

The Introductory Statistics Course: The Entity-Property-Relationship Approach

Donald B. Macnaughton*

This paper proposes five concepts for discussion at the beginning of an introductory statistics course: (1) entities, (2) properties of entities (which are roughly equivalent to variables), (3) a major goal of empirical research: to predict and control the values of variables, (4) relationships between variables as a key to prediction and control, and (5) statistical techniques for studying relationships between variables as a means to accurate prediction and control. After students have learned the five concepts they learn standard statistical topics in terms of the concepts. It is recommended that students learn the material through numerous practical examples. It is argued that the approach gives students a lasting appreciation of the vital role of the field of statistics in empirical research.

KEY WORDS: Statistical education; Prediction; Control; Role of statistics in empirical research.

1. INTRODUCTION

Two former presidents of the American Statistical Association have stated that “students frequently view statistics as the worst course taken in college” (Hogg 1991, Iman 1994). A third former president has stated that the field of statistics is in a “crisis” and the subject has become “irrelevant to much of scientific enquiry” (Box 1995). Many statisticians reluctantly agree with these remarks.

In contrast, many statisticians agree that the field of statistics is a fundamental tool of the scientific method which, in turn, plays a key role in modern society. Thus rather than being a worst course and possibly irrelevant, the introductory statistics course ought to be a friendly

introduction to the simplicity, beauty, and truth of the scientific method.

Teachers must therefore reshape the introductory course. Many teachers have already contributed to the reshaping, as noted below. This paper proposes further changes.

I focus on the introductory statistics course for students who are *not* majoring in statistics or mathematics, whom I call “non-statistics-majors”. Most students who take an introductory statistics course are members of this group. The introductory course for non-statistics-majors is important because it is the main seedbed of public opinion about the field of statistics.

Section 2 recommends two goals for the introductory statistics course. Section 3 proposes a sequence of five concepts for discussion at the beginning of an introductory course. Section 4 argues that the five concepts provide a broad and deep foundation on which to build the field of statistics. Section 5 reports on tests of the approach in three courses. Section 6 discusses considerations for teachers wishing to use the approach, and section 7 gives a summary.

2. COURSE GOALS

2.1 The Value of Emphasizing Goals

Emphasizing the goals of any undertaking helps us to define and focus on *what is most important*. Otherwise, we may take a scattershot approach, which is generally less efficient. Thus, following Hogg (1990, 1992), I recommend that all introductory statistics courses have carefully stated goals. Ask yourself—what are the goals of the introductory statistics course most familiar to you?

I have observed goal-setting exercises in which the goals were given much attention for a brief period and then forgotten—a waste of a valuable resource. Since a teacher’s chosen goals are (by the teacher’s own definition) what is most important, I recommend that teachers regularly revisit their course goals to ask (a) if the goals are still valid and (b) if the day-to-day operations are effectively serving the goals.

2.2 A Definition of “Empirical Research”

In recommending goals for the introductory statistics course I refer to “empirical research”, which thus deserves a definition:

Empirical research is any research in which data are gathered from the external world and then conclusions are drawn from the data about the external world.

*Donald B. Macnaughton is president of MatStat Research Consulting Inc, 246 Cortleigh Blvd., Toronto Ontario, Canada M5N 1P7. E-mail: donmac@matstat.com Portions of this material were presented at the Joint Statistical Meetings in San Francisco, August 11, 1993, at the Joint Statistical Meetings in Orlando, August 15, 1995, and at the Joint Statistical Meetings in Dallas, August 12, 1998. The author thanks Professor John Flowers of the School of Physical and Health Education, University of Toronto and Professors Donald F. Burrill and Alexander Even of The Ontario Institute for Studies in Education for testing the approach in their statistics courses. Their comments and their students’ comments helped substantially to clarify the ideas. The author also acknowledges insightful comments and criticisms from David R. Bellhouse, Carol J. Blumberg, Christopher Chatfield, George W. Cobb, Sir David Cox, W. Edwards Deming, David J. Finney, Robert W. Frick, Joan B. Garfield, Robert V. Hogg, Olaf E. Kraulis, Gudmund R. Iversen, William H. Kruskal, Alexander M. Macnaughton, Christine E. McLaren, David S. Moore, Thomas L. Moore, Jerry Moreno, John A. Nelder, Ingram Olkin, Allan J. Rossman, Richard L. Scheaffer, Milo A. Schield, Stephen Senn, Gary Smith, and Paul F. Velleman.

Empirical research is a key step of the scientific method, which is crucial in many areas of human endeavor, such as in science, education, business, industry, and government. (No statement of fact in any branch of science is accepted until it has been verified through careful empirical research.)

2.3 Topic-Based Goals Have a Serious Drawback

Many introductory statistics courses have what can be called “topic-based” goals. A teacher using such goals does not specify general goals, but instead simply specifies a list of statistical topics to be covered in the course—that is, the teacher specifies a syllabus. For example, a teacher of a traditional introductory course might aim to cover (in specified amounts of detail) the topics of probability theory, distribution theory, point and interval estimation, and so on. Similarly, a teacher of an activity-based course might make a list of statistical topics and then assign various activities to the students in order to cover the topics.

Unfortunately, topic-based goals have a serious drawback: By emphasizing lower-level statistical *topics*, these goals usually *fail* to emphasize what is essential, which is to help students to appreciate the vital *role* or *use* of the field of statistics. Unless students understand and appreciate the general role of statistics, knowledge of statistical topics will be both of little interest to them and of little use.

2.4 Recommended Goals

I recommend that the goals of an introductory statistics course be

1. to give students a lasting appreciation of the vital role of the field of statistics in empirical research and
2. to teach students to understand and use some useful statistical methods in empirical research.

These goals imply a commitment to the *practical application* of statistics in empirical research. Emphasis on practical applications is important because the social value of statistics is through its practical applications, and not (directly) through theory or mathematics. Thus we stand a much better chance of impressing students with the usefulness of statistics if we discuss practical applications.

(We must also discuss generalizations that link statistical ideas together, but [to maximize appreciation] practical applications deserve strong emphasis. More on generalizations below.)

The first goal refers to the *role* of statistics in empirical research. Clearly, to satisfy this goal, we need a precise sense of what the role is. I propose a definition of the role in subsection 4.7. I discuss in a paper how we might *measure* students’ appreciation of the role of statistics as an aid to experimenting with different approaches to the introductory course (1998a, app. A).

Other writers who discuss goals for the introductory statistics course include Hogg (1990, 1991), Chromiak, Hoefler, Rossman, and Tesman (1992), Cobb (1992), Iversen (1992), Watkins, Burrill, Landwehr, and Scheaffer (1992), Hoerl and Snee (1995), Gal and Garfield (1997a, pp. 2 - 5), and Moore (1997a).

3. FIVE CONCEPTS

To help define the role of statistics, let us consider a sequence of five concepts I recommend for discussion at the beginning of an introductory statistics course. To make the approach easy for teachers to use, I present each concept mainly as a highly condensed version of how it might be presented to students.

(In this paper I discuss the five concepts mostly in abstract terms because this gives readers a compact overview of the recommended approach. The abstractness of my discussion has led readers of an earlier version of the paper to criticize the approach as being too abstract for students to understand [Macnaughton 1998a, app. C]. Please note that I am not suggesting that we present the concepts to students in the compact form in which I describe them below. Instead, I recommend that the concepts be presented in terms of numerous practical examples. I further discuss presentation issues in section 6 and I illustrate how the ideas might be presented to students in a paper [1996a].)

Let us begin with what may be the most fundamental concept of human reality.

Concept 1: Entities

If you stop and observe your train of thought at this moment, you will probably agree that you think about *things*. For example, during the course of a minute or so, you might think about, among other things, a friend, an appointment, today’s weather, and an idea. Each of these things is an example of an *entity*.

Many different types of entities exist, for example

- organisms (e.g., people)
- inanimate physical objects
- events
- processes
- ideas, emotions
- societal entities (e.g., educational institutions, governments, businesses)
- symbols
- forces
- waves
- mathematical entities (e.g., sets, numbers, vectors).

Entities are fundamental units of human reality because everything in human reality is an entity.

People usually view entities as existing in two different places: in the external world and in our minds. We use the entities in our minds mainly to stand for entities in

the external world, much as we use a map to stand for its territory.

All sane people learn to use the concept of ‘entity’ when they are infants. We use the concept unconsciously as a way of organizing the multitude of stimuli that enter our minds minute by minute when we are awake.

Infants recognize that entities can be grouped into ‘types’. For example, they recognize that the entities ‘‘mother’’ and ‘‘father’’ and other similar entities in the external world (e.g., siblings) have heads with two eyes, a mouth, and usually hair on top. Infants group these entities into a type—the type we call ‘‘people’’. Similarly, infants learn to group all inanimate physical objects (beginning with small familiar objects) into a type. Infants (unconsciously) recognize that all the entities of a given type, although each is unique, have many things in common, as I discuss further in the next subsection.

Since everything in human reality is an entity, the concept of ‘entity’ pervades human thought. However, the concept is rarely present in our *conscious* thought because the concept is not often directly discussed in everyday conversation, or in empirical research, or in statistics. Direct discussion of the concept of ‘entity’ is usually omitted because, in dealing with specific issues, we rarely need to drill down all the way to the foundational concept to discuss ‘‘things’’ at such a basic level. Instead, discussions usually concern one or more specific *types* of entities, which are best referred to by their type names. (For example, medical researchers study two familiar types of entities called ‘‘human beings’’ and ‘‘diseases.’’) However, as we shall see, the concept of ‘entity’ serves as a foundation for other important concepts. Therefore, it is helpful to make the concept explicit.

At those times when entities *are* discussed in general terms in empirical research or in statistics, they may also be called members of the population, cases, individuals, instances, items, specimens, subjects, or (experimental or survey) units.

