46) to give ourselves immediately the distribution of Y . Consider, for example, the sum

$$Y = X_1 + X_2 + \dots + X_n$$

of n independent binomially distributed r.v.'s with the same probability of success p and the other parameter being equal to n_1, n_2, \dots, n_n , respectively. Thus, the moment-generating function for Y is

$$\begin{aligned} M_Y(t) &= [pe^t + (1-p)]^{n_1} [pe^t + (1-p)]^{n_2} \cdots [pe^t + (1-p)]^{n_n} \\ &= [pe^t + (1-p)]^{n_1+n_2+\dots+n_n}. \end{aligned}$$

We note that, not unexpectedly, this is the m.g.f. of a binomial r.v. with parameters $N = n_1 + n_2 + \dots + n_n$ and p .

Next, we note that the moment generating function for

$$Z = c_1 X_1 + c_2 X_2 + \dots + c_n X_n$$

where the X_j are independent normal variables with parameters μ_j and σ_j^2 , is simply

$$M_Z(t) = \exp \left(t \sum_{i=1}^n c_i \mu_i + \frac{t^2}{2} \sum_{i=1}^n (c_i \sigma_i)^2 \right). \quad (\text{A.47})$$

From (A.45), we recognize that Z must be a normal random variable with mean $\sum_{i=1}^n c_i \mu_i$ and variance $\sum_{i=1}^n c_i^2 \sigma_i^2$.

A.8 The Central Limit Theorem

We are now in a position to derive one version of the *central limit theorem* (CLT). Let us suppose we have a sample X_1, X_2, \dots, X_n of independently and identically distributed random variables with mean μ and variance σ^2 . We wish to determine, for n large, the approximate distribution of the sample mean

$$\bar{X} = \frac{X_1 + X_2 + \dots + X_n}{n}.$$

We shall examine the distribution of the sample mean when put into the standard form. Let

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} = \frac{X_1 - \mu}{\sigma\sqrt{n}} + \frac{X_2 - \mu}{\sigma\sqrt{n}} + \dots + \frac{X_n - \mu}{\sigma\sqrt{n}}.$$

Now, using the independence of the X_i 's and the fact that they are identically distributed with the same mean and variance, we can write

$$\begin{aligned} M_Z(t) &= E(e^{tZ}) = \prod_{i=1}^n E\left[\exp\left(t \frac{X_i - \mu}{\sigma\sqrt{n}}\right)\right] \\ &= \left\{E\left[\exp\left(t \frac{X_1 - \mu}{\sigma\sqrt{n}}\right)\right]\right\}^n \\ &= \left\{E\left[1 + t \frac{X_1 - \mu}{\sigma\sqrt{n}} + \frac{t^2}{2} \frac{(X_1 - \mu)^2}{\sigma^2 n} + o\left(\frac{1}{n}\right)\right]\right\}^n \\ &= \left(1 + \frac{t^2}{2n}\right)^n \rightarrow e^{t^2/2} \text{ as } n \rightarrow \infty. \end{aligned} \quad (\text{A.48})$$

But (A.48) is the m.g.f. of a normal distribution with mean zero and variance one. Thus, we have been able to show that the distribution of the sample mean of a random sample of n i.i.d. random variables with mean μ and variance σ^2 becomes "close" to the normal distribution with mean μ and variance σ^2/n as n becomes large.

Perhaps the easiest method of remembering the CLT is that if a statistic is the result of a summing process, then

$$Z = \frac{\text{statistic} - E(\text{statistic})}{\sqrt{\text{Var}(\text{statistic})}} \quad (\text{A.49})$$

is approximately normally distributed with mean 0 and variance 1.

A.9 The Gamma Distribution

Consider the *gamma function*:

$$\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} e^{-x} dx \text{ for } \alpha > 0. \quad (\text{A.50})$$

Integrating by parts, we obtain

$$\Gamma(\alpha) = (\alpha - 1)\Gamma(\alpha - 1) \text{ for } \alpha > 1. \quad (\text{A.51})$$

When $\alpha = n$, with n a positive integer, repeating (A.51) $n - 1$ times yields

$$\Gamma(n) = (n - 1)!, \quad (\text{A.52})$$

since $\Gamma(1) = \int_0^\infty e^{-x} dx = 1$. The random variable X has the *gamma distribution* with parameters α and β , if its p.d.f. is

$$f(x) = \frac{1}{\beta^\alpha \Gamma(\alpha)} x^{\alpha-1} e^{-x/\beta} \text{ for } x > 0 \quad (\text{A.53})$$

and zero elsewhere, where both constants, α and β , are positive. The mean of X is $\alpha\beta$ and the variance of X is $\alpha\beta^2$.

The gamma distribution with parameter $\alpha = 1$ is called the (*negative exponential distribution* with parameter β). That is, the exponential r.v. X has the p.d.f.

$$f(x) = \frac{1}{\beta} e^{-x/\beta} \text{ for } x > 0, \quad (\text{A.54})$$

and zero elsewhere, where $\beta > 0$. It also follows from the above that X has the mean β and variance β^2 .

The gamma distribution with parameters $\alpha = \nu/2$ and $\beta = 2$, where ν is a positive integer, is called the *chi-square* (χ_ν^2 for short) distribution with ν degrees of freedom. The chi-square random variable X has the p.d.f.

$$f(x) = \frac{1}{2^{\nu/2} \Gamma(\nu/2)} x^{\nu/2-1} e^{-x/2} \text{ for } x > 0 \quad (\text{A.55})$$

and zero elsewhere. The r.v. X has the mean ν and variance 2ν .

The m.g.f. of a gamma variate with parameters α and β can be computed in the following way:

$$\begin{aligned} M(t) &= \frac{1}{\Gamma(\alpha)\beta^\alpha} \int_0^\infty e^{tx} x^{\alpha-1} e^{-x/\beta} dx \\ &= \frac{1}{\Gamma(\alpha)\beta^\alpha} \int_0^\infty x^{\alpha-1} e^{-x(1-\beta t)/\beta} dx; \end{aligned}$$

now this integral is finite only for $t < 1/\beta$ and substituting $y = x(1 - \beta t)/\beta$ yields

$$\begin{aligned} M(t) &= \frac{1}{\Gamma(\alpha)} \int_0^\infty y^{\alpha-1} e^{-y} dy \left(\frac{1}{1 - \beta t} \right)^\alpha \\ &= \left(\frac{1}{1 - \beta t} \right)^\alpha \text{ for } t < \frac{1}{\beta}, \end{aligned} \quad (\text{A.56})$$

where the last equality follows from the form of the p.d.f. of the gamma distribution with parameters α and 1. In particular, the m.g.f. for a chi-square r.v. with ν degrees of freedom has the form

$$M(t) = \left(\frac{1}{1 - 2t} \right)^{\nu/2} = (1 - 2t)^{-\nu/2}, \quad t < 1/2. \quad (\text{A.57})$$

Suppose we consider the moment generating function of the square of a $\mathcal{N}(0, 1)$ random variable.

$$\begin{aligned} M_{Z^2}(t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^\infty e^{tz^2} e^{-z^2/2} dz \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^\infty \exp\left(-\frac{1}{2}z^2(1 - 2t)\right) dz \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^\infty e^{\frac{w^2}{2}} dw (1 - 2t)^{\frac{1}{2}} \\ &= (1 - 2t)^{\frac{1}{2}}, \end{aligned} \quad (\text{A.58})$$

where $w = z\sqrt{1 - 2t}$.

Next, let us consider the sum of squares of n random variables independently distributed as $\mathcal{N}(0, 1)$. That is, we wish to consider the moment generating function of

$$\chi^2 = \sum_{i=1}^n z_i^2.$$

Then, from (A.45) and (A.46), we have

$$M_{\chi^2}(t) = \left(\frac{1}{1 - 2t} \right)^{n/2} \quad (\text{A.59})$$

We recognize this to be a χ^2 variable with n degrees of freedom.

A.10 Conditional Density Functions

Let us return to questions of interdependence between random variables and consider briefly conditional distribution of one random variable given that another random variable has assumed a fixed value.

If two random variables X and Y are discrete and have a joint probability function $f(x, y)$, then, the *conditional probability function* of the r.v. X , given that $Y = y$, has the form

$$f(x_i|y) = P(X = x_i|Y = y) = \frac{f(x_i, y)}{f_Y(y)}, \quad (\text{A.60})$$

where

$$f_Y(y) = \sum_{\text{all values of } x_j} f(x_j, y). \quad (\text{A.61})$$

Next, let us now suppose that random variables X and Y are continuous and have joint c.d.f.

$$F(x, y) = \int_{-\infty}^x \int_{-\infty}^y f(u, v) du dv. \quad (\text{A.62})$$

We can obtain the *marginal density function* of Y via

$$f_Y(y) = \int_{-\infty}^{\infty} f(x, y) dx. \quad (\text{A.63})$$

Writing the statement of joint probability for small intervals in X and Y , we have

$$\begin{aligned} & P(x < X \leq x + \varepsilon \cap y < Y \leq y + \delta) \\ & P(y < Y \leq y + \delta)P(x < X \leq x + \varepsilon | y < Y \leq y + \delta). \end{aligned}$$

Now, exploiting the assumption of continuity of the density function, we can write

$$\begin{aligned} \int_x^{x+\varepsilon} \int_y^{y+\delta} f(x, y) dy dx &= \int_y^{y+\delta} f_Y(y) dy \int_x^{x+\varepsilon} f_{X|y}(x) dx \\ &= \varepsilon \delta f(x, y) = \delta f_Y(y) \varepsilon f_{X|y}(x). \end{aligned}$$

Here, we have used the terms f_Y and $f_{X|y}$ to denote the *marginal density function* of Y , and the *conditional density function* of X given $Y = y$, respectively. This gives us immediately

$$f_{X|y}(x) = \frac{f(x, y)}{f_Y(y)}. \quad (\text{A.64})$$

Note that this is a function of the argument x , whereas y is fixed; y is the value assumed by the random variable Y .

A.11 The Weak Law of Large Numbers

Let us now consider the set of n data drawn from some probability distribution. Prior to the experiment which yields the data, they can be treated as a sequence of n independent and identically distributed (i.i.d.) random variables X_1, X_2, \dots, X_n . Such sequence will be labeled as a *random sample* of size n . Suppose that the mean and variance of the underlying probability distribution are μ and σ^2 , respectively. Otherwise, the probability distribution is unknown. We shall find the mean and variance of the sample mean of the random sample.

It is easy to see that

$$\begin{aligned}\mu_{\bar{x}} &= \frac{E(X_1 + X_2 + \dots + X_n)}{n} = \frac{E(X_1) + E(X_2) + \dots + E(X_n)}{n} \\ &= \frac{\mu + \mu + \dots + \mu}{n} = \mu.\end{aligned}$$

In this derivation, we have not used independence or the fact that all the r.v.'s have the same distribution, only the fact that they all have the same (finite) mean. We say that \bar{X} is an *unbiased* estimator of μ .

