Preface

This Preface is useful for explaining to both the reader and the instructor my objectives in presenting the material in this text as I did. I begin by defining the intended readership and the organization of the text. This leads to a discussion of the innovations in the text and the reasons for including them. Next, I discuss the careful design that underlies the Exercises; mere "numbers" are not enough. I then summarize the required level of mathematical reasoning in the last section.

Because my objectives for the reader of this text differ from those in conventional texts, some students may require guidance to use the text most effectively. This is especially true of the Exercises. Consequently, at the end of Chapter 1, I have provided an "Addendum for the Reader," which every student is strongly encouraged to read.

The Intended Audience

This is intended as an introduction to statistical reasoning for undergraduates at the sophomore level who have an inquiring mind. By this, I mean the type of student who, if presented with a challenge, will attempt to respond and, if encouraged and aided to do so, will enjoy experimenting with new ideas.

Because I am an economist who teaches in a faculty of arts and science department as well as in a business school, my initial tendency was to use economic and business examples. But I soon recognized in the early stages of teaching my course that students were more interested in and receptive to examples to which they could relate personally. Consequently, as my objective was to engage students' minds, I decided that I had to show them that "statistical reasoning" is everywhere, and that they could not escape having to deal with statistical concepts. Statistical concepts were important in their personal lives, not just as a tool in their professional careers. The result is a book that is intended for an audience much broader than what is traditionally regarded as "economics and business." The intended audience is restricted only by the inquisitiveness of the student and the acceptance of innovation by the instructor. This book may be a challenge for some undergraduates, but a little perseverance will be highly rewarded. However, it is also true that the book's approach and use of computers means that for such a student learning statistics can actually be intellectually fun.

Later, I will discuss the level of mathematical reasoning this text involves, but first I want to define the text's objectives. The limited objective is to produce statistically literate people who will have acquired enough background to be taught specific tools in a work environment and who can read the popular press with new insight. By no means do I think that two quarters of instruction, or even two semesters of statistics, is enough to produce a "statistician"—someone able to work on his or her own within any organization. Further, I suspect no other instructor does either. Consequently, the task that I set myself was to produce a book that would enable a student to gain sufficient insight into statistical reasoning and practice to be able to understand the major strengths and weaknesses of the statistical approach and to appreciate what it is that statistical methods can and cannot achieve. In this connection and for the business school student in particular, it is my objective to produce someone who has sufficient theoretical background to be able to learn quickly and efficiently what specifically has to be mastered within any organization. Firms, agencies, and nonprofits all want to train their own people using their own methods. At the entry level, they are not hiring Ph.D.'s in statistics. What is wanted in the real world is the ability to learn fast. This text attempts to provide this facility in the context of statistical reasoning.

This perspective marks a substantial difference from the spirit that pervades many contemporary texts. The contemporary view seems to be to give the student as extensive an exposure to statistical procedures as possible, presumably with a view to providing someone who can enter a firm and begin applications immediately. But in one semester, such an extensive exposure comes at the cost of the intensive margin; the student gets a lot without much depth. The philosophy underlying this text is that given a reasonable level of depth in the analysis, students can later acquire a much more extensive, and even more intensive, exposure to statistics on their own, or in the context of the work environment. This text is not intended to be an encyclopedia of statistical techniques.

For the mathematically prepared student, the entire book can be covered comfortably in one semester, with classes meeting for 1 hour and 15 minutes. Students who are not well prepared mathematically also can complete the text (with the exception of Chapters 12 and 14 and the sections labeled "For the Student Comfortable with Calculus") within this time frame.

Some words about pacing are useful. The first two chapters in particular are fairly easy for students to grasp on their own; two lectures will suffice, and the instructor can leave the details to the text. The next three chapters are more challenging but less so than Chapters 6 on. The book is designed so that the pace begins easily but quickly builds up. This is to help students build confidence in a course of study that has a reputation for opacity and difficulty. During the earlier "easier" period, students should ensure they have mastery over the basic algebra that will be used intensively throughout the remainder of the course; I use the first few "lab sessions" for this purpose.

The most challenging chapters that require the most time are Chapters 6 through 8 and 12. These chapters develop the basic theory, and students always have difficulty in mastering theoretical concepts. Once these have been mastered the remainder can be handled by students with equanimity, in part, because from their perspective the later material is much more concerned with practical issues.