Because people use the concept of ‘entity’ almost entirely at an unconscious level, some people have difficulty grasping the fundamental role the concept plays in their thought.

Appendix A further discusses the priority of the concept of ‘entity’ in human thought, statistics, and science.

Concept 2: Properties of Entities (Variables)

Every entity has associated with it a set of attributes or *properties*. For example, all people have thousands of different properties, two of which are ‘‘height’’ and ‘‘blood group’’.

For any particular entity, each of its properties has a *value*. We usually report the value of a property with words, with symbols, or with numbers. For example, your height might be 172 centimeters, and your blood group might be A.

If we need to *know* the value of a property of an entity, we can apply an appropriate measuring instrument *to* the entity. If the instrument is measuring properly, it will return a measurement to us that is an estimate of the value of the property in the entity at the time of the measurement. For example, if we need to know the (value of the) height (property) of a person, we can apply a height-measuring instrument (e.g., a tape measure) to the person, and the instrument will return a value that is (in the specified units) an estimate of the person’s height.

I note above that people group entities into types, such as ‘‘human beings’’ and ‘‘inanimate physical objects’’. We (unconsciously) view all entities of any given type as having *exactly* the same properties although, of course, the *values* of the properties generally differ from entity to entity. For example, *all* human beings have the properties ‘‘height’’ and ‘‘blood group’’, but your values of the two properties probably differ from mine. By viewing all the entities of a given type as having exactly the same properties, we capitalize on a key feature of the external world and thereby greatly simplify our model of the external world.

Empirical researchers and statisticians usually refer to properties of entities as *variables*. That is, whenever empirical researchers or statisticians refer to a variable, they are usually referring (either specifically or generally) to some property of some type of entity. A reasonable definition of the concept of ‘variable’ is

A *variable* is a formal representation of a property of entities.

Interestingly, the concept of ‘variable’ is so ubiquitous and fundamental in statistics that some introductory courses and textbooks take the concept almost completely for granted, with little formal discussion of it. By assuming students have a satisfactory understanding of the fundamental concept of ‘variable’, these approaches befuddle a significant proportion of students right from the start.

I discuss the concept of ‘variable’ further in the paper for students (1996a) and in a Usenet post (1996b).

Concept 3: A Goal of Empirical Research: To Predict and Control the Values of Variables

An important goal of empirical research is to discover how to *predict* and *control* (with maximum accuracy) the values of properties of entities. In other words, the goal is to predict and control the values of *variables in* entities. For example, a goal of medical science (i.e., empirical research in medicine) is to discover how to accurately predict and control the state of the human body, where the state is reflected by various medical properties or variables, such as blood pressure, white blood count, and other measures of health or disease.

It appears that *all* prediction and control can be easily viewed as prediction and control of the values of proper-

ties of entities or (equivalently) prediction and control of the values of variables.

Society supports empirical research in seeking the ability to predict and control the values of variables because instances of such ability often provide substantial social or commercial benefits. For example, if medical researchers can discover how to better predict and control a person's propensity to heart attacks, this discovery will provide the social benefit of saving lives.

(From this point on in this paper I emphasize the concept of 'variable' and I refer less often to the closely related concept of 'property of an entity'. It is important to emphasize the concept of 'variable' because this concept is used throughout empirical research. This leads to an obvious question: "Why not discuss variables right from the start—is the extra concept of 'property' necessary if it is only to be supplanted by the concept of 'variable'?" My answer is that the concept of 'property' [or, equivalently, 'attribute' or 'characteristic'] is much more basic and intuitive in students' minds than the concept of 'variable'. Thus students understand the concept of 'variable' better if we carefully build it atop the intuitive foundational concept of 'property of an entity'.)

Concept 4: Relationships Between Variables as a Key to Prediction and Control

Given the goal of predicting and controlling the values of variables, an obvious question is

How can we predict and control the values of variables?

The main answer is

We can predict and control the values of variables by studying *relationships between* variables.

By a "relationship between variables" I mean the standard idea that one variable (called the *response* variable) "depends" on one or more other variables (called the *predictor* variable[s]). All statisticians are familiar with this unifying idea, although it is generally not emphasized.

Almost all prediction and control in all areas of empirical research is done on the basis of relationships between variables. For example, medical researchers have discovered that a relationship exists between the amount of fat ingested by a person and the probability that a person will have a heart attack. This relationship enables doctors and patients to predict and control heart attacks.

We can help students to appreciate the pervasiveness of the concept of a relationship between variables by discussing numerous practical examples of relationships. For example, teachers and students can discuss whether a relationship exists between

- the weight of a car and the gas mileage of a car (engineering)
- the education of a person and the income of a person (sociology)

- the net amount of force applied to a physical object and the rate of acceleration of a physical object (physics)
- the concentration of alcohol in the bloodstream of a person driving a car and the probability that a person driving a car will be involved in an accident (physiology)
- the number of hours of homework completed by a student during a course and the grade obtained by a student in a course (educational psychology)
- the style of management of a company and the amount of profit of a company (business)
- having a black cat cross one's path and having bad luck (folk beliefs).

Each of the above examples identifies a possible relationship between variables. Each of these relationships (and relationships between any other pairs or larger sets of variables) can be studied in an empirical research project. Each research project can be viewed in terms of

1. a population of entities
2. a sample of entities that is selected from the population
3. a response variable that is measured in each entity in the sample
4. one or more predictor variables that are also measured or controlled in each entity in the sample (or possibly are measured or controlled in each entity's environment)
5. (most importantly) a relationship in the entities in the population between the response variable and the predictor variable(s) that is sought or studied in the research project.

Typically we ask three main questions in an empirical research project, which are

1. Is there a relationship between the response variable and the predictor variable(s) in the entities in the population?
2. If there is a relationship, how can we best predict or control the values of the response variable in new entities from the population on the basis of the relationship?
3. How accurate will the prediction or control be?

By considering a broad set of practical examples, we can show students that most empirical (including most scientific) research projects can be easily and usefully viewed in terms of the five points and three questions given above. That is, most empirical research projects can be viewed as studies of relationships between variables, with the aim being to develop the ability to predict or control the values of the response variable in new entities from the population. This approach provides a simple yet comprehensive model of most empirical research.

(I provide support for the claims in the preceding paragraph in appendix B.)

I recommend that students' initial sense of the concept of a relationship between two variables be in terms of

- the value of one variable in entities "depends" on the value of another variable and

- as the value of one variable changes in entities, the value of the other variable changes somewhat in step with the values of the first variable.

* * *

In addition to showing students that the concept of a relationship between variables plays a central role in empirical research, we can also show mathematically minded students how the concept plays a central role on the theoretical side of science. Specifically, mathematical equations are crucial on the theoretical side of many branches of science. But most (all?) mathematical equations in science (as opposed to equations in pure mathematics) are simply statements of known or hypothesized relationships between variables (relationships between properties of entities).

In the initial discussion of concepts 1 - 4 with students, I recommend that a teacher *not* introduce any statistical concepts beyond the four concepts themselves and some graphs to illustrate relationships. I recommend this approach because I believe it is important for students to develop a good sense of the concept of a relationship between variables *before* they try to understand the complicated statistical ideas behind the study of relationships.

Subsection 6.5 further discusses the use of practical examples in an introductory statistics course. The paper for students gives a formal definition of the concept of a relationship between variables (1996a, sec. 7.10).

Concept 5: Statistical Techniques for Studying Relationships Between Variables

Once students properly understand and appreciate the usefulness of relationships between variables as a means to prediction and control, we can then bring the field of statistics out onto the stage. We can characterize the field as a set of optimal techniques to help empirical researchers study variables and relationships between variables as a means to accurate prediction and control.

When first introducing the field of statistics, it is helpful to classify statistical techniques into four functional groups

1. techniques for *detecting* relationships between variables
2. techniques for *illustrating* relationships between variables
3. techniques for *predicting* and *controlling* the values of variables on the basis of relationships between variables, and
4. miscellaneous techniques for the study of variables and relationships between variables.

The field of statistics provides many methods for carrying out these four groups of techniques, with the choice of the best method(s) depending on the particular empirical research situation at hand. I discuss some of these methods in subsection 4.3 and in appendix C. I further

discuss the four groups of techniques in the paper for students (1996a, sec. 8 - 13).

After introducing the four groups of techniques, I recommend that a teacher spend the rest of the course and subsequent courses introducing standard statistical principles and methods in terms of the techniques. I propose a syllabus for a course following this approach in subsection 6.4.

Summing Up

The preceding subsections describe a sequence of five concepts for discussion at the beginning of an introductory statistics course. The concepts are

1. entities
2. properties of entities (variables)
3. an important goal of empirical research: to predict and control the values of variables
4. relationships between variables as a key to prediction and control
5. statistical techniques for studying relationships between variables as a means to accurate prediction and control.

After introducing the five concepts, the teacher focuses the rest of the course on statistical techniques for studying variables and relationships between variables.

I call this approach to the introductory statistics course the “entity-property-relationship” (EPR) approach. I discuss evaluation of the approach in section 4, tests of the approach in section 5, and implementation considerations in section 6.

4. EVALUATION OF THE EPR APPROACH

The entity-property-relationship approach to the introductory statistics course has several important features and benefits, which I discuss in this section.

(Some of the material in this section is abstract and is presented solely to assist in evaluating the EPR approach. I do not intend that this abstract material be discussed in an introductory course.)

4.1 The Concepts of the Approach Are Easy to Understand

The concepts of entities, properties (variables), and relationships are ubiquitous in students’ (unconscious) thought. Therefore, if we carefully discuss these concepts (with sufficient practical examples), students find the concepts easy to understand.

The ease of understanding leads me to conjecture that the concepts of entities, properties, variables, and relationships can be taught at all levels of teaching statistics from late elementary school up, with only the teaching time and depth of coverage of the concepts varying at different levels.