Next we shall derive the variance of \bar{X} :

$$\begin{aligned}\sigma_{\bar{x}}^2 &= E[(\bar{X} - \mu)^2] \\ &= E\left[\left(\frac{(X_1 - \mu)}{n} + \frac{(X_2 - \mu)}{n} + \dots + \frac{(X_n - \mu)}{n}\right)^2\right] \\ &= \sum_{i=1}^n \frac{E[(X_i - \mu)^2]}{n^2} + \text{terms like } E\left[\frac{(X_1 - \mu)(X_2 - \mu)}{n^2}\right].\end{aligned}$$

Now, by independence, the expectation of the cross-product terms is zero:

$$\begin{aligned}E[(X_1 - \mu)(X_2 - \mu)] &= \int_{-\infty}^{\infty} (x_1 - \mu)(x_2 - \mu) f(x_1) f(x_2) dx_1 dx_2 \\ &= E(X_1 - \mu) E(X_2 - \mu) = 0\end{aligned}$$

(the argument for discrete distributions is analogous). Thus, we have

$$\sigma_{\bar{x}}^2 = \frac{\sigma^2}{n}.$$

We note that in the above derivation the fact that the X_i 's are identically distributed has been superfluous. Only the facts that the random variables are independent and have the same μ and σ^2 have been needed. The property that the variability of \bar{X} about the true mean μ decreases as n increases is of key importance in experimental science. We shall develop this notion further below.

Let us begin by stating the celebrated *Chebyshev's inequality*. If Y is any random variable with mean μ_y and variance σ_y^2 , then for any $\varepsilon > 0$

$$P(|Y - \mu_y| > \varepsilon) \leq \frac{\sigma_y^2}{\varepsilon^2}. \quad (\text{A.65})$$

As a practical approximation device, it is not a particularly useful inequality. However, as an asymptotic device, it is invaluable. Let us consider the case where $Y = \bar{X}$. Then, we have

$$P(|\bar{X} - \mu| > \varepsilon) \leq \frac{\sigma^2}{n\varepsilon^2}, \quad (\text{A.66})$$

or equivalently

$$P(|\bar{X} - \mu| \leq \varepsilon) > 1 - \frac{\sigma^2}{n\varepsilon^2}. \quad (\text{A.67})$$

Equation (A.67) is a form of the *weak law of large numbers*. The WLLN tells us that if we are willing to take a sufficiently large sample, then we can obtain an arbitrarily large probability that \bar{X} will be arbitrarily close to μ .³

³In fact, even a more powerful result, the *strong law of large numbers*, is available. In order to make the difference between the WLLN and SLLN more transparent, let us denote the sample mean based on a sample of size n by \bar{X}_n , so that the dependence of \bar{X} on n be emphasized. Now we can write the WLLN in the following way

$$\lim_{n \rightarrow \infty} P(|\bar{X}_n - \mu| \leq \varepsilon) = 1 \quad (\text{A.68})$$

for each positive ε . On the other hand, the SLLN states that

$$P(\lim_{n \rightarrow \infty} |\bar{X}_n - \mu| = 0) = 1. \quad (\text{A.69})$$

Loosely speaking, in the WLLN, the probability of \bar{X}_n being close to μ for only one n at a time is claimed, whereas in the SLLN, the closeness of \bar{X}_n to μ for all large n simultaneously is asserted with probability one. The rather practical advantage of the SLLN is that if $g(x)$ is some function, then

$$P(\lim_{n \rightarrow \infty} |g(\bar{X}_n) - g(\mu)| = 0) = 1. \quad (\text{A.70})$$

The WLLN and the SLLN are particular cases of convergence in probability and almost sure convergence of a sequence of r.v.'s, respectively. Let $Y_1, Y_2, \dots, Y_n, \dots$ be an infinite sequence of r.v.'s. We say that this sequence of r.v.'s converges *in probability* or *stochastically* to a random variable Y if

$$\lim_{n \rightarrow \infty} P(|Y_n - Y| > \varepsilon) = 0$$

for each positive ε . We say that the sequence $Y_1, Y_2, \dots, Y_n, \dots$ converges *almost surely* or converges *with probability one* if

$$P(\lim_{n \rightarrow \infty} |Y_n - Y| = 0) = 1.$$

A.12 The Multivariate Normal Distribution

The random vector \mathbf{X} of dimension p is said to have *multivariate normal* (or p -dimensional *multinormal* or p -variate *normal*) distribution if its p.d.f. is given by

$$f(\mathbf{x}) = |2\pi\mathbf{\Sigma}|^{-1/2} \exp\left\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})' \mathbf{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})\right\}, \quad (\text{A.71})$$

where $\boldsymbol{\mu}$ is a constant vector and $\mathbf{\Sigma}$ is a constant positive definite matrix. It can be shown that $\boldsymbol{\mu}$ and $\mathbf{\Sigma}$ are the mean vector and covariance matrix of the random vector \mathbf{X} , respectively. For short, we write that \mathbf{X} is $\mathcal{N}(\boldsymbol{\mu}, \mathbf{\Sigma})$ distributed. Now if the covariance matrix $\mathbf{\Sigma}$ is diagonal, $\mathbf{\Sigma} = \text{diag}(\sigma_{11}, \sigma_{22}, \dots, \sigma_{pp})$, the density function can be written as:

$$f(\mathbf{x}) = \prod_{i=1}^p (2\pi\sigma_{ii})^{-1/2} \exp\left\{-\frac{1}{2}(x_i - \mu_i)\sigma_{ii}^{-1}(x_i - \mu_i)\right\}. \quad (\text{A.72})$$

Thus, the elements of \mathbf{X} are then mutually independent normal random variables with means μ_i and variances σ_{ii} , $i = 1, 2, \dots, p$, respectively. If the random vector \mathbf{X} is multivariate normal, then the property that its elements are uncorrelated one with another (i.e., that $\text{Cov}(X_i, X_j) = 0$, $i \neq j$) implies their mutual independence.

A.13 The Wiener Process

A *stochastic process* $\{X_t\}$ is a collection of random variables indexed on the real variable t . Typically, t is time. Let us suppose the $\{X_t\}$ process has the property that for any collection of t values $t_1 < t_2 < \dots < t_n$, the vector random variable $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ is an n dimensional normal distribution. Then $\{X_t\}$ is a *Gaussian process*.

Next, suppose that a stochastic process $W(t)$ has the following properties

- $W(0) = 0$;
- For any t , $W(t)$ is normal with mean zero and variance t ;
- If the intervals $[t_1, t_2]$ and $[t_3, t_4]$ do not overlap, then the random variables $W(t_2) - W(t_1)$ and $W(t_4) - W(t_3)$ are stochastically independent.

Then $W(t)$ is called a *Wiener process*.

We define a *Brownian process* $S(t)$ as

$$S(t) = \mu t + \sigma W(t), \quad (\text{A.73})$$

where $W(t)$ is a Wiener process. We write this as the stochastic differential equation

$$dS(t) = \mu t + \sigma dW(t). \quad (\text{A.74})$$

If the logarithm of $S(t)$ is a Brownian process, then we say that $S(t)$ is a *geometric Brownian process*. We may write this as the stochastic differential equation

$$\frac{dS(t)}{Sdt} = \mu t + \sigma dW(t). \quad (\text{A.75})$$

Then, we have

$$\begin{aligned} S(t) &= S(0) \exp[(\mu - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}Z] \\ &= S(0) \exp(\mu t) \times \exp(\sigma\sqrt{t}Z - (\frac{1}{2}\sigma^2t)) \end{aligned} \quad (\text{A.76})$$

where the Z are independently distributed as $\mathcal{N}(0, 1)$. To find the percentile values of a security at time t , we use the percentile values from

Z_{critical}	$P(Z > Z_{\text{critical}}) = P(Z < -Z_{\text{critical}})$
2.3263	.01
1.6449	.05
1.2816	.10
1.0364	.15
.8416	.20
.6745	.25
.5244	.30
.2533	.40
0	.50

A.14 The Poisson Process and the Poisson Distribution

Let us consider a counting process described by Poisson's Four Axioms:

1. $P(1 \text{ occurrence in } [t, t + \epsilon]) = \lambda\epsilon;$
2. $P(\text{ more than 1 occurrence in } [t, t + \epsilon]) = o(\epsilon)$
where $\lim_{\epsilon \rightarrow 0} o(\epsilon)/\epsilon = 0$;
3. $P(k \text{ in } [t_1, t_2] \text{ and } m \text{ in } [t_3, t_4]) = P(k \text{ in } [t_1, t_2])P(m \text{ in } [t_3, t_4])$
if $[t_1, t_2] \cap [t_3, t_4] = \emptyset$;
4. $P(k \text{ in } [t_1, t_1 + s]) = P(k \text{ in } [t_2, t_2 + s])$ for all t_1, t_2 , and s .

Then we may write

$$\begin{aligned} P(k + 1 \text{ in } [0, t + \epsilon]) &= P(k + 1 \text{ in } [0, t])P(0 \text{ in } [t, t + \epsilon]) \\ &\quad + P(k \text{ in } [0, t])P(1 \text{ in } [t, t + \epsilon]) + o(\epsilon) \\ &= P(k + 1, t)(1 - \lambda\epsilon) + P(k, t)\lambda\epsilon + o(\epsilon) \end{aligned}$$

where $P(k, t) = P(k \text{ in } [0, t])$.

Then we have

$$\frac{P(k+1, t+\epsilon) - P(k+1, t)}{\epsilon} = \lambda[P(k, t) - P(k+1)] + \frac{o(\epsilon)}{\epsilon}. \quad (\text{A.77})$$

Taking the limit as $\epsilon \rightarrow 0$, we have the differential-difference equation:

$$\frac{dP(k+1, t)}{dt} = \lambda[P(k, t) - P(k+1, t)]. \quad (\text{A.78})$$

Taking $k = -1$, since we know that it is impossible for a negative number of events to occur, we have

$$\frac{dP(0, t)}{dt} = -\lambda P(0, t). \quad (\text{A.79})$$

So,

$$P(0, t) = \exp(-\lambda t). \quad (\text{A.80})$$

Next, we for for $P(1, t)$ via

$$\frac{P(1, t)}{dt} = \lambda[\exp(-\lambda t) - P(1, t)], \quad (\text{A.81})$$

with solution,

$$P(1, t) = \exp(-\lambda t)(\lambda t). \quad (\text{A.82})$$

Continuing in this fashion, we quickly conjecture that the general solution of (A.78) is given by

$$P(k, t) = \frac{e^{-\lambda t}(\lambda t)^k}{k!}. \quad (\text{A.83})$$

(A.83) defines the *Poisson distribution*. We can quickly compute the mean and variance both to be λt . If there are N happenings in a time interval of length T , then a natural estimate for λ is found by solving

$$\hat{\lambda}T = N. \quad (\text{A.84})$$

This gives us

$$\hat{\lambda} = \frac{N}{T}. \quad (\text{A.85})$$

Of special interest to us is the probability that no shock (event) occurs in the interval from t to $t+s$. Clearly, that is given by $P(0, s) = \exp(-\lambda s)$. This immediately enables us to write the cumulative distribution function of the time it takes to reach an event, namely,

$$F(t) = 1 - \exp(-\lambda t) \quad (\text{A.86})$$

Now, we know that the cdf of a continuous random variable is distributed as a uniform random variable on the interval from 0 to 1. This gives us a ready means for generating a simulated time until first occurrence of an event.

1. Generate u from $U(0, 1)$.
2. Set $u = 1 - \exp(-\lambda t)$.
3. $t = \log(1 - u)/\lambda$.

Once we observe an occurrence time t_1 , we start over with the new time origin set at t_1 .

Note that we now have a stochastic process. At any given time t , the probability an event happens in the interval $(t, t + \epsilon)$ is given by $\lambda \times \epsilon$. The probability is stochastically independent of the history prior to t .