The Organization of This Text

The basic outline of the material covered in the text is simple. The text begins with descriptive statistics. This beginning is motivated by the notion that given data that are presumed to be *acausal,* either through ignorance of the model underlying the data-generating mechanism, or because there are no causal links, one has to develop new specific tools to analyze such a unique type of data. The first two parts of the text do far more than merely list descriptive statistics. Students will see by numerous examples that large numbers of repetitions of an experiment give rise to uniquely shaped histograms. These insights are extended to bivariate data in Chapter 5. These early chapters stress the importance of the shape of histograms and the connection between shape and the experiment generating the data. Including a discussion of the description of bivariate data so early in the course is unusual but is pedagogically reasonable given my approach. Students are encouraged in the context of descriptive statistics to extend the ideas of the previous chapters to bivariate data.

By this stage in the text, students will realize that observations on random data involve regularities that are to be explained by some "theory of statistics," and that these regularities involve the shapes of bivariate histograms or scatter diagrams as well.

In Part Three, Probability theory is introduced as the theory of statistics and, more precisely, as the theory of histograms. Distributions enable students to see in detail the theoretical analysis underlying the idea of a probability distribution as an analogue of sample histograms.

Chapter 9, Elementary Sampling Theory, opens Part Four and provides the bridge between probability and distribution theory in Part Three and the observations discussed in Parts One and Two. The link between theory and observation culminates in Chapters 10 and 11, on estimation and hypothesis testing, respectively. Part Five opens with a discussion of the theory of regression in Chapter 13 and the analysis of variance in Chapter 14; this discussion complements the purely descriptive analysis contained in Chapter 5. Chapter 15 offers a retrospective on the major lessons of the text.

The Innovations in This Text

A basic feature of the style of exposition of this text is one of discovery for students. Each new step poses new questions that require answers, and students are guided in discovering the answers. Each new step provides a reasonable solution to a practical problem; results are not presented as *obiter dicta,* or incidental observations, to be remembered and not questioned. Consequently, the exposition explores the alternative solutions to a problem and then settles on a reasonable resolution as a useful procedure until a better idea comes up. Statistical procedures are presented as practical attempts to solve practical problems—not as arbitrary rules to be blindly followed by the uninitiated.

The approach taken in this text is *process* oriented, not *fact* oriented. For example, I want students to be aware of the process of convergence, not just the existence of such a property. I am more concerned that students have some intuitive feel for the process involved in the central limit theorem than that they merely know of the existence of such a phenomenon.

One key feature is to get students to realize the unique aspects of statistical theory. The first of the two most important aspects is that statistics *per se* deal with random phenomena, phenomena which by their definition are not causally related. Such phenomena are very different from all that students have been exposed to during their prior academic careers. Statistics are in essence *acausal*. Keeping this fact firmly in mind helps students to deal later in the text with the subtle issues of how to use statistical analysis to evaluate scientific hypotheses. For example, knowing that statistics are *acausal* enables students to recognize that correlation is neither necessary nor sufficient for causality. Causality can only have meaning in the context of the theory of physical phenomena, or the theory of economic phenomena, or the theory of political phenomena, and so on.

Consequently, the text begins by considering data for which we have no explanation at all; no model; and no ability to predict values, either from observing other variables or from observing the past. Having no theory, no model, we are forced to consider what we can learn from such phenomena. From performing calculations of descriptive statistics and then observing the histograms of a large number of observations on a variety of phenomena, or different experiments, students learn pragmatically several lessons. They discover from the examples that for continuous variables the histograms become smooth; that the same experiment repeated gives rise to the same shape of histogram, or relative frequency distribution for discrete data; and that different experiments give rise to different shapes of histograms, or relative frequency plots. A wide variety of examples of distributions are presented. In each case, the observed shape of a distribution is linked to the conditions of the experiment that generated it. And this is, of course, the key to statistical analysis. Students, having begun with the idea that when observing random data there was nothing to explain, now see that there are a lot of regularities and that a theory is needed to explain them. The stage is set for the development of probability theory and the derivation of density functions that is to come in the second part of the text.

Within the first two parts of the text, the first four moments play a substantial role in formalizing the ideas of the shape of a histogram and a distribution function. The focus is on enabling students to relate the shape of histograms to the values of standardized moments. Thus, students should be able to provide a reasonable guess as to the values of the first four moments from observing a smooth histogram and, in turn, provide a reasonable approximation of a smooth histogram from a given set of moments. This experience in calculating the first four sample moments and in interpreting the results provides a natural introduction to the analysis of distributions that is to come in the theoretical part of the text.