4.2 The Approach Provides a Broad and Deep Foundation for Discussion of Statistical Concepts

The second of the five concepts in section 3 introduces the fundamental concept of ‘variable’ in terms of the concepts of ‘entity’ and ‘property’. The fourth concept introduces the fundamental idea of a relationship between properties (relationship between variables). It is easy to see that several other fundamental statistical concepts are built atop the concepts of ‘entity’, ‘property’, or ‘relationship’. For example

- *set*: a fundamental type of entity in human reality that consists of zero or more entities of a specified type
- *population*: a “large” set of entities upon which interest is focused
- *sample*: a set of entities selected from a population, or a set of values of properties of entities selected from a population
- *event*: a fundamental type of entity in human reality
- *probability*: a particular property of an event that reflects how often the event occurs or is thought likely to occur
- *distribution*: a mathematical entity that succinctly describes in probabilistic terms what the different values of a property are in the different entities in a population or sample
- *model (equation)*: a mathematical statement of a relationship between properties of entities
- *parameter*: a general name for a property of a population, distribution, or model equation (or for the *value* of such a property)
- *estimation* (of the value of a property or the value of a parameter)
- *hypothesis* ([*a*] of the existence of a relationship between properties or [*b*] about the value of a parameter)
- *statistical test*: a technique for providing an objective measure of the weight of a body of evidence in support of a hypothesis
- *statistic*: any of various well-defined properties of a sample whose values are obtained by performing mathematical operations on the values of one or more properties of the entities in the sample; often used to estimate the values of parameters or in performing statistical tests.

Note that each of the fundamental statistical concepts in the above list is built atop the concepts of ‘entity’, ‘property’, or ‘relationship’, or is built atop concepts that are themselves built atop the three concepts. Furthermore, the concepts of entities, properties, and relationships appear to be among the most basic concepts of human reality. Thus the EPR approach provides a broad and deep foundation for discussion of statistical concepts.

4.3 The Approach Unifies the Main Statistical Methods

The discussion of concept 5 in section 3 identifies four groups of techniques that statistical methods can perform to help empirical researchers study relationships between variables, namely

1. techniques for *detecting* relationships between variables
2. techniques for *illustrating* relationships between variables
3. techniques for *predicting* and *controlling* the values of variables on the basis of relationships between variables, and
4. miscellaneous techniques for the study of variables and relationships between or among variables.

This raises the question: Which of the currently available statistical methods can actually perform these four groups of techniques? The following nineteen statistical methods can perform one or more of the foregoing four groups of techniques:

- general linear model (*t*-test, analysis of variance, linear regression, multiple comparison methods, hierarchical methods, variance components analysis, multivariate analysis of variance, multivariate linear regression)
- generalized linear model
- response surface methods
- exploratory data analysis
- time series analysis
- survey analysis
- survival analysis
- categorical analysis
- graphical methods
- Bayesian methods
- nonlinear regression
- neural networks
- discriminant analysis
- nonparametric methods
- probit analysis
- logistic regression
- correlation analysis
- structural and path analysis
- univariate analysis.

Upon consideration, many statisticians will agree that the above list of statistical methods contains almost all of the currently popular statistical methods. Many statisticians will also agree that the *only* techniques that the statistical methods in the list can perform are given in the four groups of techniques at the beginning of this subsection.

(Furthermore, statistical methods can usually be characterized as performing only the *first three* groups of techniques—the fourth group is only rarely needed to characterize the use of statistics in actual empirical research. I provide support for the claims in the preceding

paragraph and I explain why certain statistical methods are excluded from the list in appendix C.)

Since the nineteen statistical methods in the list are fully explained (at a high level) by the four groups of statistical techniques that are emphasized in the EPR approach, the approach unifies the main statistical methods. That is, we can *interpret each new statistical method* we teach to students *in terms of the same set of simple concepts*: entities, properties, variables, and relationships between variables. This unification makes the field of statistics substantially easier for students to understand.

4.4 The Approach Links Well With Scientific Explanation

I discuss the close relationship between the EPR approach and scientific “explanation” in a Usenet post (1997).

4.5 The Concepts of the Approach Are Fundamental

The preceding three subsections respectively suggest that the concepts of entities, properties, variables, and relationships between variables are

- a foundation for several of the fundamental statistical concepts
- central to many of the main statistical methods and
- central to scientific explanation.

This suggests that the concepts of the EPR approach are more fundamental than many of the other concepts that are traditionally discussed in statistics courses.

4.6 The Fundamental Concepts Should Be Taught First

Concepts in a body of knowledge are usually easiest to understand and remember if they are developed in a logical order beginning with the most fundamental. This is especially true if the fundamental concepts are intuitive, as are the concepts of entities, properties, variables, and relationships. Therefore, I recommend that teachers cover the concepts of entities, properties, variables, and relationships *first*, before introducing even the most rudimentary of the traditional statistical topics. By covering the fundamental unifying concepts first, we lay a broad and deep foundation upon which to build the traditional topics.

4.7 The Approach Suggests the Role of Statistics

In subsection 2.4 I recommend that the first goal of the introductory statistics course be to give students a lasting appreciation of the vital role of the field of statistics in empirical research. What *is* this role? I propose the following definition:

The main *role* of the field of statistics in empirical research is to provide efficient general techniques to help researchers study variables and relationships between variables (relationships between properties of

entities), mainly as a means to predicting and controlling the values of variables.

4.8 The Approach Gives Students a Lasting Appreciation of Statistics

Consider:

- the EPR approach is aimed specifically at satisfying a main goal of empirical research—the goal of accurate prediction and control
- most students are directly interested in prediction and control (of variables of interest to them)
- the concept of prediction and control on the basis of relationships between variables is not difficult to understand
- the approach demonstrates that the field of statistics plays a broad role across all empirical research
- the approach is developed in a logical sequence from fundamental concepts (with numerous practical examples).

These points suggest that the EPR approach gives students a lasting appreciation of the field of statistics and its vital role in empirical research.

4.9 The Concepts Are Old But the Approach Is New

The concepts of entities, properties, and relationships are not new. Indeed, researchers and statisticians use these concepts implicitly throughout their thinking and discussion.

Surprisingly, however, the fundamental concepts of entities, properties, variables, and relationships are almost never carefully discussed in a unified approach in introductory statistics courses. I believe that the unfortunate omission of careful unified discussion of these concepts is the main reason why the field of statistics is so widely misunderstood.

(Some leaders in statistical education have already independently adopted an important aspect of the EPR approach in that they emphasize relationships between variables in their introductory courses. For example, using an idea developed by Gudmund Iversen, George Cobb teaches two introductory courses, both of which start with relationships—one devoted to experimental design and applied analysis of variance and the other devoted to applied regression [G. Cobb, personal communication, August 21, 1996]. Similarly, Robin Lock teaches an introductory course devoted to time series analysis—i.e., methods for studying relationships between variables when an important predictor variable is “time” [Cobb 1993, sec. 3.1].)

4.10 Compatibility With Other Approaches

Several helpful new approaches to teaching the introductory statistics course have recently been proposed. As suggested by Moore (1997a), these approaches fall neatly

into two distinct groups: *conceptual* approaches and *pedagogical* approaches.

Under the *conceptual* approaches a particular set of statistical concepts is emphasized and other statistical concepts are de-emphasized. The EPR approach is a conceptual approach. Other conceptual approaches include:

- emphasis on data analysis and de-emphasis of theoretical concepts (especially de-emphasis of probability theory, distribution theory, and the theory of statistical tests) (Cobb 1992; Moore 1992a, 1992b)
- emphasis on exploratory data analysis and de-emphasis of confirmatory data analysis (Tukey 1977; Velleman and Hoaglin 1992)
- emphasis on statistical reasoning and de-emphasis of statistical methods (computations) (Ruberg 1990; Bradstreet 1996)
- emphasis on the Bayesian approach to statistics (usually restricted to more mathematically literate students) (Blackwell 1969; DeGroot 1986; Albert 1996, 1997; Berry 1996, 1997; Berry and Lindgren 1996; Antelman 1997; Moore 1997b, 1997c)
- emphasis on (a) process (which is a particular type of entity) and (b) the minimization of the variation in selected properties of a process, mainly through relationships between variables, generally with less emphasis on formal statistical methods (Snee 1993; Hoerl and Snee 1995; Britz, Emerling, Hare, Hoerl, and Shade 1997)
- emphasis on nonparametric statistical methods (Iman and Conover 1983; Noether 1992)
- emphasis on resampling (Boomsma and Molenaar 1991; Simon and Bruce 1991; Smith and Gelfand 1992; Wonacott 1992; Albert 1993; Simon 1993, 1994; Willemain 1994; Pollack, Fireworker, and Borenstein 1995; Hesterberg 1998)
- emphasis on probability (Falk and Konold 1992).