A.14.1 Simulating Bear Jumps

Let us suppose we have a bank account of one million pesos which is growing at the rate of 20% per year. Unfortunately, a random devaluation of currency in the amount of 10% occurs on the average of once a year. A random devaluation of currency in the amount of 20% occurs on the average of once every five years. What will be the value of the bank account six months in the future in present day pesos? If there are no devaluations, the answer is

$$S(10) = 1,000,000 \exp(.5 \times .2) = 1,105,171.$$

On the other hand, we have to deal with the concatenation of our “sure thing” bank account with the concatenation of two Poisson bear jump mechanisms. The 10% jumps have $\lambda_{10\%}$ given by:

$$\lambda_{10\%} = \frac{1}{1} = 1. \quad (\text{A.87})$$

The 20% jumps have $\lambda_{20\%}$ given by

$$\lambda_{20\%} = \frac{1}{5} = 0.2. \quad (\text{A.88})$$

Returning to (A.84), to handle the 10% jumps, we generate a uniform random variate u_1 on $[0,1]$ we have the following downjump multiplier table:

1. If $u_1 < \exp(-1 \times 0.5) = 0.60653$, use multiplier 1.00.
2. If $0.60653 \leq u_1 < 0.90980$, use multiplier 0.9.
3. If $0.909080 \leq u_1 < 0.98561$, use multiplier 0.81.
4. If $0.98561 \leq u_1 < 0.99824$, use multiplier 0.729.
5. If $0.99824 \leq u_1 < 0.99982$, use multiplier 0.6561.
6. If $0.99824 \leq u_1$, use multiplier 0.59049.

To handle the 20% bear jumps, we generate a uniform random variate u_2 . We then have the following downjump multiplier table.

1. If $u_1 < \exp(-.2 \times 0.5) = 0.90484$, use multiplier 1.00
2. If $0.90484 \leq u_2 < 0.99532$, use multiplier 0.8
3. If $0.99532 \leq u_2 < 0.99985$, use multiplier 0.64.
4. If $0.99985 \leq u_2$, use multiplier 0.512.

Many standard software packages have automatic Poisson generators. To use such a routine, one simply enters λT and a random number of bear jumps, from 0 to infinity, is generated.

A.15 Parametric Simulation

For the standard lognormal model for stock growth, we have

$$S(t) = S(0) \exp[(\mu - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}Z]. \quad (\text{A.89})$$

Then, from (A.90), we have, for all t and Δt

$$r(t + \Delta t, t) = \frac{S(t + \Delta t)}{S(t)} = \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) \Delta t + Z\sigma\sqrt{\Delta t} \right]. \quad (\text{A.90})$$

Defining $R(t + \Delta t, t) = \log(r(t + \Delta t, t))$, we have

$$R(t + \Delta t, t) = \left(\mu - \frac{\sigma^2}{2} \right) \Delta t + \epsilon\sigma\sqrt{\Delta t}.$$

Then

$$E[R(t + \Delta t, t)] = \left(\mu - \frac{\sigma^2}{2} \right) \Delta t. \quad (\text{A.91})$$

We will take μ s to be given on an annual basis. Then, if the data are taken at N points separated by Δt , let the sample mean \bar{R} be defined by

$$\bar{R} = \frac{1}{N} \sum_{i=1}^N R(i) \quad (\text{A.92})$$

By the strong law of large numbers, the sample mean \bar{R} converges almost surely to its expectation $(\mu - \sigma^2/2)\Delta t$. Next, we note that

$$[R(t + \Delta t, t) - E(R(t + \Delta t, t))]^2 = \epsilon^2\sigma^2\Delta t, \quad (\text{A.93})$$

so

$$\text{Var}[R(t + \Delta t, t)] = E[R(t + \Delta t, t) - \left(\mu - \frac{\sigma^2}{2} \right) \Delta t]^2 = \sigma^2\Delta t. \quad (\text{A.94})$$

$$s_R^2 = \frac{1}{N-1} \sum_{i=1}^N (R(i) - \bar{R})^2. \quad (\text{A.95})$$

The most utilized estimation technique in statistics is the *method of moments*. By this procedure, we replace the mean by the sample mean, the variance by the sample variance, etc. Doing this in (A.94), we have

$$\hat{\sigma}^2 = \frac{s_R^2}{\Delta t}. \quad (\text{A.96})$$

Then, from (A.92) we have

$$\hat{\mu} = \frac{\bar{R}}{\Delta t} + \frac{\hat{\sigma}^2}{2}. \quad (\text{A.97})$$

Having estimated from historical data μ and σ , we can now simulate the value of our stock using (A.90) for any desired time horizon.

A.15.1 Simulating a Geometric Brownian Walk

First we start with a simple geometric Brownian case. The program is incredibly simple. We need only assume available a uniform generator, a normal generator, and a sorting routine.

Simulation of Portfolios

1. Enter $S(0), T, \hat{\mu}$, and $\hat{\sigma}$.
2. Repeat 10,000 times.
3. Generate normal observation Z with mean zero and variance 1.
4. Set $S(T) = S(0) \exp[(\mu - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}Z]$.
5. End repeat.
6. Sort the 10,000 end values.
7. Then $F(v) = \frac{\text{Number of sorted values} \leq v}{10,000}$.

We will now add on the possibility of two types of bear jumps: a 10% downturn on the average of once every year ($\lambda_1 = 1$) and a 20% downturn on the average of once every five years ($\lambda_2 = .2$).

Simulation With Jumps

1. Enter $S(0), T, \hat{\mu}, \hat{\sigma}, \lambda_1, \lambda_2$.
2. Repeat 10,000 times.

3. Generate normal observation Z with mean zero and variance 1.
4. Set $S(T) = S(0) \exp[(\mu - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}Z]$.
5. Generate a Poisson variate m_1 from $Po(\lambda_1 \times T)$.
6. Replace $S(T)$ by $.9^{m_1} \times S(T)$
7. Generate a Poisson variate m_2 from $Po(\lambda_2 \times T)$.
8. Replace $S(T)$ by $.8^{m_2} \times S(T)$
9. End repeat.
10. Sort the 10,000 end values.
11. Then $F(v) = \frac{\text{Number of sorted values} \leq v}{10,000}$.

A.15.2 The Multivariate Case

Now, we assume that we have p securities to be considered in a portfolio with weights $\{c_j\}$ which are non-negative and sum to one. We estimate the $\{\mu_j\}$ and the $\{\sigma\}$ precisely as in the one security case. Accordingly, for the j th security, we let

$$R_j(i) = \log \left(\frac{S_j(i)}{S_j(i-1)} \right).$$

Then

$$\bar{R}_j = \frac{1}{N} \sum_{i=1}^N R_j(i)$$

and

$$s_{R_j}^2 = \frac{1}{N-1} \sum_{i=1}^N (R_j(i) - \bar{R}_j)^2.$$

So we have an estimates for σ_j^2 and μ_j , namely

$$\hat{\sigma}_j^2 = \frac{s_{R_j}^2}{\Delta t} \quad (\text{A.98})$$

and

$$\hat{\mu}_j = \frac{\bar{R}_j}{\Delta t} + \frac{\hat{\sigma}_j^2}{2}. \quad (\text{A.99})$$

Now, we must also account for the correlation between the growths of stocks:

$$\hat{\sigma}_{jm} = \frac{1}{N-1} \sum_{i=1}^N (R_j(i) - \bar{R}_j)(R_m(i) - \bar{R}_m). \quad (\text{A.100})$$

Then, we have as our estimated covariance matrix

$$\hat{\Sigma} = \begin{pmatrix} \hat{\sigma}_{11} & \hat{\sigma}_{12} & \dots & \hat{\sigma}_{1p} \\ \hat{\sigma}_{12} & \hat{\sigma}_{22} & \dots & \hat{\sigma}_{2p} \\ \dots & \dots & \dots & \dots \\ \hat{\sigma}_{1p} & \hat{\sigma}_{2p} & \dots & \hat{\sigma}_{pp} \end{pmatrix}, \quad (\text{A.101})$$

For a covariance matrix, we may obtain a *Cholesky decomposition*

$$\hat{\Sigma} = \mathbf{L} \mathbf{L}^T, \quad (\text{A.102})$$

where \mathbf{L} is a lower triangular matrix, frequently referred to as the matrix square root of $\hat{\Sigma}$. Subroutines for obtaining the matrix square root are available in most standard matrix compatible software such as Matlab, Splus, and SAS.

Let us generate p independent normal variates with mean 0 and variance one, putting them into a row vector, $\mathbf{Z} = (z_1, z_2, \dots, z_p)$. Then, we compute the row vector

$$\mathbf{V} = \mathbf{Z} \mathbf{L}^T = (v_1, v_2, \dots, v_p).$$

Then, the price of the j th stock in the joint simulation at time T is given by

$$S_j(t) = S_j(0) \exp[(\mu_j - \frac{1}{2}\sigma_j^2)T + v_j\sqrt{T}]. \quad (\text{A.103})$$

Then, for the p -dimensional case:

Portfolio Simulation

1. Enter $\{S_j(0)\}_{j=1}^p$, T , $\{\hat{\mu}_j\}$, and \mathbf{L} .
2. Repeat 10,000 times.
3. Generate p independent normal variates with mean 0 and variance one, putting them into a row vector, $\mathbf{Z} = (z_1, z_2, \dots, z_p)$.
4. Compute the row vector

$$\mathbf{V} = \mathbf{Z} \mathbf{L}^T = (v_1, v_2, \dots, v_p).$$

5. For each of the p stocks, compute

$$S_j(t) = S_j(0) \exp[(\mu_j - \frac{1}{2}\sigma_j^2)T + v_j\sqrt{T}].$$

6. For the i th repeat save the $(S_1(T), S_2(T), \dots, S_p(T))$ as the row vector \mathbf{S}_i .
7. End repeat.

The multivariate simulation with the two bear jump processes added becomes:

Simulation Multivariate With Jumps

1. Enter $\{S_j(0)\}_{j=1}^p$, T , $\{\hat{\mu}_j\}$, \mathbf{L} , λ_1 and λ_2 .
2. Repeat 10,000 times.
3. Generate a Poisson variate m_1 from $Po(\lambda_1 \times T)$.
4. Replace $S_j(0)$ by $.9^{m_1} \times S_j(T)$
5. Generate a Poisson variate m_2 from $Po(\lambda_2 \times T)$.
6. Replace $S_j(0)$ by $.8^{m_2} \times S_j(T)$
7. Generate p independent normal variates with mean 0 and variance one, putting them into a row vector, $\mathbf{Z} = (z_1, z_2, \dots, z_p)$.
8. Compute the row vector

$$\mathbf{V} = \mathbf{Z}\mathbf{L}^T = (v_1, v_2, \dots, v_p).$$

9. For each of the p stocks, compute

$$S_j(T) = S_j(0) \exp[(\mu_j - \frac{1}{2}\sigma_j^2)T + v_j\sqrt{T}].$$

10. For the i th repeat save the $(S_1(T), S_2(T), \dots, S_p(T))$ as the row vector $\mathbf{S}_i(T)$.
11. End repeat.

Suppose we have a portfolio consisting of p stocks each weighted by $c_j \geq 0$ such that $\sum c_j = 1$. Then to obtain a simulation of the portfolio, we must look at the 10,000 values from the above simulations, of the form

$$P_i(T) = \sum_{j=1}^p c_j S_{i,j}(T). \quad (\text{A.104})$$

Here, the i refers to the number of the simulation. We can then form the simulation of the portfolio results by sorting the $P_i(T)$ and obtaining the cumulative distribution function. This gives us

A Simulation Based Portfolio

1. Enter $S(T)$, $\{c_j\}$.
2. For all i from 1 to 10,000, find $P_i(T) = \sum_{j=1}^p c_j S_{i,j}(T)$.
3. Sort the $P_i(T)$.
4. $F(v) = \frac{\text{Number of sorted values } P_i(t) \leq v}{10,000}$.