In Chapter 5, I introduce bivariate data and correlation as a *descriptive* measure of linear "association." The correlation coefficient is developed as the "first cross product moment," and as such is a natural extension of the moments discussed in Chapter 4. Further, the correlation coefficient is another measure of shape, in this case of the distribution of bivariate data. This material provides the summary statistics explained by the theory developed later in the text. The analogy between the development of descriptive bivariate statistics and the descriptive statistics developed in prior chapters is strongly emphasized; this analogy continues to apply to the corresponding development of the theory in the chapters on bivariate distributions, ANOVA, and regression.

All distributions are derived while emphasizing the relationship between experiment and the shape of the distribution. These results mirror the observed empirical results from Chapter 3. Students are led through "labs" that enable them to experiment with generating alternative distributions and comparing outcomes to sample histograms. After deriving the theoretical moments of a probability distribution and relating them to the parameters of the distribution, students will perceive a triangle involving, for a given distributional class, the shape of the distribution, the parameters, and the four moments. This analysis provides a clear link between the lessons learned in Part Two on sample moments and the later chapters on inference in Part Four. In Chapter 7, I introduce expectation as a generalization of the calculation of theoretical moments and explore many of its interesting properties.

Chapter 9, Elementary Sampling Theory, provides the essential link between the theory developed in Part Three and the actual observations discussed in Parts One and Two. Sampling theory is the mechanism by which theory is used to provide the necesary interpretation of observed data, and it provides the rationale for the inferences drawn in Chapters 10 and 11 on estimation and hypotheses testing, respectively.

Chapters 10 and 11 focus on the properties of estimators and the principles of hypothesis testing. These chapters contain a few unique features. One, for example, is to discuss why the 10%, 5%, and 1% test sizes in tables were chosen in the first place and that with modern computers we do not need to be bound by such conventions. Indeed, this text contains no formal tables, except as examples, because all the information normally contained in the numerous tables can more easily be obtained from computer algorithms. Another unique subject is the trade-off between Type I and II errors and the role of one's preferences in making these decisions. Experiments enable readers on their own to observe the relationship between the size of a hypothesized difference, the sample size, the chosen α level, the power, and the size of the error variance in formulating hypotheses tests.

Chapters 13 and 14, which discuss regression and ANOVA, respectively, are prefaced by a unique discussion on the generation of bivariate and conditional distributions in Chapter 12. Conditional probability is given a lot of weight because it provides the basis for the theory of regression and ANOVA. The insights provided in Chapter 12 facilitate understanding the principles of regression and the nebulous connection between regression and causality. It is at this stage that the more conventional view of statistics as a way of modeling "models observed with error" comes to the forefront. The role of the model in the analysis and the effect it has on drawing inferences from the data are emphasized. At this stage, students can begin to appreciate through simple examples the intimate interaction between probability theory and the formulation of scientific models.

The last chapter, "Retrospective," is novel. Its objective is to review in broad terms the basic concepts of the text and to readdress the issues raised in the first two chapters concerning the role of statistics in decision making, policy implementation, and science. Chapter 15 attempts to enable students to begin to appreciate the enormous gains in understanding that derive from applying the theory of statistics to observations of data. Interpretation is the key to understanding, and interpretation comes from our mastery of the underlying principles of inference.

A single case study is carried throughout the text; every chapter contains a reference to the case study and indicates how the contents of the chapter relate to the problem posed. I hope by this device to provide for students a motivating example, a unifying framework, and the ability to observe within a single context the development of the tools discussed in the text. This device also will enable students to observe the growth in their knowledge and understanding.

The Roles of the Exercises

The Exercises play several distinct roles in this text. Much material, useful insights, and interesting applications are in these exercises. If you do not see a favorite procedure in the main body of the text, look in the index and you may well find it in one of the exercise problems. The Exercises challenge better prepared students, leading them to investigate developments of the material and examine the ideas with greater rigor.

Learning is facilitated by including worked examples in the Exercises. Each worked example is headed by a statement of its objective. These exercises extend and illustrate the material and enable students to experiment on their own. Statistics can only be learned through doing and examining the repetition of exercises that exemplify stochastic variation in all its forms. This is achieved in a seamless manner in this text.