In contrast to the conceptual approaches, the *pedagogical* approaches to teaching the introductory statistics course emphasize new ways of *teaching* any set of statistical concepts. The pedagogical approaches include:

- more interaction between the teacher and the students instead of straight lectures (Mosteller 1988; Zahn 1994)
- application of the principles of Total Quality Management to the design and management of the course (Hogg and Hogg 1995; Wild 1995)
- use of multimedia, film, or video, to teach the concepts (Moore 1993; Ottaviani 1996; Velleman and Moore 1996; Cobb 1997; Cryer and Cobb 1997; Doane, Mathieson, and Tracy 1997; Newton and Harvill (1997); Macnaughton 1998b; Velleman 1998)
- use of demonstrations, activities, or projects to teach the concepts (Jowett and Davies 1960; Hunter 1977; Hansen 1980; Carlson 1989; Bisgaard 1991; Halvorsen and Moore 1991; Bryce 1992; McKenzie 1992; Roberts 1992; Sylwester and Mee 1992; Zahn 1992; Gunter

1993, 1996; Garfield 1995; Magel 1996; Rossman 1996; Scheaffer, Gnanadesikan, Watkins, and Witmer 1996)

- use of students working in groups instead of working individually (Dietz 1993; Garfield 1995)
- use of cases to teach the concepts (Chatterjee, Handcock, and Simonoff 1995; Czitrom and Spagon 1997; Parr and Smith 1998; Peck, Haugh, and Goodman 1998)
- use of news stories to teach the concepts (Snell and Finn 1992; Snell 1999)
- use of improved methods for assessing students (Eltinge 1992; Cobb 1993; Garfield 1994; Gal and Garfield 1997b; Goldman, McKenzie, and Sevin 1997)
- use of concept maps to teach the concepts (Schau and Mattern 1997a, 1997b)
- use of a computerized library of examples that students can analyze to aid in teaching the concepts (Hand, Daly, Lunn, McConway, and Ostrowski 1994; Velleman, Hutcheson, Meyer, and Walker 1996; Pearl, Notz, and Stasny 1996; OzData 1999; StatLib 1999)
- use of real (or at least realistic) data in examples (as opposed to data that have no obvious practical implications) (Hunter 1977; Cobb 1987; Moore and Roberts 1989; Willett and Singer 1992; Macnaughton 1998a)
- emphasis on improving students' writing ability
- use of computers or calculators to generate data displays, to illustrate and simulate statistical ideas, and to perform statistical analyses.

Many introductory statistics teachers now use some combination of the above conceptual and pedagogical approaches. (The main disagreement among teachers is only about the relative emphasis that each approach deserves.)

(It is possible to classify the use of multimedia, film, video, computers, and calculators as “technological” approaches, rather than as “pedagogical” approaches. However, it seems more reasonable to view technology as a means to better pedagogy rather than as an end in itself.)

There is a simple relationship between the EPR conceptual approach to the introductory statistics course and the other approaches—the EPR approach can be effectively used in conjunction with any (or any group) of them.

Moore (1997a) reviews several of the new approaches to statistical education. Cox (1998) comments on some general aspects of statistical education. Gordon and Gordon (1992) and Hoaglin and Moore (1992) give papers by leading statistics educators about teaching statistics. Hawkins, Jolliffe, and Glickman (1992) give a general discussion of teaching statistical concepts.

4.11 Responses to Criticisms of the EPR Approach

I discuss several insightful criticisms of the EPR approach in a paper (1998a, app. C - G) and in some Usenet posts (1998c).

5. TESTS OF THE EPR APPROACH

Three teachers have tested the entity-property-relationship approach in their statistics courses using a draft textbook (Macnaughton 1986). They commented that

... students found the book enjoyable and easy to understand.

Using a unique approach, Macnaughton has provided a comprehensive first-rate introduction to the material. I would highly recommend the book for use in introductory statistics courses ...

- *Professor Alexander Even, The Ontario Institute for Studies in Education*

... students obtained a good understanding of the basic principles of statistical analysis.

... [the approach] substantially simplifies the material without sacrificing important concepts.

- *Professor John Flowers, School of Physical and Health Education, University of Toronto*

The absence of overt mathematics enables the underlying principles of scientific research ... to be more directly apprehended by persons who have ... weak grounding in mathematics. ...

... Students' comments have been uniformly favorable

... the book is to be commended to the instructor.

- *Professor Donald F. Burrill, The Ontario Institute for Studies in Education*

6. IMPLEMENTING THE EPR APPROACH

Until textbooks based on the EPR approach become generally available, a teacher wishing to use the approach in an introductory course can use the paper for students (1996a) to reinforce class discussion of the five introductory concepts. As noted above, after covering the five concepts, I recommend that the teacher spend the rest of the course covering standard statistical principles and methods in terms of the concepts. Some considerations for implementing the approach follow.

6.1 Motivating Students on the First Day

The first day of class is important because it can set a positive attitude toward the course in students' (and teachers') minds. What should be the very first statistical idea we introduce to students?

I recommend that the first idea be that students are going to learn how to make accurate predictions. We can promise students they will learn how to make accurate predictions using scientific methods that are recognized throughout science and empirical research as being the very best methods available.

For example, we can promise students they will learn how to accurately (but generally not perfectly) predict

- the mark they will get on the final
- their average annual income over the next several years
- their longevity
- whether it will rain tomorrow
- just about anything else of interest (if it can be reliably measured).

(Along with prediction methods, I recommend that the introductory course devote substantial attention to the more complicated methods of exercising accurate control through formal experimentation. However, for simplicity, I recommend that discussion of control and experimentation be omitted at the very beginning—the promise of accurate predictions seems quite enough to engage students. I discuss experimentation further in subsection 6.4.)

If we promise students on the first day of class that they will learn how to make accurate predictions, we arouse their curiosity and we set the stage for development of the five concepts discussed in section 3. The promise also sets the course in a practical direction, which is more likely to impress most students than if we begin with mathematical discussion.

If we promise students on the first day of class that they will learn how to make accurate predictions, we must (for the course to be successful) later deliver on the promise. In particular, the intelligent student will be interested in whether we can demonstrate *practical* methods for making *accurate* predictions. Fortunately, statistical methods (broadly viewed) are universally accepted among knowledgeable empirical researchers as the most practical and most accurate methods available for making predictions (and exercising control).

6.2 What Topics Should Follow the Five Concepts?

Depending on the level of the students, the five concepts I discuss in section 3 can be properly introduced in one to eight class sessions. I recommend in section 3 that after covering the five concepts the teacher spend the rest of the course expanding concept 5 (statistical techniques for studying relationships between variables) by covering various standard topics selected from the field of statistics. (I recommend covering topics that are more frequently used in empirical research first.) The next two subsections discuss two ways of implementing the approach.

6.3 An Easy Implementation

Perhaps the easiest way to implement the EPR approach is to follow discussion of the five concepts with material selected from an already existing introductory statistics course. The teacher can use the five concepts to introduce and unify the material. This enables the teacher to use the EPR approach in conjunction with an already-existing course with only a minimum amount of modification to the course.

6.4 Implementation in Phases

Another more unified way of implementing the EPR approach is to break the course into five phases: an introductory phase, a practical-experience phase, a generalization phase, a specific-methods phase (optional), and a mathematics phase (also optional).

Introductory Phase. In this first phase the teacher introduces the five concepts discussed in section 3 (using numerous practical examples).

Practical-Experience Phase. To solidify students' understanding of the five concepts, the teacher follows the introductory phase with a practical-experience phase in which students obtain hands-on experience with specific statistical methods.

I recommend that the teacher begin the practical-experience phase with discussion of a commonly occurring simple type research project—the observational research project that studies the relationship between two continuous variables. Possibly using the material in the paper for students (1996a) as an introduction, the teacher can discuss how to design an observational research project to study the relationship between two continuous variables, how to use statistical techniques to analyze data from such a research project to determine if a relationship is present between the variables, how to use scatterplots to illustrate such a relationship, and how to use the model equation derived from such a relationship to make predictions. To reinforce the discussion, I recommend that students be given computer assignments to detect and study (practical) relationships between pairs of continuous variables in various sets of data.

If time permits, the bivariate case can be extended to the multiple regression case.

Next, the teacher can discuss “experiments” and the associated statistical methods as a powerful tool for studying *causal* relationships between variables. This topic is important because the study of causal relationships through experiments is the soul of empirical research. Discussion can be in terms of designing simple experiments to study causal relationships, detecting causal relationships in experimental data, illustrating causal relationships, and predicting and controlling on the basis of causal relationships. To reinforce the discussion, I recommend that students be given computer assignments to detect and study simple (practical) causal relationships between variables.

If time permits, the fully randomized one-way case can be extended to the multi-way case, repeated measurements, blocking, analysis of covariance, and so on.

I recommend that the length of the practical-experience phase be adjusted to allow enough time for the teacher to properly cover the material in the next (generalization) phase.

Generalization Phase. After the practical-experience phase, I recommend a generalization phase to cover im-

portant topics that span almost the entire field of statistics. The recommended topics are

1. *Types of Variables.* This topic introduces students to a standard typology for variables, with four mutually exclusive and exhaustive categories: continuous, discrete-ordinal, discrete-nominal, and binary. (This topic is in preparation for the following topic.)
2. *Overview of Statistical Methods.* This topic provides a high-level overview of statistical methods (such as those listed in subsection 4.3 and appendix C), with discussion of the conditions under which each method is applicable. Most of the nineteen methods listed in subsection 4.3 are best summarized in a table, with rows of the table indexing the four possible types of the response variable and with columns indexing the four possible types of the predictor variable(s) in the research. Each cell in the body of the table contains the names of the statistical methods that may be used when the response variable and predictor variables are of the indicated two types. I recommend that students *not* be required to memorize the table, but rather the table be provided as an aid to them if they later need to choose the appropriate method to analyze the results of an empirical research project.
3. *Underlying Assumptions of Statistical Methods.* This topic discusses how every statistical method is based on certain assumptions (which are different for different methods), and why therefore (to avoid later possible embarrassment) a researcher should always verify that the underlying assumptions are adequately satisfied before attempting to draw conclusions from the use of a statistical method.
4. *Design of Empirical Research Projects.* This topic discusses key considerations in the design of empirical research projects, with emphasis on (a) the importance of eliminating possible alternative explanations of the results (Macnaughton 1986), (b) the importance of maximizing the power of the statistical tests (appendix B in the paper for students [1996a]), and (c) the importance of maximizing prediction and control accuracy (appendix E in the paper for students [1996a]).
5. *Statistical Thinking.* This topic gives students practice in understanding, criticizing, and writing reports of empirical research.