The *Portfolio Simulation* may then be used for a host of purposes. We might use it as an assist, for example, in deciding whether one wished to replace one stock in the portfolio by another.

We wish to make it clear that, due to the fact that one is using historical estimates of growth and volatility rather than the actual values, rather than the true ones (which we cannot know), the portfolio optimization based on such simulations should only be used as an exploratory and speculative tool. It is not a magic bullet, nor do we claim it to be. It is a useful technique for what might happen.

A.16 Resampling Simulation

Let us suppose we have a data base showing the year to year change in a stock or a stock index. We can then obtain a data base of terms like

$$R_i = \log\left(\frac{S(t_i)}{S(t_{i-1})}\right).$$

In other words, we know that

$$S(t_i) = S(t_{i-1}) \times \exp(R_i).$$

Suppose we have a data base of n such terms, $\{R_1, R_2, \dots, R_n\}$. Let us make the (frequently reasonable) assumption that the ups and downs of the stock or the index in the past are a good guide to the ups and downs in the future. It would not be a good idea, if we wished to forecast the value of the stock five years in advance, randomly to sample (with replacement) five of the R_i 's, say, $\{R_3, R_{17}, R_{20}, R_{20}, R_{31}\}$ and use

$$\hat{S}(5) = S(0) \times \exp[R_3 + R_{17} + R_{20} + R_{20} + R_{31}].$$

On the other hand, if we wished to obtain, not a point estimate for $S(T)$, but an estimate for the distribution of possible values of $S(5)$, experience shows that this frequently can be done as follows:

Portfolio Resampling

1. Enter $S(0)$, T , and the $\{R_i\}$.
2. Repeat 10,000 times
3. For pass i , randomly sample with replacement T values from $\{R_1, R_2, \dots, R_n\}$, say, $\{R_{i1}, R_{i2}, R_{i3}, R_{i4}, \dots\}$.
4. Compute

$$SS(T) = S(0) \times \exp[R_{i1} + R_{i2} + R_{i3} + R_{i4} + \dots]$$

Clearly we use T resampled R values to obtain.

5. Obtain the empirical cumulative distribution function from the resulting 10,000 values of $SS(T)$. That is, compute

$$F_T(v) = \frac{\text{Number of sorted values } \{SS(T)\} \leq v}{10,000}. \quad (\text{A.105})$$

By looking at the simulations for, say, five, ten, twenty and forty years in the future, an investor can examine the historically based outcomes of buying a security in the light of his/her anticipated needs.

A.16.1 The Multivariate Case

Here, for each of p stocks, we compute

$$R_{i,j} = \log\left(\frac{S(t_{i,j})}{S(t_{i,j-1})}\right).$$

Multivariate Portfolio Resampling

We then proceed very much as in the parametric case:

1. Enter $\{S_j(0)\}_{j=1}^k$, T , and the $\{R_{i,j}\}_{i,j}$.
2. Repeat 10,000 times
3. For pass i , randomly sample with replacement T values from the length of the historical list, say, $l_{i,1}, l_{i,2}, \dots, l_{i,T}$.
4. For each stock j

$$SS_{i,j}(T) = S_j(0) \times \exp[R_{l_{(i,1)},j} + R_{l_{(i,2)},j} + \dots + R_{l_{(i,T)},j}]$$

5. Store $(SS_{i,1}, SS_{i,2}, \dots, SS_{i,j})$ as a row vector \mathbf{SS}_i .
6. End Repeat.

Now for a portfolio of p stocks, balanced according to

$$P(t) = \sum_{i=1}^p c_i S_i(t), \quad (\text{A.106})$$

where the weights are non-negative and sum to one, we simply use the,

An Algorithm for Resampling Portfolios

1. Enter $\mathbf{SS}(T)$, $\{c_j\}$.
2. For all i from 1 to 10,000, find $P_i(T) = \sum_{j=1}^p c_j S_{i,j}(T)$.
3. Sort the $P_i(T)$.
4. $F(v) = \frac{\text{Number of } P_i(T) \leq v}{10,000}$.

We may then proceed to obtain simulations of the portfolio value at a given time. Because of the correlations between stock values, it is essential that, when we randomly select a year, we sample the annual growth factors of the stocks in the portfolio for that year.

A.17 A Portfolio Case Study

Next, we take the 90 stocks in the S&P 100 that were in business prior to 1991. The data base we shall use will utilize the 12 years 1990–2001, utilizing monthly data, both for the estimation of the parameters characterizing the simple geometric Brownian model parameters (without Poissonian jumps), namely the $\{\mu_j\}_{j=1}^{j=80}$ and the covariance matrix $\{\sigma_{i,j}\}$, and for obtaining resampling months. We look below at the result of one of many possible optimization criteria that might have been considered. We find the allocation of an investment in the portfolio amongst the 90 stocks maximizing the one year lower 20 percentile, with the constraint that no stock has more than 5% of the portfolio share.

**Table A.2. Portfolio Allocation from S&P 100.
Maximizing One Year 20 Percentile
with Max 5% in Any Stock.**

id	permno	ticker	u	σ	par alloc	npar alloc
1	10104	ORCL	0.44	0.65	0.05	0.05
2	10107	MSFT	0.39	0.48	0.05	0.05
3	10145	HON	0.12	0.44	0.00	0.00
4	10147	EMC	0.48	0.61	0.00	0.01
5	10401	T	0.05	0.40	0.00	0.00
6	10890	UIS	0.36	0.68	0.00	0.01
7	11308	KO	0.07	0.31	0.00	0.00
8	11703	DD	0.05	0.28	0.00	0.00
9	11754	EK	-0.11	0.35	0.00	0.00
10	11850	XOM	0.12	0.17	0.00	0.00
11	12052	GD	0.19	0.25	0.05	0.05
12	12060	GE	0.23	0.26	0.00	0.00
13	12079	GM	0.08	0.36	0.00	0.00
14	12490	IBM	0.32	0.34	0.05	0.00
15	13100	MAY	0.07	0.29	0.00	0.00
16	13856	PEP	0.13	0.28	0.00	0.00
17	13901	MO	0.13	0.32	0.00	0.00
18	14008	AMGN	0.32	0.39	0.05	0.05
19	14277	SLB	0.12	0.36	0.00	0.00
20	14322	S	0.05	0.36	0.00	0.00
21	15560	RSH	0.27	0.49	0.05	0.05
22	15579	TXN	0.40	0.56	0.00	0.01
23	16424	G	0.09	0.33	0.00	0.00
24	17830	UTX	0.21	0.35	0.00	0.00
25	18163	PG	0.16	0.31	0.04	0.00
26	18382	PHA	0.12	0.30	0.00	0.00
27	18411	SO	0.14	0.24	0.05	0.05
28	18729	CL	0.25	0.33	0.05	0.05
29	19393	BMY	0.20	0.26	0.01	0.03
30	19561	BA	0.05	0.35	0.00	0.00
31	20220	BDK	0.06	0.38	0.00	0.00
32	20626	DOW	0.07	0.30	0.00	0.00
33	21573	IP	0.07	0.36	0.00	0.00
34	21776	EXC	0.17	0.33	0.05	0.05
35	21936	PFE	0.26	0.28	0.05	0.05

Appendix B

Statistical Tables

Tables of the Normal Distribution

z	Values of $\frac{1}{\sqrt{2\pi}} \int_{-\infty}^z e^{-\frac{t^2}{2}} dt$									
	.0	.01	.02	.03	.04	.05	.06	.07	.08	.09
.0	.50000	.50399	.50798	.51197	.51595	.51994	.52392	.52790	.53188	.53586
.1	.53983	.54380	.54776	.55172	.55567	.55962	.56356	.56749	.57142	.57535
.2	.57926	.58317	.58706	.59095	.59483	.59871	.60257	.60642	.61026	.61409
.3	.61791	.62172	.62552	.62930	.63307	.63683	.64058	.64431	.64803	.65173
.4	.65542	.65910	.66276	.66640	.67003	.67364	.67724	.68082	.68439	.68793
.5	.69146	.69497	.69847	.70194	.70540	.70884	.71226	.71566	.71904	.72240
.6	.72675	.72907	.73237	.73565	.73891	.74215	.74537	.74857	.75175	.75490
.7	.75804	.76115	.76424	.76730	.77036	.77337	.77637	.77935	.78230	.78524
.8	.78814	.79103	.79389	.79673	.79955	.80234	.80511	.80785	.81057	.81327
.9	.81594	.81859	.82121	.82381	.82639	.82894	.83147	.83398	.83646	.83891
1.0	.84134	.84375	.84614	.84849	.85083	.85314	.85543	.85769	.85993	.86214
1.1	.86433	.86650	.86864	.87076	.87286	.87493	.87698	.87900	.88100	.88298
1.2	.88493	.88686	.88877	.89065	.89251	.89435	.89617	.89796	.89973	.90147
1.3	.90320	.90490	.90658	.90824	.90988	.91149	.91309	.91466	.91621	.91774
1.4	.91924	.92073	.92220	.92364	.92507	.92647	.92785	.92922	.93056	.93189
1.5	.93319	.93448	.93574	.93699	.93822	.93943	.94062	.94179	.94295	.94408
1.6	.94520	.94630	.94738	.94845	.94950	.95053	.95154	.95254	.95352	.95449
1.7	.95543	.95637	.95728	.95818	.95907	.95994	.96080	.96164	.96246	.96327
1.8	.96407	.96485	.96562	.96638	.96712	.96784	.96858	.96928	.96995	.97062
1.9	.97128	.97193	.97257	.97320	.97381	.97441	.97500	.97558	.97615	.97670
2.0	.97725	.97778	.97831	.97882	.97932	.97982	.98030	.98077	.98124	.98169
2.1	.98214	.98257	.98300	.98341	.98382	.98422	.98461	.98500	.98537	.98574
2.2	.98610	.98645	.98679	.98713	.98745	.98778	.98809	.98840	.98870	.98899
2.3	.98928	.98956	.98983	.99010	.99036	.99061	.99086	.99111	.99134	.99158
2.4	.99180	.99202	.99224	.99245	.99266	.99286	.99305	.99324	.99343	.99361
2.5	.99379	.99396	.99413	.99430	.99446	.99461	.99477	.99492	.99508	.99520
2.6	.99534	.99547	.99560	.99573	.99585	.99598	.99609	.99621	.99632	.99643
2.7	.99653	.99664	.99674	.99683	.99693	.99702	.99711	.99720	.99728	.99736
2.8	.99744	.99752	.99760	.99767	.99774	.99781	.99788	.99795	.99801	.99807
2.9	.99813	.99819	.99825	.99831	.99836	.99841	.99846	.99851	.99856	.99861
3.0	.99865	.99869	.99874	.99878	.99882	.99886	.99889	.99893	.99898	.99900
3.1	.99903	.99906	.99910	.99913	.99916	.99918	.99921	.99924	.99926	.99929
3.2	.99931	.99934	.99936	.99938	.99940	.99942	.99944	.99946	.99948	.99950
3.3	.99952	.99953	.99955	.99957	.99958	.99960	.99961	.99962	.99964	.99965
3.4	.99966	.99968	.99969	.99970	.99971	.99972	.99973	.99974	.99975	.99976
3.5	.99977	.99978	.99978	.99979	.99980	.99981	.99981	.99982	.99983	.99983
3.6	.99984	.99985	.99985	.99986	.99986	.99987	.99987	.99988	.99988	.99989
3.7	.99989	.99990	.99990	.99990	.99991	.99991	.99992	.99992	.99992	.99992
3.8	.99993	.99993	.99993	.99994	.99994	.99994	.99994	.99995	.99995	.99995
3.9	.99995	.99995	.99996	.99996	.99996	.99996	.99996	.99996	.99997	.99997

⁰Empirical Model Building: Data, Modeling, and Reality, Second Edition. James R. Thompson ©2011 John Wiley & Sons, Inc. Published 2011 by John Wiley & Sons.