The exercises contain a far greater number of worked examples than is customary. These worked examples provide students with a detailed set of instructions for carrying out a procedure; a set of examples for answering the various types of questions normally placed in the text; and the tools for enabling them to experiment with the concepts. When assigning exercises, look for the nearest worked example that will provide a "blueprint" for doing the exercise that you have assigned. Routinely assigning the worked examples is a good practice; there is much to be learned in them. In Section l.7, I have instructed students to pay particular attention to the worked examples. Indeed, if no other exercises are done, these should be. The benefits from working through them cannot be exaggerated. In this connection, the instructor is recommended to bring Section 1.7 to his or her students' attention, because it will aid their study efforts dramatically; my own students declare that it is "a must read again and again" section.

The exercises attached to each chapter are in three sections: Calculation Practice, Exploring the Tools, and Applications. Each section provides training in a different skill.

Calculation Practice

The first section, "Calculation Practice," contains problems designed to give students manipulative skills and facility in calculation. The initial exercises in this first set involve very simple numbers so that students can easily carry out the calculations by hand. It is only after gaining some feel for the calculations that the exercises move to using the computer and more "realistic" numbers. The idea is that to understand what is involved in statistical calculations students must have some manipulative practice themselves before they can understand what the computer is doing when used with more realistic data.

The numbers are kept very simple at first, so that "arithmetic" is no barrier to understanding the objective. The computer is next introduced with easy-to-use menu commands to facilitate using more realistic numbers. The next stage for more complicated formulae, such as that for the correlation coefficient, is to pull together the individual terms in the correlation coefficient "by hand." This is done before allowing students to use a packaged computer routine and then not until they have verified that the packaged routine does indeed produce the same numerical result as their own calculations.

The text contains an easy-to-use program on a CD that is a simplified student version of S-Plus. S-Plus comes with a series of laboratory routines as well as a fully developed statistical analysis package with extensive graphing capabilities. Almost all of the calculations, graphing, and routines needed by students can be accessed through a menu-driven GUI interface. S-Plus has been chosen as the main vehicle for using the computer to carry out the calculations because it provides excellent graphics; students can generate their own distributions and histograms; and they can be taught very easily how to *use only those tools that are needed at the moment that they are needed.* Some summary instructions for using S-Plus are contained in Appendix B.

That I chose to use S-Plus as the vehicle for the programming should be irrelevant to the instructor and the student. I have set up the computer procedures so that readers can quickly and effortlessly carry out the exercises with a minimum of instruction. In no sense must students spend time "learning some computer program" that they may never see again, even though S-Plus is a superb tool for the professional statistician as well as for the neophyte.

I have tried to prevent the use of the computer being a barrier to understanding and a source of both complexity and frustration, as is so often the case. I also wanted to make sure that if students perform any calculation *exactly* as I have done, they are guaranteed to obtain *precisely* the same result as that indicated in the text. I have restricted the use of the available algorithms to a very limited set because I am reluctant to expose students to the use of packaged computer programs to perform complex operations, such as regression analysis or ANOVA, until they have done it by hand a few times. Black boxes of any type are to be avoided because they lead the student into mindless manipulations of data with no intuition for what is involved. The manner in which the computer algorithmns have been implemented means that neither the instructor nor the student need be aware of what the underlying routines are; the computer has become an extension of the student's analysis and merely provides shortcuts to tedious routine calculations, each of which is thoroughly understood.

Exploring the Tools

The second set of questions, "Exploring the Tools," is designed to help students explore the properties and limitations of the statistical tools learned in each chapter. For example, students should acquire a feel for the effect from adding or deleting variable values on the calculation of sample moments or for when the distribution of values is special. This is the section in which the student is encouraged to experiment in trying to gain a deeper understanding of the concepts. "Labs" here provide an exceptional introduction to random variation and statistical analysis. Every distribution, statistic, and test discussed in the text can be simulated and its statistical properties examined by computer.

My intent is to provide a vehicle for engaging students' interest and enabling them to develop insight into statistical concepts. As it will be very easy for students to experiment with the labs on their own, I am confident that this is what they will do.

Applications

The last section of questions, "Applications," provides students with applications of the tools. There are some unique features here as well. For example, greater stress is placed in this section on the formulation of the problem, the questions that have to be answered to proceed with the analysis. Most of the questions concentrate on the *interpretation* of the results obtained from the statistical procedures, rather than the mere mechanical application of the procedures.

In large part, the questions posed here address the issues normally ignored, or assumed away, in the standard textbook question, which often restricts its attention to the mechanics of substituting numbers into formulae and reducing the resulting expression. The guiding principle in this text is something like "If you were employed as a statistician for a firm, or an agency, what sort of questions would your boss pose to you?" Most such questions would state the problem and ask what data we need to solve it and how we proceed. The real question is, "What is the statistical question?" To enhance this questioning attitude, questions are sometimes asked in the later chapters that do not require use of any of the tools in the current chapter; the idea is to encourage the student to think about the problem before writing down a formula.