In discussing topics 2, 3, and 4 above I recommend that the teacher *not* discuss many details of the topics—the important points are the topics themselves, not their (complicated) details, which can be covered in later courses. A reasonable ending-point for an introductory statistics course is at the end of the generalization phase.

Specific-Methods Phase. For students who are likely to perform empirical or statistical research in their careers, I recommend that the preceding three phases be extended (in courses following the introductory course) with de-

tailed discussion of selected statistical methods, such as those listed in subsection 4.3 and appendix C.

For each method in subsection 4.3, I recommend that the following topics be covered (when applicable):

- *Design*: how to design a research project to study relationships between variables using the method
- *Power*: how to use a statistical package to compute the power of statistical tests for detecting relationships between variables using the method
- *Data Checking*: how to use a statistical package to examine the data from a research project using the method in order to identify and (when appropriate) correct anomalies in the data prior to analysis (using methods for studying univariate and possibly bivariate distributions of the values of variables)
- *Assumptions*: how to analyze the description and results of a research project in order to determine whether the underlying assumptions of the method are sufficiently satisfied to permit drawing conclusions from the use of the method
- *Detection*: how to use a statistical package to analyze the results of a research project in order to detect relationships between variables using the method
- *Illustration*: how to use a statistical package to illustrate relationships between variables using the method
- *Prediction and Control*: how to use a statistical package to analyze the results of a research project in order to derive model equations for relationships between variables using the method, as an aid to prediction or control
- *Reporting*: how to write reports of empirical research projects that use the method.

(A similar set of topics can be used to discuss the statistical methods listed in appendix C.)

Except for statistics or mathematics majors, I recommend that the use of mathematics be avoided in the specific-methods phase. Instead, I recommend that attention be focused on designing research projects and on correctly interpreting the relevant output from a statistical package.

Mathematics Phases. For students pursuing careers in statistics or mathematics, I recommend that discussion in the earlier phases be interwoven and extended with discussion of the underlying mathematics. I illustrate an approach to such discussion in a paper about computing sums of squares in unbalanced analysis of variance (1998d). I recommend that statistics majors be made as aware of the role of statistics in empirical research as students in other disciplines. Otherwise, the tail may become disconnected from the dog.

6.5 Practical Examples

As noted in section 3, and following Hunter (1977, pp. 16-17) and Willett and Singer (1992, p. 91), I recommend that any implementation of the EPR approach discuss each important concept in terms of numerous *practical examples*. Practical examples can appear in lectures,

textbooks, class discussions, exercises, activities, and projects.

I believe that the most important examples in an introductory statistics course are the examples of empirical research projects. For an example of an empirical research project to be judged “practical” I recommend (following Scheaffer 1992, p. 69) that the following question have an affirmative answer:

Does an understanding of the relationship between variables under study in the example have an obvious significant social or commercial benefit? That is, does the example provide some clear basis for action?

Examples of empirical research projects that satisfy this practicality criterion are easy to find in most fields of empirical research. For example, research in medicine to study the relationship between AIDS symptoms and a new treatment for AIDS is clearly very practical according to the criterion. That is, AIDS research (if it finds a new relationship between relevant variables) provides a clear basis for action in the treatment of AIDS. Similarly, research to study relationships between variables that help make computers more efficient or less expensive is also (in a commercial sense) very practical since (if it finds a new relevant relationship) it provides a clear basis for action in manufacturing computers.

Examples that fail to satisfy the practicality criterion seriously detract from the field of statistics because they associate the field with problems that appear to be frivolous (or at best inconsequential). For example, a research project that studies the relationship between people’s forearm lengths and their foot lengths is a “frivolous” research project, since students can see no obvious practical use of knowledge of this relationship.

(Interestingly, if one looks hard enough, there are practical uses of most relationships between variables. For example, the relationship between forearm length and foot length is of interest in orthopedics and physical anthropology. However, a somewhat complicated explanation is needed before students can see the practicality of the relationship in either of those fields. Most students are unimpressed by such complicated and obscure applications.)

If the students in a particular introductory statistics course are all majoring in the same discipline, and if that discipline performs empirical research, we can almost certainly make the greatest impression on these students if we choose significant practical examples of empirical research projects from that discipline.

We can also impress students if we choose practical examples of research projects that use response variables that students themselves are directly interested in predicting and controlling, such as variables reflecting student grades, student health, student happiness, student expenses, and student income.

It is surprising how many examples of empirical research projects in statistics textbooks are not practical. And when one studies such examples and asks “Would an enlightened empirical researcher every actually *do* the research project discussed in the example?” the answer is (for various reasons) often a clear “No”.

(The frequent use of impractical examples by some statistics textbook writers is one reason for insisting that teachers and textbook writers use *real data* in examples. Real data guarantee that at least one empirical researcher has judged the research to be in some sense practical. Unfortunately, insisting on real data does not ensure practicality according to the criterion I recommend above. Furthermore, insisting on real data imposes a significant burden on the teacher, since practical real data that are appropriate for an introductory course are often hard to find. This burden can be avoided by using made-up data as long as the teacher ensures that the examples are practical and realistically described. I further discuss real versus *realistic* data in a paper [1998a, sec. 6.4].)

Unfortunately, not all examples of empirical research projects used in an introductory course can be practical—particularly examples in activities and projects where, for example, it is useful for students to experiment with paper helicopters because real helicopters (or other reasonable entities for study) are too expensive and unwieldy (Rogers 1986, Box 1992, Santy and Einwalter 1997). When an example cannot be practical, I recommend that the teacher carefully show students how the example relates to other examples that *are* practical. This lessens the chance of students thinking that the field of statistics specializes in dealing with frivolous or obscure problems.

I discuss the use of examples further in a paper (1998a, sec. 6).

6.6 Instantiation and Generalization

Once a concept has been introduced and studied through a sufficient number of practical examples, I recommend that the teacher cement the appropriate generalizations about the concept in students’ minds. This helps students to use the concept in new situations. For example, once students understand (through sufficient practical examples) the concept of a relationship between variables, the teacher can make the generalization that most empirical research projects can be viewed as studying relationships between variables.

After stating a generalization, I recommend that the teacher assign exercises in which students identify details of the generalization in specific instances. For example, after discussing the generalization that most empirical research projects can be viewed as studying relationships between variables, I recommend that the teacher assign exercises in which students answer the following questions about various empirical research projects:

1. What is the population of entities under study?

2. What is the size of the sample and how was the sample selected from the population?
3. What is the response variable that was measured in each entity in the sample?
4. What is (are) the predictor variable(s) that was (were) measured in each entity in the sample (or measured in each entity’s environment)?
5. In plain language or in graphical terms (as opposed to mathematical terms), what is the relationship between the response variable and the predictor variable(s) that was sought, discovered, or studied in the research project?

Such exercises are important because once students have interpreted a sufficient number of diverse instances as easily fitting within a generalization, they recognize the unifying power of the generalization. In particular, once students have interpreted a sufficient number of diverse research projects as studies of relationships between variables, they recognize that most empirical research projects can be easily interpreted from this comprehensive simplifying point of view.

6.7 Feedback Systems

How many explanations, examples, exercises, or activities should a teacher provide or assign to ensure that students understand a particular generalization? This depends, of course, on the generalization and on the nature of the students and is often difficult to determine at the front line of teaching—especially if a teacher is using a new approach. To reduce this difficulty, I recommend that teachers use feedback systems to assess whether students understand each main concept and generalization.

Some effective feedback systems for assessing students’ understanding are

- minute papers (which are ungraded and may be anonymous) in which students briefly report their understanding, their questions, and the muddiest point to help the teacher evaluate a lecture or lesson (Mosteller 1988)
- graded quizzes and exercises that test students’ understanding of the ideas (I give examples in the paper for students [1996a])
- two-way discussions between the teacher and students about the ideas.

Gal and Garfield (1997b, pt. 2) give four interesting essays by statistics educators about assessing students’ understanding of statistical ideas.

6.8 The Use of Mathematics in the Introductory Statistics Course

As suggested in subsection 6.4, except in courses aimed at statistics or mathematics majors, I recommend that discussion of the underlying mathematics of statistics (e.g., probability theory, distribution theory, theory of statistical tests) be omitted from the introductory statistics course. This recommendation is motivated by the needs of

the typical user of statistical methods, who is interested in the field of statistics only to the extent that it can help him or her to detect and study relationships between variables (or perform equivalent functions under another name). And like the typical automobile driver who needs transportation, but who cares little about the mechanical details of the engine, the typical user of statistical methods needs help studying relationships between variables, but cares little about the mathematical details of the help. Instead, the user's attention is directed toward the substantive area of empirical research within which he or she is working (e.g., a particular branch of medicine). Thus the less we engage (and confuse) potential users with the complicated mathematical details of statistical methods, and the more we teach them how to properly *use* the methods in empirical research, the greater the chance we have of turning potential users into users.

The paper for students (1996a) illustrates one approach to showing due deference to the underlying mathematics without getting immersed in complicated details.

6.9 Univariate Distributions

Traditionally, the introductory statistics teacher spends a substantial amount of time near the beginning of the course covering the topic of univariate distributions of the values of variables. (The coverage generally includes ways of summarizing and illustrating univariate distributions and may also include the mathematics of univariate distributions.) Since many statistical ideas depend on the concept of a univariate distribution, it is clearly mandatory to cover this topic at some point in students' statistical careers—but where?

Except in courses for statistics majors, I recommend that discussion of univariate distributions be *omitted* at the beginning of the introductory course. This is because students find the topic to be boring and of little obvious use.

(Nor is the topic of univariate distributions *necessary* at the beginning. That is, almost all practical examples of the use of statistics in empirical research do not study univariate distributions but instead study relationships between variables. As suggested by the discussion in section 3, one does not need to understand univariate distributions in order to attain a proper basic understanding of relationships between variables.)