Tables of the Chi-Square Distribution

Critical Values of $P = 1 - \frac{1}{2^{\nu/2}\Gamma(\nu/2)} \int_0^{\chi^2} x^{\nu/2-1} e^{-x/2} dx$

ν	0.100	0.250	0.500	0.750	0.900	0.950	0.975	0.990	0.995	0.998	0.999
1	0.016	0.102	0.455	1.323	2.706	3.841	5.024	6.635	7.879	9.550	10.828
2	0.211	0.575	1.986	2.773	4.605	5.991	7.378	9.210	10.597	12.429	13.816
3	0.584	1.213	2.966	4.108	6.251	7.815	9.348	11.345	12.888	14.796	16.266
4	1.064	1.923	3.957	5.985	7.779	9.488	11.143	13.277	14.860	16.924	18.467
5	1.610	2.675	4.351	6.626	9.236	11.070	12.893	15.086	16.750	18.907	20.515
6	2.204	3.455	5.348	7.841	10.645	12.592	14.449	16.812	18.548	20.791	22.458
7	2.833	4.255	6.346	9.037	12.017	14.067	16.013	18.475	20.278	22.601	24.322
8	3.490	5.071	7.344	10.219	13.362	15.507	17.535	20.090	21.955	24.352	26.124
9	4.168	5.899	8.343	11.389	14.684	16.919	19.023	21.666	23.589	26.056	27.877
10	4.865	6.737	9.342	12.549	15.987	18.307	20.483	23.209	25.188	27.722	29.588
11	5.578	7.584	10.341	13.701	17.275	19.675	21.920	24.725	26.757	29.954	31.264
12	6.304	8.438	11.340	14.845	18.549	21.026	23.397	26.217	28.300	30.957	32.900
13	7.042	9.299	12.340	15.984	19.812	22.362	24.736	27.688	29.819	32.538	34.528
14	7.790	10.185	13.399	17.117	21.064	23.685	26.119	29.141	31.319	34.091	36.123
15	8.547	11.087	14.339	18.245	22.307	24.996	27.488	30.578	32.801	35.628	37.697
16	9.312	11.912	15.338	19.369	23.542	26.298	28.845	32	34.287	37.146	39.252
17	10.085	12.792	16.338	20.489	24.769	27.587	30.191	33.409	35.718	38.648	40.700
18	10.865	13.675	17.338	21.005	26.989	28.869	31.526	34.805	37.156	40.136	42.312
19	11.651	14.562	18.338	22.718	27.204	30.144	32.852	36.191	38.582	41.610	43.820
20	12.443	15.452	19.337	23.828	28.412	31.410	34.170	37.566	39.997	43.072	45.315
21	13.240	16.344	20.397	24.935	29.615	32.671	35.479	38.932	41.401	44.522	46.797
22	14.041	17.240	21.397	26.039	30.813	33.924	36.781	40.289	42.796	45.962	48.268
23	14.848	18.137	22.397	27.141	32.007	35.172	38.076	41.638	44.181	47.391	49.728
24	15.659	19.037	23.397	28.241	33.196	36.415	39.364	42.980	45.559	48.812	51.179
25	16.473	19.939	24.397	29.399	34.382	37.652	40.646	44.314	46.928	50.223	52.620
26	17.292	20.843	26.396	30.455	36.563	38.885	41.923	45.642	48.290	51.627	54.052
27	18.114	21.749	26.356	31.528	36.741	40.113	43.195	46.963	49.645	53.023	55.476
28	18.939	22.657	27.356	32.620	37.916	41.397	44.461	48.278	50.993	54.411	56.892
29	19.768	23.567	28.356	33.711	39.087	42.557	45.722	49.588	52.336	55.792	58.301
30	20.599	24.478	29.356	34.800	40.256	43.773	46.979	50.892	53.672	57.167	59.703

Index

Adjuvant chemotherapy, 30–36
AIDS, 80–108
Alamo, battle of, 56–57
Alt., F. B., 243
Ancient Israel, 39–46
Arino, O., 110
Aristotle, 194
Arrow, K. J., xvi, 290–293, 303
ASA Committee on AIDS, 92
Atkinson, E. N., 172, 173
Axelrod, D., 110

Baggerly, K. A., 38
Baggett, L. S., 385, 386
Balinski, M. L., 303
Bartoszyński, R., 38, 172, 173
Battle of the Bulge, 60
Bayes Venn diagram, 395
Bear jumps, 354–356, 414–415
Bernoulli, D., 303
Binomial distribution, 397
Black, F., 385
Black–Scholes Hedge using binomial trees, 369–374
Black–Scholes Hedge using differential equations, 374–376
Black–Scholes: some limiting cases, 376–381
Black–Scholes: search for risk neutrality, 364–384
Bootstrap, 111–130, 316–317
Boswell, S. B., 172
Box, G. E. P., 271
Box–Hunter design, 257–259
Box plot, 189–190

Bridges, E., 157, 172
Broad Street Pump, 78
Brown, B. W., 38, 172, 173

Capital asset pricing model, 339–341
Capital market line, 333–335
Capture recapture, 322–324
Casella, G., 228
Cauchy density, 250
Centers for Disease Control, 96
Central limit theorem, 405
Chaos theory, xv, 193–208
Cholera, 77–79
Circle of the elites, 105
Clements, B., 79
Cleveland, W. S., 172
Cochran, W., 305
Coldman, A., 31
Compound test, 236–239
Condensation progression, 268
Conditional density functions, 407–408
Contamination model for SPC, 231–236
Corn, 111–122
Cox, D. D., 172
Craps: an introduction to probability, 387–391

Darwin, C., 111–122
Data augmentation algorithm, 213–219
Data based wargaming, 48–63
David and Bathsheba, 273
Demming, E., xv, 248
Deming's theorem, 232
Diagnostic example, 396
Dirichlet problem, 139–142

Discrete logistic model, 197
 Diversification, 343
 Dwass, M., 38

Eddy, W., 175
 Effect of bathhouse closings, 93
 Efficient market hypothesis (EMH), 328–341
 Efron, B., 38, 130, 172
 Elliot, M. N., 172, 268, 271
 Elliptic differential equations, 142–143
 EM algorithm, 210–213
 Engel, J. H., 74
 Enigma code, 131
 Ensor, K. B., 157, 172, 206, 208
 Estimating μ and σ , 346–347
 Estimating total population size, 321–324
 Exploratory data analysis, xii, xiv, xv, 175–192

Findlay, M. C., 303, 386
 Fisher, R., 38, 111–122, 130
 Fisher–Anderson iris data, 268
 Five figure summary, 188
 Fokker–Planck equation, 138–139
 Ford, H., 2
 Fortified defense, 56–63
 Forward estimation, 156–170
 Franklin, B., 37
 Freireich, J., 209
 French Revolution, 193
 Fuzzy logic, xv
 Fwu, C., 271

Gambler's ruin, 134–143
 Game independent of the odds, 365–368
 Gamma distribution, 406–407
 Gay bathhouses, 85–91
 Gehan, E., 209
 Geometric Brownian motion for stocks, 345–346
 George, E. I., 228
 Gibbs sampler, 219–224
 Go, K., 109
 Goldie, J., 31
 Goldman–Sachs, 299
 Graunt, J., 44–47, 74
 Graunt's life table, 45–47

Haberman, R., 74
 Hamilton, A., xv
 Hamilton's rule, 294–295
 Hanning smooth, 178–184
 Hattin, battle of, 59
 Hepatitis-B, 92
 Heterosexual AIDS, 83
 High activity population, 90
 Homosexual transmission, 83
 Hull, J. C., 385
 Hunter, J. S., 257–259, 271
 Huntington, E., 298

Idealized simulation paradigm, 144
 Immune system, 68–69
 Inflation, 6
 IQ testing, 128–129
 Individual Retirement Account, 2
 Irish potato famine, xiii
 Isandlwana, battle of, 56
 Ito's Lemma, 344–345
 Iwo Jima, 54–56

Jefferson, T., xv
 Jefferson's rule, 296
 John Quincy Adams's rule, 298
 Juran, J. M., 248

Kahn, H., v, xii, 64–65, 194
 Kaposi's sarcoma, 82
 Kimmel, M., 110
 Kinsey, A., 325
 King of the mountain algorithm, 244–246
 Koronacki, J., 38, 248, 272
 Kriegspiel, 48–51
 Kuhn, T., 38

Lanchester, F. W., 51–63
 Lattes, R., 146
 Lawera, M., 248, 271
 Lawera and Thompson algorithm, 259–262
 Leprosy, 75–77
 Lorenz, E. N., 197–205
 Lorenz weather model, 201–204

Malthus, T., 18–28, 38
 McBride, C., 172
 McBride, R., 91–92
 Markowitz's efficient frontier, 329–333, 386

MaxMedian rule, 359–364
Mead, R., 271
Mean absolute deviation (MAD), 315
Mean update algorithm, 265–270
Mean update estimation of mode, 266
Median (3R) smooth, 184–187
Mandelbrot set, 195–196
Metastases, 29
Mode finding when cluster scales are very different, 269
Moment generating functions, 401–402
Monte Carlo, xiii, 131–143
Morgenstern, O., 303
Mortgage, 16
Moses, xii
Mosteller, F., 305
Multivariate normal distribution, 411

Negatively correlated portfolios, 351–354
Nelder, J. A., 271
Nelder–Mead algorithm, 253–256
Newton's method, 249
Normal distribution, 402–405

Oncological SIMEST example, 164–170
Orwell, G., 193

Parametric simulation, 415–420
Pareto, V., 303
Pareto's maxim, 9
Peters, E. E., 208
Platonic world, xii
Poisson process, 412–414
Polio, 79–80
Postmodernism, 193
Press, W. H., 271
Problem with standards, 313–316
Pyramid clubs, 69–72

Quality Control Task Force, 104

Randomly selected portfolios beating CML, 335
Rank testing with higher dimensional SPC data, 239–243
Rayleigh, W. J., 190
Reading scores, 123–128
Rejewski, M., v, 131
Resampling and histograms, 317

Resampling simulation, 420
Robust estimation for location in higher dimensions, 243–246
Rorke's Drift, battle of, 56–57
Rostenkowski–Packwood, 18
Rotello, G., 82, 109
Rubin, R., 299

Saga of Happy Valley, 212–213
Scholes, M., 385
Scott, D. W., 148, 272
Secant method, 8
Security market line, 335–327
Sharpe's measure, 340
Sharpe, W., 328–335, 386
Sharpe diagonal model, 337–339
Sharpe's super efficient frontier, 333–335
Shilts, R., 83, 109
Shreider, Y. A., 146
SIMDAT, 151–156
SIMDAT–SIMEST stopping rule, 169–170
SIMEST, 37, 147–170
SIMEST conjecture, 252
Simon, J. L., 38, 130, 228, 326
Simulated annealing, 262–264
Simulation problem solving, 132
Smallpox, 77
Smith, C., 386
Smoothed shot of moon crater, 184
Smoothing, 176–187
Snow, J., 77–79
Social Security, 2
SPC, basic notions of, 230
Statistical process control, 229–248
Stem and leafplot, 187
Stewart, I., 208
Stiglitz, J., 105
Stivers, D., 206, 208
St. Petersburg paradox, 274–283
St. Petersburg trust, 287–290
Stratego, 50
Stratification, 309–311
Structures for data in high dimensions, 264–265
Sucharski, H. 58–60