Finally, the remaining exercises in the "Applications" section relate to the case study discussed at the end of every chapter. Students should be able to see the relevance of the chapter's contents to the case study and, by concentrating on a single case study throughout the text, recognize the context and the links to the concepts presented in other chapters. Students will quickly begin to perceive the growth in their knowledge of statistical thinking.

Why I Wrote the Text This Way

My reasons for developing the text in this way arose out of my frustration in observing the teaching of statistics with seemingly little long-term impact, even for students who obtained good grades. Examining the issue, I discovered that most students studied statistics by engaging in a sophisticated form of memorization; that is, they would memorize the formulae and create rules of thumb for the insertion of the formulae into word problems. This enabled them to earn a good grade in most statistics classes and without understanding very much at all. Within six weeks of the end of the course, the formulae are forgotten and the whole exercise has been a waste of everyone's time. Statistics is not the only discipline that tends to be taught in this manner for the nonmajor; much of mathematics is taught in the same way and with equally depressing results. For example, I have run experiments for years on the mathematical knowledge of students and found that while they can instantaneously shout out the formulae for common derivatives and integrals, the simplest problem defeats them if not put into the context that they had in their mathematics class.

I feel that many modern statistical texts encourage memorization and treat "the study of statistics" as a sequence of cookbook recipes. I appreciate the fact that for many instructors their chosen approach to teaching statistics is to provide a series of statistical recipes. The argument is that an introductory undergraduate course cannot do anything else because the students have neither the knowledge nor the skills to follow a more insightful course. Although I can appreciate the concerns and constraints of such instructors, whose constraints often include the stricture to "cover regression analysis at any cost," I believe that such an approach is ultimately neither challenging to the student nor rewarding to the instructor. The alternative approach is to be prepared to cover less material but to cover the material that is included more intensively.

This text is for those instructors who have been looking for a text that tries to stress the understanding of the basics and the development of "statistical intuition." This is a text for the instructor who is willing to give up covering a long list of "applied topics" in exchange for a greater appreciation by students of the elementary ideas and the fundamental logic underlying all the applications that they are ever likely to meet. Although it is true that in using a text of this type fewer applications will be covered, I trust that what is covered will be more intuitively appealing to students. I am hopeful that students will be encouraged not to memorize undigested formulae as the only way to pass statistics tests but to recognize that the material can be understood, and that there is a logical and understandable explanation for everything that is done.

A very simple example will illustrate my intent and explain an aspect of the text's organizational structure. In most texts, one statistic that is presented very early is the sample variance, s^2. Before s^2, the mean has been presented and, if other moments are mentioned later, they will be defined in terms of division by sample size, n, not $(n-1)$. The student is sometimes told, he "will learn the reason for this exception later," or there is a footnote to the effect that the problem has something to do with "degrees of freedom." But most students are unlikely to understand these comments. Here is yet another "mysterious formula" that has to be memorized, because its logic escapes the student.

In this text I take a different approach. Although I am developing sample moments as measures of properties of histograms, the second moment is no exception to the general rule of moments; $m_2(x)$ is defined by division by n, just as is the case for the first and all other moments. The pattern makes sense; there are no strange exceptions. But, when I get to Section 7.5 on expectations where the student learns that $m_2(x)$ is biased for σ^2, it is not only appropriate at that time to tell students about s^2, they can discover it for themselves.

The Mathematics Requirements for the Text

The formal mathematics required in the text are a little beyond basic algebra to begin and build from there. Calculus concepts are used, but little in the way of formalism. The intent is to build intuition about the relevant concepts, even without a careful explanation of the subtleties of mathematical analysis. I start with easy material and provide students with the facility to improve their manipulative skills. As students proceed through the text and gain confidence in their ability to handle the mathematical concepts and tools, the requirements are raised. To avoid letting students fall into the mindless manipulation of symbols without understanding, developing intuition into the mathematics used is stressed. Little is required beyond what a student should have on entering a first calculus course, and certainly no more than would be known by the end of a first calculus course. Currently, most students in economics or business statistics courses have had at least one calculus course.