If we omit univariate distributions at the *beginning* of the introductory course, where should we discuss them? For teachers using the phased approach described in subsection 6.4, I recommend discussing univariate distributions as the first method in the Specific Methods phase of the course. Univariate distributions belong first in the Specific Methods phase because the concepts of univariate distributions are helpful in discussing all the other statistical methods.

(The concepts of univariate distributions are especially helpful in understanding statistical power, in checking data for anomalies, in examining data to determine whether the underlying assumptions of a statistical method are satisfied enough to justify the use of the method, and in specifying the estimated accuracy of predictions or control based on a model derived from empirical research.)

I discuss the treatment of univariate distributions further in a paper (1998a, sec. 9.1 and app. G) and in some Usenet posts (1998c).

6.10 Implementation With Computer Support

Appendix D discusses future computer systems for studying relationships between variables. Such systems will make it substantially easier for teachers to convey statistical concepts to students.

7. SUMMARY

Under the entity-property-relationship approach, we present the following five concepts to students at the beginning of the introductory statistics course:

1. entities
2. properties of entities (variables)
3. an important goal of empirical research: to predict and control the values of variables
4. relationships between variables as a key to prediction and control
5. statistical techniques for studying relationships between variables as a means to accurate prediction and control.

To facilitate understanding, the concepts are presented to students in terms of numerous practical examples.

After students have learned the five concepts, standard statistical principles and methods are developed in terms of the concepts, again with emphasis on practical examples.

The EPR approach is broad, and the concepts of the approach are fundamental. The approach gives students a lasting appreciation the vital role of the field of statistics in empirical research.

APPENDIX A: THE PRIORITY OF THE CONCEPT OF 'ENTITY'

We can see further evidence that the concept of 'entity' pervades human thought by considering the role of nouns in human speech. All nouns are simply names for entities (some of which exist in the external world and some of which do not). Since most human sentences contain at least one noun, most human sentences (thoughts) refer directly to entities.

An alternative approach to using the concept of 'entity' is to use an entity-less fog of unattached properties, but this approach seems much less viable, and perhaps

impossible because it is generally necessary to link individual values of various properties together in an analysis. It is the concept of ‘entity’ that does the linking.

Since we cannot easily abandon the concept of ‘entity’, and since several other important concepts can be built atop the concept, a reasonable approach to describing empirical research, science, and statistics is to begin with the concept of ‘entity’ as a primitive. Another approach is to begin with the concept of ‘property’ as a primitive and then to define an entity as being a cohesive cluster of properties. However, the property-first approach has the problem that it presupposes the concept of ‘entity’, because both properties and clusters (i.e., sets) are entities.

APPENDIX B: DO RESEARCH PROJECTS STUDY RELATIONSHIPS?

Some empirical researchers and statisticians are unaware that the concept of a relationship between variables can be used as widely as this paper claims. (Probably more than one-half of all empirical research projects are *not* viewed by their authors as studying relationships between variables.) Therefore, some readers may wonder if it is possible to view most empirical research projects as studying relationships between variables (as a means to prediction and control of the values of variables). These readers are encouraged to seek examples of research projects that apparently cannot be viewed as being studies of relationships and then to reconsider these examples in terms of the five questions in subsection 6.6. One finds that many such examples *can* be viewed as studying one or more relationships between variables, and this fresh point of view substantially increases understanding.

The remainder of this appendix discusses four aspects of empirical research projects that may at first appear *not* to be studying relationships between variables, but which can nevertheless be usefully viewed as studying relationships.

B.1 Physics or Chemistry Experiments

Suppose a physics or chemistry experiment has discovered that if ingredients A and B are mixed together under conditions C, then ingredient D appears. Does this experiment study a relationship between variables? Yes. The predictor variables are “amount of ingredient A”, “amount of ingredient B”, and the variables that reflect the conditions C. The response variable is “amount of ingredient D that is produced”. Viewing the experiment in terms of a relationship between variables has the benefit of reminding us that in addition to being interested in the fact that ingredient D is produced, we also wish to know how much of ingredient D is produced for given values of A, B, and C. This knowledge gives us better prediction or control capability of D.

(The population of entities in the example is all the cases or instances ever in which the ingredients A and B

are brought together under conditions C, and the sample is the set of those instances that occur in the experiment.)

B.2 Parameter Estimation

Some readers may wonder whether a research project that performs “parameter estimation” can be viewed as studying a relationship between variables. One type of parameter-estimation research project is intended to determine (i.e., estimate) the values of one or more “population parameters”, where a population parameter is simply some property of a population, such as the average of the values of some property of the entities in the population. (In more general terms, the procedure is studying the univariate distribution of the parameter.) Any standard procedure for estimating the value of a population parameter can be usefully viewed as the study of a degenerate case of a relationship between variables in which the response variable is present, but *the set of predictor variables is empty*. This point of view is consistent at a deep level with the underlying mathematics. That is, the mathematics of any standard population-parameter-estimation procedure (or any procedure for studying univariate distributions) is (rigorously) the limiting case of the mathematics of some particular standard method of studying relationships between variables, when the number of predictor variables is reduced to zero.

Note that if we adopt the foregoing point of view of parameter estimation, we are naturally led to ask in particular cases *why* the set of predictor variables is empty. Asking that question may prompt us to find some predictor variables. If predictor variables are found that are related to the response variable, knowledge of the relationship will give us a more complete understanding of the response variable than we would obtain by studying that variable in isolation. In particular, the knowledge will give us the ability to make more accurate predictions or the ability to exercise more accurate control of the values of the response variable.

A second type of parameter-estimation research project can be performed to determine (i.e., estimate) the values of one or more parameters of a model (equation) of a relationship between variables. Because model equations are directly related to the study of relationships between variables, this second type of parameter estimation is simply a particular aspect of the study of relationships between variables.

B.3 Interval Estimation

Some readers may wonder whether a research project that performs “interval estimation” can be viewed as studying relationships between variables. Here, given the links just established between parameter estimation and relationships between variables, it is easy to see similar links between interval estimation and relationships between variables. The links occur because the intervals in

interval estimation are simply intervals in which we determine (with a stated level of “confidence”) that the associated parameters probably lie.

B.4 Statistical Tests

Many instances of the commonly used statistical tests can be usefully viewed as tests about relationships between variables. For example, consider a research project that uses a two-sample t -test to determine if there is evidence of a difference between female and male medical patients in their response to a particular medical treatment. The response variable in this research project is the measured value of the “response” to the treatment for each patient who participated in the research project. And the predictor variable is the variable “gender”, which reflects an important property of the patients. (“Gender” is a variable in the sense that for any patient in the research project “gender” has a value, namely, *female* or *male*, and the values *vary* somewhat from patient to patient.) Thus we can view this research project (and the associated t -test) as an attempt to see if there is a relationship in patients between the variables “gender” and “response”. That is, the research project is an attempt to determine whether response “depends” on gender.

Similarly, in a research project in which the entities are cross-classified in two or more different ways, each classification (i.e., each subscript or margin) represents a different predictor variable. These predictor variables may reflect properties of the entities that are under study, or they may reflect properties of the entities’ environment.

Of course, in the case of a particular treatment that is *applied* to the entities, the associated predictor variable reflects a property of the entities *after* they have received the treatment. That is, the predictor variable reflects the amount of the treatment—possibly zero—that an entity received.

The preceding points together with consideration of the standard statistical tests suggest that we can view many instances of the standard tests (e.g., the t -test, many tests of analysis of variance, regression analysis, categorical analysis, time series analysis, etc.) as techniques for detecting evidence of a relationship (in the population of entities under study) between the response variable and one or more of the predictor variables that are studied in a research project.

But even if we agree that we can view many instances of statistical tests as techniques for detecting evidence of relationships between variables, a more fundamental question remains: Is it empirically *useful* to view these statistical tests this way?

To answer that question, note that it is precisely instances of the concept of a relationship between variables in entities that empirical researchers are usually interested in detecting when they use statistical tests in research projects. That is, (using the definition of a relationship

between variables in subsection 7.10 of the paper for students [1996a]) researchers are usually precisely interested in determining whether the expected value in entities of some variable y depends on the values of certain other variables x_1, \dots, x_p . Researchers are interested in making this determination because if they find such a relationship, then we (as society) can confidently use the knowledge of the relationship to make predictions or exercise control.

(Another valid way of viewing some statistical tests is to say that they are techniques for detecting differences between subpopulations. Thus in the t -test example we can say that the test determines whether there is a difference between the subpopulation of females and the subpopulation of males in their response to the treatment. However, saying that certain statistical tests are techniques for detecting differences between subpopulations is not as general as saying that the tests are techniques for detecting relationships between variables. Specifically, all statistical tests that can be viewed as detecting differences between subpopulations can also easily be viewed as detecting relationships between variables by [1] viewing the relevant property that is measured in each entity in the sample as the response variable, and [2] by viewing whatever distinguishes the different subpopulation groups in the research project from each other as constituting the [set of] predictor variable[s]. On the other hand, not all statistical tests that can be viewed as detecting relationships between variables can be viewed as detecting differences between subpopulations. Specifically, in some statistical tests that detect relationships between variables [especially in observational research projects as opposed to experiments, and especially when continuous predictor variables are used], no subpopulation groups exist, and thus the statistical tests cannot be viewed as detecting differences between subpopulations, not even subpopulations of one.)

APPENDIX C: SUPPORT FOR CLAIMS IN SUBSECTION 4.3

In subsection 4.3 I list nineteen statistical methods and then I make two claims about these methods. This appendix provides support for the claims and discusses related matters.

The first claim is that the list of nineteen methods contains almost all of the currently popular statistical methods. This is supported by the fact that if we survey empirical research projects that use statistical methods, we will find that a large majority of research projects use as their main statistical method(s) one or more of the methods in the list.