Tanner, M. A., 228
Tapia, R. A., 109, 173, 271, 272
Taylor, M. S., 173

Thompson, J. R., 38, 74, 109, 110, 146, 172, 173, 208, 271, 272, 303, 385, 386

Thucydides, xii

Tiling with hexagons, 50

Time indexed distribution of portfolio value, 348–351

Tracking polls, 308–309

Treynor's measure, 341

Tukey, J. W. xiii, xv, 38, 175–192, 305

Uniform distribution, 399–400

“U.S. as AIDS country zero,” 95–103

Volterra models of predation, 65–67, 207

von Bortkiewicz, L., 171

von Neumann, J., 131, 303

von Neumann and Morgenstern utility, 282–285

von Reiswitz, G. L. 48

Wald, A., 249, 272

War deaths vs. AIDS deaths, 93

Warehouse inventory, 311–312

Warsaw Uprising (1944), 274

Waxman, H., 84–85

Weak law of large numbers, 409–410

West, R. W., 95, 110

Westerplatte, battle of, 58–60

Wiener process, 411–412

Williams, E. E., 303, 385, 386

Young, H. P., 303

WILEY SERIES IN PROBABILITY AND STATISTICS
ESTABLISHED BY WALTER A. SHEWHART AND SAMUEL S. WILKS

Editors: *David J. Balding, Noel A. C. Cressie, Garrett M. Fitzmaurice, Harvey Goldstein, Iain M. Johnstone, Geert Molenberghs, David W. Scott, Adrian F. M. Smith, Ruey S. Tsay, Sanford Weisberg*

Editors Emeriti: *Vic Barnett, J. Stuart Hunter, Joseph B. Kadane, Jozef L. Teugels*

The *Wiley Series in Probability and Statistics* is well established and authoritative. It covers many topics of current research interest in both pure and applied statistics and probability theory. Written by leading statisticians and institutions, the titles span both state-of-the-art developments in the field and classical methods.

Reflecting the wide range of current research in statistics, the series encompasses applied, methodological and theoretical statistics, ranging from applications and new techniques made possible by advances in computerized practice to rigorous treatment of theoretical approaches.

This series provides essential and invaluable reading for all statisticians, whether in academia, industry, government, or research.

† ABRAHAM and LEDOLTER · Statistical Methods for Forecasting
AGRESTI · Analysis of Ordinal Categorical Data, *Second Edition*
AGRESTI · An Introduction to Categorical Data Analysis, *Second Edition*
AGRESTI · Categorical Data Analysis, *Second Edition*
ALTMAN, GILL, and McDONALD · Numerical Issues in Statistical Computing for the Social Scientist
AMARATUNGA and CABRERA · Exploration and Analysis of DNA Microarray and Protein Array Data
ANDĚL · Mathematics of Chance
ANDERSON · An Introduction to Multivariate Statistical Analysis, *Third Edition*
* ANDERSON · The Statistical Analysis of Time Series
ANDERSON, AUQUIER, HAUCK, OAKES, VANDAELE, and WEISBERG · Statistical Methods for Comparative Studies
ANDERSON and LOYNES · The Teaching of Practical Statistics
ARMITAGE and DAVID (editors) · Advances in Biometry
ARNOLD, BALAKRISHNAN, and NAGARAJA · Records
* ARTHANARI and DODGE · Mathematical Programming in Statistics
* BAILEY · The Elements of Stochastic Processes with Applications to the Natural Sciences
BAJORSKI · Statistics for Imaging, Optics, and Photonics
BALAKRISHNAN and KOUTRAS · Runs and Scans with Applications
BALAKRISHNAN and NG · Precedence-Type Tests and Applications
BARNETT · Comparative Statistical Inference, *Third Edition*
BARNETT · Environmental Statistics
BARNETT and LEWIS · Outliers in Statistical Data, *Third Edition*
BARTOSZYNSKI and NIEWIADOMSKA-BUGAJ · Probability and Statistical Inference
BASILEVSKY · Statistical Factor Analysis and Related Methods: Theory and Applications
BASU and RIGDON · Statistical Methods for the Reliability of Repairable Systems
BATES and WATTS · Nonlinear Regression Analysis and Its Applications
BECHHOFER, SANTNER, and GOLDSMAN · Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons
BEIRLANT, GOEGEBEUR, SEGERS, TEUGELS, and DE WAAL · Statistics of Extremes: Theory and Applications

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

† BELSLEY · Conditioning Diagnostics: Collinearity and Weak Data in Regression

† BELSLEY, KUH, and WELSCH · Regression Diagnostics: Identifying Influential Data and Sources of Collinearity

BENDAT and PIERSOL · Random Data: Analysis and Measurement Procedures, *Fourth Edition*

BERNARDO and SMITH · Bayesian Theory

BERRY, CHALONER, and GEWEKE · Bayesian Analysis in Statistics and Econometrics: Essays in Honor of Arnold Zellner

BHAT and MILLER · Elements of Applied Stochastic Processes, *Third Edition*

BHATTACHARYA and WAYMIRE · Stochastic Processes with Applications

BIEMER, GROVES, LYBERG, MATHIOWETZ, and SUDMAN · Measurement Errors in Surveys

BILLINGSLEY · Convergence of Probability Measures, *Second Edition*

BILLINGSLEY · Probability and Measure, *Third Edition*

BIRKES and DODGE · Alternative Methods of Regression

BISGAARD and KULAHCI · Time Series Analysis and Forecasting by Example

BISWAS, DATTA, FINE, and SEGAL · Statistical Advances in the Biomedical Sciences: Clinical Trials, Epidemiology, Survival Analysis, and Bioinformatics

BLISCHKE AND MURTHY (editors) · Case Studies in Reliability and Maintenance

BLISCHKE AND MURTHY · Reliability: Modeling, Prediction, and Optimization

BLOOMFIELD · Fourier Analysis of Time Series: An Introduction, *Second Edition*

BOLLEN · Structural Equations with Latent Variables

BOLLEN and CURRAN · Latent Curve Models: A Structural Equation Perspective

BOROVKOV · Ergodicity and Stability of Stochastic Processes

BOSQ and BLANKE · Inference and Prediction in Large Dimensions

BOULEAU · Numerical Methods for Stochastic Processes

BOX · Bayesian Inference in Statistical Analysis

BOX · Improving Almost Anything, *Revised Edition*

BOX · R. A. Fisher, the Life of a Scientist

BOX and DRAPER · Empirical Model-Building and Response Surfaces

* BOX and DRAPER · Evolutionary Operation: A Statistical Method for Process Improvement

BOX and DRAPER · Response Surfaces, Mixtures, and Ridge Analyses, *Second Edition*

BOX, HUNTER, and HUNTER · Statistics for Experimenters: Design, Innovation, and Discovery, *Second Edition*

BOX, JENKINS, and REINSEL · Time Series Analysis: Forecasting and Control, *Fourth Edition*

BOX, LUCEÑO, and PANIAGUA-QUIÑONES · Statistical Control by Monitoring and Adjustment, *Second Edition*

BRANDIMARTE · Numerical Methods in Finance: A MATLAB-Based Introduction

† BROWN and HOLLANDER · Statistics: A Biomedical Introduction

BRUNNER, DOMHOF, and LANGER · Nonparametric Analysis of Longitudinal Data in Factorial Experiments

BUCKLEW · Large Deviation Techniques in Decision, Simulation, and Estimation

CAIROLI and DALANG · Sequential Stochastic Optimization

CASTILLO, HADI, BALAKRISHNAN, and SARABIA · Extreme Value and Related Models with Applications in Engineering and Science

CHAN · Time Series: Applications to Finance with R and S-Plus®, *Second Edition*

CHARALAMIDES · Combinatorial Methods in Discrete Distributions

CHATTERJEE and HADI · Regression Analysis by Example, *Fourth Edition*

CHATTERJEE and HADI · Sensitivity Analysis in Linear Regression

CHERNICK · Bootstrap Methods: A Guide for Practitioners and Researchers, *Second Edition*

CHERNICK and FRIIS · Introductory Biostatistics for the Health Sciences

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

CHILÈS and DELFINER · Geostatistics: Modeling Spatial Uncertainty

CHOW and LIU · Design and Analysis of Clinical Trials: Concepts and Methodologies, *Second Edition*

CLARKE · Linear Models: The Theory and Application of Analysis of Variance

CLARKE and DISNEY · Probability and Random Processes: A First Course with Applications, *Second Edition*

* COCHRAN and COX · Experimental Designs, *Second Edition*

COLLINS and LANZA · Latent Class and Latent Transition Analysis: With Applications in the Social, Behavioral, and Health Sciences

CONGDON · Applied Bayesian Modelling

CONGDON · Bayesian Models for Categorical Data

CONGDON · Bayesian Statistical Modelling, *Second Edition*

CONOVER · Practical Nonparametric Statistics, *Third Edition*

COOK · Regression Graphics

COOK and WEISBERG · An Introduction to Regression Graphics

COOK and WEISBERG · Applied Regression Including Computing and Graphics

CORNELL · A Primer on Experiments with Mixtures

CORNELL · Experiments with Mixtures, Designs, Models, and the Analysis of Mixture Data, *Third Edition*

COVER and THOMAS · Elements of Information Theory

COX · A Handbook of Introductory Statistical Methods

* COX · Planning of Experiments

CRESSIE · Statistics for Spatial Data, *Revised Edition*

CRESSIE and WIKLE · Statistics for Spatio-Temporal Data

CSÓRGÓ and HORVÁTH · Limit Theorems in Change Point Analysis

DANIEL · Applications of Statistics to Industrial Experimentation

DANIEL · Biostatistics: A Foundation for Analysis in the Health Sciences, *Eighth Edition*

* DANIEL · Fitting Equations to Data: Computer Analysis of Multifactor Data, *Second Edition*

DASU and JOHNSON · Exploratory Data Mining and Data Cleaning

DAVID and NAGARAJA · Order Statistics, *Third Edition*

* DEGROOT, FIENBERG, and KADANE · Statistics and the Law

DEL CASTILLO · Statistical Process Adjustment for Quality Control

DEMARIS · Regression with Social Data: Modeling Continuous and Limited Response Variables

DEMIDENKO · Mixed Models: Theory and Applications

DENISON, HOLMES, MALLICK and SMITH · Bayesian Methods for Nonlinear Classification and Regression

DETTE and STUDDEN · The Theory of Canonical Moments with Applications in Statistics, Probability, and Analysis

DEY and MUKERJEE · Fractional Factorial Plans

DILLON and GOLDSTEIN · Multivariate Analysis: Methods and Applications

DODGE · Alternative Methods of Regression

* DODGE and ROMIG · Sampling Inspection Tables, *Second Edition*

* DOOB · Stochastic Processes

DOWDY, WEARDEN, and CHILKO · Statistics for Research, *Third Edition*

DRAPER and SMITH · Applied Regression Analysis, *Third Edition*

DRYDEN and MARDIA · Statistical Shape Analysis

DUDEWICZ and MISHRA · Modern Mathematical Statistics

DUNN and CLARK · Basic Statistics: A Primer for the Biomedical Sciences, *Third Edition*

DUPUIS and ELLIS · A Weak Convergence Approach to the Theory of Large Deviations

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

EDLER and KITSOS · Recent Advances in Quantitative Methods in Cancer and Human Health Risk Assessment

* ELANDT-JOHNSON and JOHNSON · Survival Models and Data Analysis

ENDERS · Applied Econometric Time Series

† ETHIER and KURTZ · Markov Processes: Characterization and Convergence

EVANS, HASTINGS, and PEACOCK · Statistical Distributions, *Third Edition*

EVERITT · Cluster Analysis, *Fifth Edition*

FELLER · An Introduction to Probability Theory and Its Applications, Volume I, *Third Edition*, Revised; Volume II, *Second Edition*