The manner in which the mathematics are developed in the text reflects my observation that for the vast majority of students the weak link is algebra. In short, students have relatively less difficulty with calculus concepts than they do with algebraic concepts and procedures. Consequently, I have spent some considerable space in the text on the development of the algebraic tools that are needed and relatively less on the review of calculus concepts. Those with a strong mathematics background might feel that the text is too easy in the beginning relative to the level at which it ends. However, given the uneven nature of students' preparation—that is, too little familiarity with algebra—this is not so.

All the mathematics used in the text are explained in Appendix A. Exercises are provided on all the material both in the Exercises at the end of each chapter and in Appendix A. The introduction starts at a low level that should be easy for all students, although "Σ notation" seems to be a perennial difficulty. Further, there may be gaps in a particular student's training, so that beginning slowly with a comprehensive review should be an encouraging introduction. In addition, by beginning slowly, those who are likely to be mathematically challenged by the later portions of the text will have had the opportunity to gain some confidence before tackling the tougher sections. The material builds throughout the text in complexity to include basic ideas of *continuity*, *limit*, *differentiation*, and *integral*, mainly as the concept of the "area under a curve." Little formalism is invoked but intuition is stressed.

The rationale for this approach is that the material to be taught necessitates some calculus intuition. Some understanding of the concepts of limits, continuity, differentiability, and integration are essential for students to grasp the idea of a density function, the relationship between the density function and the probability function, the meaning of expectation, and so on. However, in contrast to some approaches to the study of statistics, this material can be made intelligible without a heavy reliance on "full mathematical rigor." If the student goes on to study statistics more seriously and therefore in greater depth, then at that time there will be a need for enhanced rigor. However, I believe that such a student will proceed much faster and with less difficulty in a more mathematically demanding course, having already obtained the basic insights that I try to inculcate in this text.

Acknowledgments

The author of any complex technical book inevitably owes a debt of gratitude to many people, and I am no exception. Numerous graduate students and undergraduates have contributed to the final form of the text to the ultimate benefit of the reader. To all of these, I owe my heartfelt thanks.

However, there are some key individuals who have contributed in special ways to the publication that has been so long in coming. A long-standing thanks is due to George Lobell, an editor of great talent, who originally urged me to write the text and who patiently waited for many years while research took precedence over completion. I hope that he feels that the wait has been adequately rewarded.

Clearly, we are all in the debt of Joe Newton and Jane Harvill, who wrote the labs that are such an important component of the development of the material.

It is customary to thank one's spouse, but in my case, my heartfelt thanks cannot adequately express the debt that I, and the reader, owe Kate Ramsey. Not only did she keep me focused on completion and alleviate my frequent disappointments, but she was also instrumental in creating the computer instructions for all the exercises. She pulled together and edited the numerous data sets and spent hours proofreading the text for accuracy and consistency of notation. She never let imprecise statements, especially in the exercises, go by without challenge. Quite literally, the text would not have been finished without her help and support.

Three graduate students in particular played a critical role in editing the text and creating many of the exercises and their solutions—Alannah Orrison, Enrique Schroth, and Guillermo Felices. I am deeply grateful for their imaginative contributions. Very special thanks also go to Eliane Catilina, an instructor at the University of Virginia, who made useful comments and was one of the first instructors to use the text in the classroom. Because of her great teaching talent, the student response to the text was a most successful experiment.

At a late stage in writing the manuscript I added the Bayesian supplement (on the accompanying CD) and was fortunate to persuade Arnold Zellner to read it through. The reader will benefit from Arnold's insightful comments, even though I was not able to incorporate all his excellent suggestions. My thanks to Arnold for his continuing support of my academic efforts over these many years. I am also grateful for insightful comments and the case study provided by Elias Grivoyannis.

I am highly indebted to the following reviewers who reviewed drafts of the manuscript and made many helpful suggestions: Hongshik Ahn, SUNY Stony Brook; Anthony D. Becker, St. Olaf College; Michael A. Boozer, Yale University; Bruce Cooil, Vanderbilt University; Jerry Dwyer, Clemson University; Mark Eakin, University of Texas, Arlington; Yongmiao Hong, Cornell University; James Stapleton, Michigan State University; Ebenge Usip, Youngstown State University; Keith Womer, University of Mississippi.

Lastly, but essential for publication, I would like to express my appreciation to my editor, Curt Hinrichs. Curt had the courage to support a book that was not designed for the mass market. He was instrumental in bringing together all the complex interactions between text, labs, computer instructions, data, computer software, supplemental chapters, web sites, and the creation and provision of compact discs bound with the text to a successful resolution; no mean feat.

I hope only that the efforts of all these talented people will be rewarded by the pedagogical success of the text.