I call the statistical methods in the list “response-variable” methods. My criterion for calling a method a response-variable method is

A statistical method is a *response-variable* method if it focuses on a single response

variable and zero or more predictor variables.

All the methods in the list can be easily viewed as satisfying this criterion.

(I discuss the study of univariate distributions—i.e., the group of response-variable methods in which a researcher uses a single response variable and *zero* predictor variables—in subsection 6.9 and appendix B.2.)

The criterion above states that response-variable methods use a *single* response variable. Some multivariate methods (e.g., multivariate regression, multivariate analysis of variance) still have easily identifiable response and predictor variables, but they have *multiple* response variables. (These methods are used only rarely in real empirical research.) The set of response variables used in these methods can usually be reasonably viewed as a unity—as a single response variable that is also a *vector*. Thus I view these methods as being included in the response-variable methods. This inclusion is also reasonable from an empirical research point of view in that the response “variable” in an empirical research project ought to represent some unity or property, even if it is a vector consisting of several individual variables. If the response “variable” is just a random conglomeration of properties (of, of course, the same entities), there is no obvious empirical sense in using it as the response “variable” in an analysis.

In considering the nineteen response-variable methods the question arises whether we should speak of relationships “between” the variables or relationships “among” the variables. I recommend using the preposition “between” when referring to relationships studied by the response-variable methods. This is because in any individual use of one of these methods there is only a single response variable (which may on rare occasion be a vector) under study, and the relationship under study is *between* that variable and the predictor variable(s).

Which statistical methods are *not* response-variable methods? Cluster analysis, multidimensional scaling, factor analysis, principal components analysis, and a few other less frequently used methods are not response-variable methods because they do not focus (either explicitly or implicitly) on a specific single response variable. These methods still study variables (properties of entities) and in a loose sense they also study relationships *between* or *among* variables. But since none of these methods focuses on a specific response variable, I view them as a separate group, which I call the “no-response-variable” statistical methods.

The no-response-variable methods are used in total in probably less than one percent of reported empirical research projects that use statistical methods. Thus although the no-response-variable methods are important in a few research projects, I believe they are not important topics for discussion in an introductory statistics course.

On the other hand, the correlation coefficient and the chi-square statistic for association in a contingency table are (implicitly) *included* in the list of nineteen response-variable methods even though when we use these two methods it is not necessary to specify which variable is the response variable and which variable(s) is (are) the predictor variable(s). In the case of these “mathematically symmetrical” statistical procedures it is true that the *mathematics* is oblivious to any distinction between response and predictor variables. However, in the *underlying empirical research* (assuming it is real research) one of the variables is almost always identifiable as the response variable and the other variable(s) is (are) identifiable as predictor variable(s). Thus although it is an interesting side issue that the mathematics of some statistical procedures is symmetrical, the underlying empirical research (assuming it is real or realistic) usually still has response and predictor variables, which it is helpful to identify as such.

In exploratory data analysis one is “looking around” in data, often without a particular response variable in mind. Thus the question arises whether we should view exploratory data analysis as a response-variable method or as a no-response-variable method. One answer is that if an exploratory data analysis is to be put to any practical use, a response variable and zero or more predictor variables will usually be (implicitly or explicitly) determined. Thus exploratory data analysis, when it is put to practical use, can usually be viewed as a response-variable method.

* * *

In subsection 4.3 I list the following four groups of techniques that statistical methods can perform:

1. techniques for *detecting* relationships between variables
2. techniques for *illustrating* relationships between variables
3. techniques for *predicting* and *controlling* the values of variables on the basis of relationships between variables, and
4. miscellaneous techniques for the study of variables and relationships between or among variables.

Then I claim that the *only* techniques that the nineteen response-variable statistical methods can perform are the four groups of techniques. When considering this claim, note that (to emphasize the practical functions of the statistical techniques) the EPR approach views parameter-estimation techniques as a subset of prediction and control techniques. The paper for students gives further discussion of the four groups of techniques (1996a, sec. 8 - 13).

* * *

The four groups of techniques listed above are *broader* than the nineteen response-variable statistical methods I list in subsection 4.3. In particular, we can view the no-response-variable methods as falling in the

category of miscellaneous techniques for the study of variables and relationships between or among variables, since all these methods clearly study variables and relationships *among* variables. This holds true even though the no-response-variable methods do not study relationships between a single response variable and one or more predictor variables in the standard sense.

APPENDIX D: FUTURE SYSTEMS FOR STUDYING RELATIONSHIPS BETWEEN VARIABLES

Nowadays if a researcher wishes to use statistical methods to study relationships between variables, the researcher must be highly skilled in the use of those methods (even though statistical software can do most of the arithmetic). And when unskilled researchers try to use these methods, they often, through misunderstanding, make serious blunders. To prevent these blunders, and to help improve the quality of empirical research, it seems likely that statisticians will develop expert software that will interactively guide an unskilled researcher through the steps of properly designing a research project, performing it, and interpreting the results. I call this type of software “research guidance software”.

To help researchers design a research project, research guidance software will help them select the response and predictor variables and help them determine the important details of the design. An important goal will be to design research projects

- that are unequivocal
- that have minimum cost
- whose statistical tests have maximum power (appendix B in the paper for students [1996a])
- whose subsequent predictions or control will have maximum accuracy (appendix E in the paper for students [1996a]).

To help researchers perform a research project, the software will be capable of controlling the instruments used in the research project and capable of obtaining values of variables directly from the instruments.

To help researchers interpret the results of a research project, the software will automatically guide researchers through the phases of detecting relationships between variables, displaying relationships, and prediction or control, as discussed in the paper for students (1996a, sec. 8 - 13).

Research guidance software will operate in a hypertext computer system, with automatic real-time generation of custom hypertext and graphics that reflect specific details of the research project under study. (The hypertext will be generated from pre-written templates.) To bring the user interface up to the visual resolution of a printed textbook, the software and video system will be capable of simultaneously displaying the equivalent of at least two full pages of highly legible text and graphics from a stan-

dard textbook. In addition, the system will be capable of displaying video segments in which a narrator discusses and illustrates a concept.

(As a prolific writer of notes in the margins of the books I read, I recommend that the software provide a convenient way for students to “write notes in the margin” of its output [including output of the help system]. I also recommend that the software be able to preserve these notes and details of their linkages when it is upgraded.)

Until recently, most researchers and teaching facilities lacked access to computer hardware that could run the type of system described in the preceding paragraphs. Thus the market for research guidance software was too small to justify the development cost of a comprehensive system. However, nowadays the necessary computer hardware is within most researchers’ and many teaching facilities’ budgets.

Furthermore, it is now easier to develop a research guidance software system because such a system can use an existing statistical package as its data analysis engine. (At least one leading statistical package vendor has developed an “output delivery system” that can provide all output from its statistical procedures in formats that are easily used as input by other programs [SAS Institute 1999].) This enables research guidance software developers to concentrate on other important tasks, such as creating the research-project-design and high-level-analysis modules, and writing the many necessary text templates for the system.

The default path through any output from a research guidance software system will be on the highest conceptual level, which will focus the user on the important points. However, customized full information about the underlying details (including tutorials in the associated statistical concepts) will be only a keystroke or two away.

Research guidance software will enable researchers to get much closer to their data because, in interpreting the results of a research project, the software will automatically examine many different graphical views of the data and will present the most interesting views (as defined by research in human perception) to the researcher. (Most researchers would be incapable of generating and examining all the necessary views manually.)

Research guidance software will be designed by teams of statisticians and editors who, in the past, would have developed statistics textbooks. The quality of the writing, graphics, and layout will be equal to that of a superior textbook.

For a research guidance software system to be complete, its developers must codify a complete general research design and data analysis strategy, a difficult but not impossible task. (The codification will involve extensive consultation with experts in research design and data analysis and extensive testing with inexperienced users.) The approach must be general enough and well-enough

presented to enable a well-motivated neophyte to properly design, perform, and analyze efficient simple empirical research projects.

Rudimentary research guidance software systems have begun to appear. Silvers, Herrmann, Godfrey, Roberts, and Cerys (1994) discuss one such system and give references to important earlier systems. Gale and Pregibon (1984) were the pioneers.

The inevitability of useful research guidance software stems from the fact that today's best research project designers and data analysts work by simply proceeding through a complicated (often subconscious, sometimes vague) decision tree. Thus three steps are necessary:

1. decision trees must be elicited from master designers and data analysts
2. decision trees must have any vagueness removed (or perhaps parameterized, with consensually chosen default values for the parameters)
3. software must be developed that will (in the manner of a wise and friendly consultant, and using different levels of detail depending on the user's knowledge) guide a researcher through the decision tree.

Once these steps are completed, the ability of research guidance software to help researchers perform research projects will approach (and likely someday surpass) that of an expert.

Research guidance software amounts to copying both statistical thought and research thought from experts' minds and from textbooks into the computer. Once captured in the computer, the thought can be automatically customized by the computer for the situation at hand, and the researcher or student can actively interact with the thought. Customized interaction guarantees better understanding.

The teaching component will be an important part of any research guidance software system. In the better systems the teaching component will be easy to understand, comprehensive, and compliant with current norms of statistical practice.

Meeker presents a similar vision for statistical technology in subsection 2.2 of an article by Moore, Cobb, Garfield, and Meeker (1995).