FISHER and VAN BELLE · Biostatistics: A Methodology for the Health Sciences

FITZMAURICE, LAIRD, and WARE · Applied Longitudinal Analysis, *Second Edition*

* FLEISS · The Design and Analysis of Clinical Experiments

FLEISS · Statistical Methods for Rates and Proportions, *Third Edition*

† FLEMING and HARRINGTON · Counting Processes and Survival Analysis

FUJIKOSHI, ULYANOV, and SHIMIZU · Multivariate Statistics: High-Dimensional and Large-Sample Approximations

FULLER · Introduction to Statistical Time Series, *Second Edition*

† FULLER · Measurement Error Models

GALLANT · Nonlinear Statistical Models

GEISSER · Modes of Parametric Statistical Inference

GELMAN and MENG · Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives

GEWEKE · Contemporary Bayesian Econometrics and Statistics

GHOSH, MUKHOPADHYAY, and SEN · Sequential Estimation

GIESBRECHT and GUMPERTZ · Planning, Construction, and Statistical Analysis of Comparative Experiments

GIFI · Nonlinear Multivariate Analysis

GIVENS and HOETING · Computational Statistics

GLASSERMAN and YAO · Monotone Structure in Discrete-Event Systems

GNANADESIKAN · Methods for Statistical Data Analysis of Multivariate Observations, *Second Edition*

GOLDSTEIN · Multilevel Statistical Models, *Fourth Edition*

GOLDSTEIN and LEWIS · Assessment: Problems, Development, and Statistical Issues

GOLDSTEIN and WOOFF · Bayes Linear Statistics

GREENWOOD and NIKULIN · A Guide to Chi-Squared Testing

GROSS, SHORTLE, THOMPSON, and HARRIS · Fundamentals of Queueing Theory, *Fourth Edition*

GROSS, SHORTLE, THOMPSON, and HARRIS · Solutions Manual to Accompany Fundamentals of Queueing Theory, *Fourth Edition*

* HAHN and SHAPIRO · Statistical Models in Engineering

HAHN and MEEKER · Statistical Intervals: A Guide for Practitioners

HALD · A History of Probability and Statistics and their Applications Before 1750

HALD · A History of Mathematical Statistics from 1750 to 1930

† HAMPEL · Robust Statistics: The Approach Based on Influence Functions

HANNAN and DEISTLER · The Statistical Theory of Linear Systems

HARMAN and KULKARNI · An Elementary Introduction to Statistical Learning Theory

HARTUNG, KNAPP, and SINHA · Statistical Meta-Analysis with Applications

HEIBERGER · Computation for the Analysis of Designed Experiments

HEDAYAT and SINHA · Design and Inference in Finite Population Sampling

HEDEKER and GIBBONS · Longitudinal Data Analysis

HELLER · MACSYMA for Statisticians

HERITIER, CANTONI, COPT, and VICTORIA-FESER · Robust Methods in Biostatistics

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

HINKELMANN and KEMPTHORNE · Design and Analysis of Experiments, Volume 1: Introduction to Experimental Design, *Second Edition*

HINKELMANN and KEMPTHORNE · Design and Analysis of Experiments, Volume 2: Advanced Experimental Design

HOAGLIN, MOSTELLER, and TUKEY · Fundamentals of Exploratory Analysis of Variance

* HOAGLIN, MOSTELLER, and TUKEY · Exploring Data Tables, Trends and Shapes

* HOAGLIN, MOSTELLER, and TUKEY · Understanding Robust and Exploratory Data Analysis

HOCHBERG and TAMHANE · Multiple Comparison Procedures

HOCKING · Methods and Applications of Linear Models: Regression and the Analysis of Variance, *Second Edition*

HOEL · Introduction to Mathematical Statistics, *Fifth Edition*

HOGG and KLUGMAN · Loss Distributions

HOLLANDER and WOLFE · Nonparametric Statistical Methods, *Second Edition*

HOSMER and LEMESHOW · Applied Logistic Regression, *Second Edition*

HOSMER, LEMESHOW, and MAY · Applied Survival Analysis: Regression Modeling of Time-to-Event Data, *Second Edition*

HUBER · Data Analysis: What Can Be Learned From the Past 50 Years

HUBER · Robust Statistics

† HUBER and RONCHETTI · Robust Statistics, *Second Edition*

HUBERTY · Applied Discriminant Analysis, *Second Edition*

HUBERTY and OLEJNIK · Applied MANOVA and Discriminant Analysis, *Second Edition*

HUITEMA · The Analysis of Covariance and Alternatives: Statistical Methods for Experiments, Quasi-Experiments, and Single-Case Studies, *Second Edition*

HUNT and KENNEDY · Financial Derivatives in Theory and Practice, *Revised Edition*

HURD and MIAMEE · Periodically Correlated Random Sequences: Spectral Theory and Practice

HUSKOVA, BERAN, and DUPAC · Collected Works of Jaroslav Hajek—with Commentary

HUZURBAZAR · Flowgraph Models for Multistate Time-to-Event Data

IMAN and CONOVER · A Modern Approach to Statistics

JACKMAN · Bayesian Analysis for the Social Sciences

† JACKSON · A User's Guide to Principle Components

JOHN · Statistical Methods in Engineering and Quality Assurance

JOHNSON · Multivariate Statistical Simulation

JOHNSON and BALAKRISHNAN · Advances in the Theory and Practice of Statistics: A Volume in Honor of Samuel Kotz

JOHNSON and BHATTACHARYYA · Statistics: Principles and Methods, *Fifth Edition*

JOHNSON, KEMP, and KOTZ · Univariate Discrete Distributions, *Third Edition*

JOHNSON and KOTZ · Distributions in Statistics

JOHNSON and KOTZ (editors) · Leading Personalities in Statistical Sciences: From the Seventeenth Century to the Present

JOHNSON, KOTZ, and BALAKRISHNAN · Continuous Univariate Distributions, Volume 1, *Second Edition*

JOHNSON, KOTZ, and BALAKRISHNAN · Continuous Univariate Distributions, Volume 2, *Second Edition*

JOHNSON, KOTZ, and BALAKRISHNAN · Discrete Multivariate Distributions

JUDGE, GRIFFITHS, HILL, LÜTKEPOHL, and LEE · The Theory and Practice of Econometrics, *Second Edition*

JUREČKOVÁ and SEN · Robust Statistical Procedures: Asymptotics and Interrelations

JUREK and MASON · Operator-Limit Distributions in Probability Theory

KADANE · Bayesian Methods and Ethics in a Clinical Trial Design

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

KADANE AND SCHUM · A Probabilistic Analysis of the Sacco and Vanzetti Evidence
 KALBFLEISCH and PRENTICE · The Statistical Analysis of Failure Time Data, *Second Edition*

KARIYA and KURATA · Generalized Least Squares
 KASS and VOS · Geometrical Foundations of Asymptotic Inference

† KAUFMAN and ROUSSEEUW · Finding Groups in Data: An Introduction to Cluster Analysis
 KEDEM and FOKIANOS · Regression Models for Time Series Analysis
 KENDALL, BARDEN, CARNE, and LE · Shape and Shape Theory
 KHURI · Advanced Calculus with Applications in Statistics, *Second Edition*
 KHURI, MATHEW, and SINHA · Statistical Tests for Mixed Linear Models

* KISH · Statistical Design for Research
 KLEIBER and KOTZ · Statistical Size Distributions in Economics and Actuarial Sciences
 KLEMELÄ · Smoothing of Multivariate Data: Density Estimation and Visualization
 KLUGMAN, PANJER, and WILLMOT · Loss Models: From Data to Decisions,
Third Edition
 KLUGMAN, PANJER, and WILLMOT · Solutions Manual to Accompany Loss Models:
 From Data to Decisions, *Third Edition*
 KOSKI and NOBLE · Bayesian Networks: An Introduction
 KOTZ, BALAKRISHNAN, and JOHNSON · Continuous Multivariate Distributions,
 Volume 1, *Second Edition*
 KOTZ and JOHNSON (editors) · Encyclopedia of Statistical Sciences: Volumes 1 to 9
 with Index
 KOTZ and JOHNSON (editors) · Encyclopedia of Statistical Sciences: Supplement
 Volume
 KOTZ, READ, and BANKS (editors) · Encyclopedia of Statistical Sciences: Update
 Volume 1
 KOTZ, READ, and BANKS (editors) · Encyclopedia of Statistical Sciences: Update
 Volume 2
 KOVALENKO, KUZNETZOV, and PEGG · Mathematical Theory of Reliability of
 Time-Dependent Systems with Practical Applications
 KOWALSKI and TU · Modern Applied U-Statistics
 KRISHNAMOORTHY and MATHEW · Statistical Tolerance Regions: Theory,
 Applications, and Computation
 KROESE, TAIMRE, and BOTEV · Handbook of Monte Carlo Methods
 KROONENBERG · Applied Multiway Data Analysis
 KULINSKAYA, MORGENTHALER, and STAUDTE · Meta Analysis: A Guide to
 Calibrating and Combining Statistical Evidence
 KUROWICKA and COOKE · Uncertainty Analysis with High Dimensional Dependence
 Modelling
 KVAM and VIDAKOVIC · Nonparametric Statistics with Applications to Science
 and Engineering
 LACHIN · Biostatistical Methods: The Assessment of Relative Risks, *Second Edition*
 LAD · Operational Subjective Statistical Methods: A Mathematical, Philosophical, and
 Historical Introduction
 LAMPERTI · Probability: A Survey of the Mathematical Theory, *Second Edition*
 LANGE, RYAN, BILLARD, BRILLINGER, CONQUEST, and GREENHOUSE ·
 Case Studies in Biometry
 LARSON · Introduction to Probability Theory and Statistical Inference, *Third Edition*
 LAWLESS · Statistical Models and Methods for Lifetime Data, *Second Edition*
 LAWSON · Statistical Methods in Spatial Epidemiology, *Second Edition*
 LE · Applied Categorical Data Analysis
 LE · Applied Survival Analysis
 LEE · Structural Equation Modeling: A Bayesian Approach

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

LEE and WANG · Statistical Methods for Survival Data Analysis, *Third Edition*
 LePAGE and BILLARD · Exploring the Limits of Bootstrap
 LEYLAND and GOLDSTEIN (editors) · Multilevel Modelling of Health Statistics
 LIAO · Statistical Group Comparison
 LINDVALL · Lectures on the Coupling Method
 LIN · Introductory Stochastic Analysis for Finance and Insurance
 LINHART and ZUCCHINI · Model Selection
 LITTLE and RUBIN · Statistical Analysis with Missing Data, *Second Edition*
 LLOYD · The Statistical Analysis of Categorical Data
 LOWEN and TEICH · Fractal-Based Point Processes
 MAGNUS and NEUDECKER · Matrix Differential Calculus with Applications in
 Statistics and Econometrics, *Revised Edition*
 MALLER and ZHOU · Survival Analysis with Long Term Survivors
 MALLOWS · Design, Data, and Analysis by Some Friends of Cuthbert Daniel
 MANN, SCHAFER, and SINGPURWALLA · Methods for Statistical Analysis of
 Reliability and Life Data
 MANTON, WOODBURY, and TOLLEY · Statistical Applications Using Fuzzy Sets
 MARCHETTE · Random Graphs for Statistical Pattern Recognition
 MARDIA and JUPP · Directional Statistics
 MARKOVICH · Nonparametric Analysis of Univariate Heavy-Tailed Data: Research and
 Practice
 MARONNA, MARTIN and YOHAI · Robust Statistics: Theory and Methods
 MASON, GUNST, and HESS · Statistical Design and Analysis of Experiments with
 Applications to Engineering and Science, *Second Edition*
 McCULLOCH, SEARLE, and NEUHAUS · Generalized, Linear, and Mixed Models,
 Second Edition
 McFADDEN · Management of Data in Clinical Trials, *Second Edition*
 * McLACHLAN · Discriminant Analysis and Statistical Pattern Recognition
 McLACHLAN, DO, and AMBROISE · Analyzing Microarray Gene Expression Data
 McLACHLAN and KRISHNAN · The EM Algorithm and Extensions, *Second Edition*
 McLACHLAN and PEEL · Finite Mixture Models
 McNEIL · Epidemiological Research Methods
 MEEKER and ESCOBAR · Statistical Methods for Reliability Data
 MEERSCHAERT and SCHEFFLER · Limit Distributions for Sums of Independent
 Random Vectors: Heavy Tails in Theory and Practice
 MENGERSEN, ROBERT, and TITTERINGTON · Mixtures: Estimation and
 Applications
 MICKEY, DUNN, and CLARK · Applied Statistics: Analysis of Variance and
 Regression, *Third Edition*
 * MILLER · Survival Analysis, *Second Edition*
 MONTGOMERY, JENNINGS, and KULAHCI · Introduction to Time Series Analysis
 and Forecasting
 MONTGOMERY, PECK, and VINING · Introduction to Linear Regression Analysis,
 Fourth Edition
 MORGENTHALER and TUKEY · Configural Polysampling: A Route to Practical
 Robustness
 MUIRHEAD · Aspects of Multivariate Statistical Theory
 MULLER and STOYAN · Comparison Methods for Stochastic Models and Risks
 MURRAY · X-STAT 2.0 Statistical Experimentation, Design Data Analysis, and
 Nonlinear Optimization
 MURTHY, XIE, and JIANG · Weibull Models
 MYERS, MONTGOMERY, and ANDERSON-COOK · Response Surface Methodology:
 Process and Product Optimization Using Designed Experiments, *Third Edition*

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

MYERS, MONTGOMERY, VINING, and ROBINSON · Generalized Linear Models.
 With Applications in Engineering and the Sciences, *Second Edition*

† NELSON · Accelerated Testing, Statistical Models, Test Plans, and Data Analyses

† NELSON · Applied Life Data Analysis

NEWMAN · Biostatistical Methods in Epidemiology

OCHI · Applied Probability and Stochastic Processes in Engineering and Physical Sciences

OKABE, BOOTS, SUGIHARA, and CHIU · Spatial Tesselations: Concepts and Applications of Voronoi Diagrams, *Second Edition*

OLIVER and SMITH · Influence Diagrams, Belief Nets and Decision Analysis

PALTA · Quantitative Methods in Population Health: Extensions of Ordinary Regressions

PANJER · Operational Risk: Modeling and Analytics

PANKRATZ · Forecasting with Dynamic Regression Models

PANKRATZ · Forecasting with Univariate Box-Jenkins Models: Concepts and Cases

PARDOUX · Markov Processes and Applications: Algorithms, Networks, Genome and Finance

* PARZEN · Modern Probability Theory and Its Applications

PEÑA, TIAO, and TSAY · A Course in Time Series Analysis

PIANTADOSI · Clinical Trials: A Methodologic Perspective

PORT · Theoretical Probability for Applications

POURAHMADI · Foundations of Time Series Analysis and Prediction Theory

POWELL · Approximate Dynamic Programming: Solving the Curses of Dimensionality, *Second Edition*

PRESS · Bayesian Statistics: Principles, Models, and Applications

PRESS · Subjective and Objective Bayesian Statistics, *Second Edition*

PRESS and TANUR · The Subjectivity of Scientists and the Bayesian Approach

PUKELSHEIM · Optimal Experimental Design

PURI, VILAPLANA, and WERTZ · New Perspectives in Theoretical and Applied Statistics

† PUTERMAN · Markov Decision Processes: Discrete Stochastic Dynamic Programming

QIU · Image Processing and Jump Regression Analysis

* RAO · Linear Statistical Inference and Its Applications, *Second Edition*

RAO · Statistical Inference for Fractional Diffusion Processes

RAUSAND and HØYLAND · System Reliability Theory: Models, Statistical Methods, and Applications, *Second Edition*

RAYNER · Smooth Tests of Goodness of Fit: Using R, *Second Edition*

RENCHER · Linear Models in Statistics

RENCHER · Methods of Multivariate Analysis, *Second Edition*

RENCHER · Multivariate Statistical Inference with Applications

* RIPLEY · Spatial Statistics

* RIPLEY · Stochastic Simulation

ROBINSON · Practical Strategies for Experimenting

ROHATGI and SALEH · An Introduction to Probability and Statistics, *Second Edition*

ROLSKI, SCHMIDLI, SCHMIDT, and TEUGELS · Stochastic Processes for Insurance and Finance

ROSENBERGER and LACHIN · Randomization in Clinical Trials: Theory and Practice

ROSS · Introduction to Probability and Statistics for Engineers and Scientists

ROSSI, ALLENBY, and McCULLOCH · Bayesian Statistics and Marketing

† ROUSSEEUW and LEROY · Robust Regression and Outlier Detection

ROYSTON and SAUERBREI · Multivariate Model Building: A Pragmatic Approach to Regression Analysis Based on Fractional Polynomials for Modeling Continuous Variables

* RUBIN · Multiple Imputation for Nonresponse in Surveys

RUBINSTEIN and KROESE · Simulation and the Monte Carlo Method, *Second Edition*

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

RUBINSTEIN and MELAMED · Modern Simulation and Modeling

RYAN · Modern Engineering Statistics

RYAN · Modern Experimental Design

RYAN · Modern Regression Methods, *Second Edition*

RYAN · Statistical Methods for Quality Improvement, *Third Edition*

SALEH · Theory of Preliminary Test and Stein-Type Estimation with Applications

SALTELLI, CHAN, and SCOTT (editors) · Sensitivity Analysis

* SCHEFFE · The Analysis of Variance

SCHIMEK · Smoothing and Regression: Approaches, Computation, and Application

SCHOTT · Matrix Analysis for Statistics, *Second Edition*

SCHOUTENS · Levy Processes in Finance: Pricing Financial Derivatives

SCHUSS · Theory and Applications of Stochastic Differential Equations

SCOTT · Multivariate Density Estimation: Theory, Practice, and Visualization

* SEARLE · Linear Models

† SEARLE · Linear Models for Unbalanced Data

† SEARLE · Matrix Algebra Useful for Statistics

† SEARLE, CASELLA, and McCULLOCH · Variance Components

SEARLE and WILLETT · Matrix Algebra for Applied Economics

SEBER · A Matrix Handbook For Statisticians

† SEBER · Multivariate Observations

SEBER and LEE · Linear Regression Analysis, *Second Edition*

† SEBER and WILD · Nonlinear Regression

SENNOTT · Stochastic Dynamic Programming and the Control of Queueing Systems

* SERFLING · Approximation Theorems of Mathematical Statistics

SHAFER and VOVK · Probability and Finance: It's Only a Game!

SHERMAN · Spatial Statistics and Spatio-Temporal Data: Covariance Functions and Directional Properties

SILVAPULLE and SEN · Constrained Statistical Inference: Inequality, Order, and Shape Restrictions

SINGPURWALLA · Reliability and Risk: A Bayesian Perspective

SMALL and MCLEISH · Hilbert Space Methods in Probability and Statistical Inference

SRIVASTAVA · Methods of Multivariate Statistics

STAPLETON · Linear Statistical Models, *Second Edition*

STAPLETON · Models for Probability and Statistical Inference: Theory and Applications

STAUDTE and SHEATHER · Robust Estimation and Testing

STOYAN, KENDALL, and MECKE · Stochastic Geometry and Its Applications, *Second Edition*

STOYAN and STOYAN · Fractals, Random Shapes and Point Fields: Methods of Geometrical Statistics

STREET and BURGESS · The Construction of Optimal Stated Choice Experiments: Theory and Methods

STYAN · The Collected Papers of T. W. Anderson: 1943–1985

SUTTON, ABRAMS, JONES, SHELDON, and SONG · Methods for Meta-Analysis in Medical Research

TAKEZAWA · Introduction to Nonparametric Regression

TAMHANE · Statistical Analysis of Designed Experiments: Theory and Applications

TANAKA · Time Series Analysis: Nonstationary and Noninvertible Distribution Theory

THOMPSON · Empirical Model Building: Data, Models, and Reality, *Second Edition*

THOMPSON · Sampling, *Third Edition*

THOMPSON · Simulation: A Modeler's Approach

THOMPSON and SEBER · Adaptive Sampling

THOMPSON, WILLIAMS, and FINDLAY · Models for Investors in Real World Markets

TCIAO, BISGAARD, HILL, PEÑA, and STIGLER (editors) · Box on Quality and Discovery: with Design, Control, and Robustness

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.

TIERNEY · LISP-STAT: An Object-Oriented Environment for Statistical Computing and Dynamic Graphics

TSAY · Analysis of Financial Time Series, *Third Edition*

UPTON and FINGLETON · Spatial Data Analysis by Example, Volume II: Categorical and Directional Data

† VAN BELLE · Statistical Rules of Thumb, *Second Edition*

VAN BELLE, FISHER, HEAGERTY, and LUMLEY · Biostatistics: A Methodology for the Health Sciences, *Second Edition*

VESTRUP · The Theory of Measures and Integration

VIDAKOVIC · Statistical Modeling by Wavelets

VINOD and REAGLE · Preparing for the Worst: Incorporating Downside Risk in Stock Market Investments

WALLER and GOTWAY · Applied Spatial Statistics for Public Health Data

WEERAHANDI · Generalized Inference in Repeated Measures: Exact Methods in MANOVA and Mixed Models

WEISBERG · Applied Linear Regression, *Third Edition*

WEISBERG · Bias and Causation: Models and Judgment for Valid Comparisons

WELSH · Aspects of Statistical Inference

WESTFALL and YOUNG · Resampling-Based Multiple Testing: Examples and Methods for *p*-Value Adjustment

* WHITTAKER · Graphical Models in Applied Multivariate Statistics

WINKER · Optimization Heuristics in Economics: Applications of Threshold Accepting

WONNACOTT and WONNACOTT · Econometrics, *Second Edition*

WOODING · Planning Pharmaceutical Clinical Trials: Basic Statistical Principles

WOODWORTH · Biostatistics: A Bayesian Introduction

WOOLSON and CLARKE · Statistical Methods for the Analysis of Biomedical Data, *Second Edition*

WU and HAMADA · Experiments: Planning, Analysis, and Parameter Design Optimization, *Second Edition*

WU and ZHANG · Nonparametric Regression Methods for Longitudinal Data Analysis

YANG · The Construction Theory of Denumerable Markov Processes

YOUNG, VALERO-MORA, and FRIENDLY · Visual Statistics: Seeing Data with Dynamic Interactive Graphics

ZACKS · Stage-Wise Adaptive Designs

* ZELLNER · An Introduction to Bayesian Inference in Econometrics

ZELTERMAN · Discrete Distributions—Applications in the Health Sciences

ZHOU, OBUCHOWSKI, and McCLISH · Statistical Methods in Diagnostic Medicine, *Second Edition*

*Now available in a lower priced paperback edition in the Wiley Classics Library.

†Now available in a lower priced paperback edition in the Wiley-Interscience Paperback Series.