REFERENCES

- Albert, J. (1993), "Teaching Bayesian Statistics Using Sampling Methods and Minitab," *The American Statistician*, 47, 182-191.
- (1996), *Bayesian Computation Using Minitab*, Belmont, CA: Duxbury.
- (1997), "Teaching Bayes' Rule: A Data-Oriented Approach" (with discussion), *The American Statistician*, 51, 247-253.
- Antelman, G. (1997), *Elementary Bayesian Statistics*, Cheltenham, England: Edward Elgar.
- Berry, D. A. (1996), *Statistics: A Bayesian Perspective*, Belmont, CA: Duxbury.
- (1997), "Teaching Elementary Bayesian Statistics with Real Applications in Science" (with discussion), *The American Statistician*, 51, 241-246.
- Berry, D. A., and Lindgren, B. W. (1996), *Statistics: Theory and Methods* (2nd ed.), Belmont, CA: Duxbury.
- Biggaard, S. (1991), "Teaching Statistics to Engineers," *The American Statistician*, 45, 274-283.
- Boomsma, A., and Molenaar, I. W. (1991), "Resampling with More Care" (with discussion), *Chance* 4 (4), 25-31.
- Blackwell, D. (1969), *Basic Statistics*, New York: McGraw-Hill.
- Box, G. E. P. (1992), "Teaching Engineers Experimental Design with a Paper Helicopter," *Quality Engineering*, 4 (3), 453-459.
- (1995), "Scientific Statistics - The Way Ahead" (abstract), in *1995 Abstracts: Summaries of Papers Presented at the Joint Statistical Meetings*, Alexandria, VA: American Statistical Association, 38.
- Bradstreet, T. E. (1996), "Teaching Introductory Statistics Courses So That Nonstatisticians Experience Statistical Reasoning," *The American Statistician*, 50, 69-78.
- Britz, G., Emerling, D., Hare, L., Hoerl, R., and Shade, J. (1997), "How to Teach Others to Apply Statistical Thinking," *Quality Progress*, 30 (6) (June), 67-79.
- Bryce, G. R. (1992), "Data Driven Experiences in an Introductory Statistics Course for Engineers Using Student Collected Data," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 155-160.
- Carlson, R. R. (1989), "A Paper Clip Experiment," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 105-106.
- Chatterjee, S., Handcock, M. S., and Simonoff, J. S. (1995), *A Casebook for a First Course in Statistics and Data Analysis*, New York: John Wiley.
- Chromiak, W., Hoefler, J., Rossman, A., and Tesman, B. (1992), "A Multidisciplinary Conversation on the First Course in Statistics," in *Statistics for the Twenty-First Century, MAA Notes No. 26*, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 26-36.
- Cobb, G. W. (1987), "Introductory Textbooks: A Framework for Evaluation," *Journal of the American Statistical Association*, 82, 321-339.
- (1992), "Teaching Statistics," in *Heeding the Call for Change, MAA Notes No. 22*, ed. L. Steen, Washington, DC: Mathematical Association of America, pp. 3-43.
- (1993), "Reconsidering Statistics Education: A National Science Foundation Conference," *Journal of Statistics Education* [online], 1 (1). Available at <http://www.amstat.org/publications/jse/v1n1/cobb.html>

- (1997), *An Electronic Companion to Statistics* (book and CD-ROM), New York: Cogito Learning Media.
- Cox, D. R. (1998), "Statistics for the Millenium: Some Remarks on Statistical Education," *The Statistician*, 47, 211-213.
- Cryer, J. D., and Cobb, G. W. (1997), *An Electronic Companion to Business Statistics* (book and CD-ROM), New York: Cogito Learning Media.
- Czitrom, V., and Spagon, P. D. (eds.) (1997), *Statistical Case Studies for Industrial Process Improvement*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- DeGroot, M. H. (1986), *Probability and Statistics* (2nd ed.), Reading, MA: Addison-Wesley.
- Dietz, E. J. (1993), "A Cooperative Learning Activity on Methods of Selecting a Sample," *The American Statistician*, 47, 104-108.
- Doane, D. P., Mathieson, K. D., and Tracy, R. L. (1997), *Visual Statistics Part I* (book and four computer disks), Boston: McGraw-Hill.
- Eltinge, E. M. (1992), "Diagnostic Testing for Introductory Statistics Courses," in *Statistics for the Twenty-First Century, MAA Notes No. 26*, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 56-65.
- Falk, R., and Konold, C. (1992), "The Psychology of Learning Probability," in *Statistics for the Twenty-First Century, MAA Notes No. 26*, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 151-164.
- Gal, I., and Garfield, J. B. (1997a), "Curricular Goals and Assessment Challenges in Statistics Education," in *The Assessment Challenge in Statistics Education*, eds. I. Gal and J. B. Garfield, Amsterdam: IOS Press, pp. 1-13.
- (eds.) (1997b), *The Assessment Challenge in Statistics Education*, Amsterdam: IOS Press.
- Gale, W. A., and Pregibon, D. (1984), "REX: An Expert System for Regression Analysis," in *Proceedings of COMPSTAT 84*, eds. T. Havranek, Z. Sidak, and M. Novak, Heidelberg: Physica-Verlag, pp. 242-248.
- Garfield, J. (1994), "Beyond Testing and Grading: Using Assessment to Improve Student Learning," *Journal of Statistics Education* [online], 2 (1). Available at <http://www.amstat.org/publications/jse/v2n1/garfield.html>
- (1995), "How Students Learn Statistics," *International Statistical Review*, 63, 25-34.
- Goldman, R. N., McKenzie, J. D., Jr., and Sevin, A. D. (1997), "The BCASA Conference on Assessment in Statistics Courses," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 285-290.
- Gordon, F., and Gordon, S. (eds.) (1992), *Statistics for the Twenty-First Century, MAA Notes No. 26*, Washington, DC: Mathematical Association of America.
- Gunter, B. (1993), "Through a Funnel Slowly with Ball Bearing and Insight to Teach Experimental Design," *The American Statistician*, 47, 265-269.
- (1996), "The MISD/MMSTC Statistical DOE Project." Available at <http://www.macomb.k12.mi.us/math/web1.htm>
- Halvorsen, K. T., and Moore, T. L. (1991), "Motivating, Monitoring, and Evaluating Student Projects," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 20-25.
- Hand, D. J., Daly, F., Lunn, A. D., McConway, K. J., and Ostrowski, E. (eds.) (1994), *A Handbook of Small Data Sets* (with computer disk), London: Chapman and Hall.
- Hansen, J. L. (1980), "Using Physical Demonstrations When Teaching Data Analysis," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 61-62.
- Hawkins, A., Jolliffe, F., and Glickman, L. (1992), *Teaching Statistical Concepts*, London: Longman.
- Hesterberg, T. C. (1998), "Simulation and Bootstrapping for Teaching Statistics," to appear in *American Statistical Association Proceedings of the Section on Statistical Education*.
- Hoaglin, D. C., and Moore, D. S. (eds.) (1992), *Perspectives on Contemporary Statistics, MAA Notes No 21*, Washington, DC: Mathematical Association of America.
- Hoerl, R., and Snee, R. (1995), "Redesigning the Introductory Statistics Course (Report No. 130)," Madison, WI: University of Wisconsin Center for Quality and Productivity Improvement.
- Hogg, R. V. (1990), "Statisticians Gather to Discuss Statistical Education," *Amstat News*, Number 169, November 1990, 19-20.
- (1991), "Statistical Education: Improvements Are Badly Needed," *The American Statistician*, 45, 342-343.
- (1992), "Towards Lean and Lively Courses in Statistics," in *Statistics for the Twenty-First Century, MAA Notes No. 26*, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 3-13.
- Hogg, R. V., and Hogg, M. C. (1995), "Continuous Quality Improvement in Higher Education," *International Statistical Review*, 63, 35-48.
- Hunter, W. G. (1977), "Some Ideas About Teaching Design of Experiments, with 2⁵ Examples of Experiments Conducted by Students," *The American Statistician*, 31, 12-17.
- Iman, R. L. (1994), "The Importance of Undergraduate Statistics," *Amstat News*, no. 215, December 1994, 6.
- Iman, R. L., and Conover, W. J. (1983), *A Modern Approach to Statistics*, New York: John Wiley.

- Iversen, G. R. (1992), "Mathematics and Statistics: An Uneasy Marriage," in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 37-44.
- Jowett, G. H., and Davies, H. M. (1960), "Practical Experimentation as a Teaching Method in Statistics" (with discussion), *Journal of the Royal Statistical Society*, Ser. A, 123, 10-35.
- Macnaughton, D. B. (1986), *How to Predict the Unknown: An Introduction to Scientific Thinking with Statistics*. Textbook, in preparation.
- (1996a), "The Entity-Property-Relationship Approach to Statistics: An Introduction for Students." Available at <http://www.matstat.com/teach/>
- (1996b), "Re: EPR Approach to Intro Stat: Relationships Between Variables (response to comments by Robert Frick)." Published in EdStat-L and sci.stat.edu on July 28, 1996. Available at <http://www.matstat.com/teach/p0015.htm>
- (1997), "EPR Approach and Scientific 'Explanation' (response to comments by Robert Frick)." Published in EdStat-L and sci.stat.edu on July 23, 1997. Available at <http://www.matstat.com/teach/p0026.htm>
- (1998a), "Eight Features of an Ideal Introductory Statistics Course." Available at <http://www.matstat.com/teach/>
- (1998b), "Review of ActivStats 2.0," *The Statistics Teacher Network*, no. 47, Winter (March) 1998, 4-6. Available at <http://www.matstat.com/teach/p0030.htm>
- (1998c), [responses to comments about the paper "Eight Features of an Ideal Introductory Statistics Course"]. Published on various dates in EdStat-L and sci.stat.edu beginning on July 23, 1998. Available at <http://www.matstat.com/teach/>
- (1998d), "Which Sums of Squares Are Best in Unbalanced Analysis of Variance?" Available at <http://www.matstat.com/ss/>
- Magel, R. C. (1996), "Increasing Student Participation in Large Introductory Statistics Classes," *The American Statistician*, 50, 51-56.
- McKenzie, J. D., Jr. (1992), "The Use of Projects in Applied Statistics Courses," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 142-146.
- Moore, D. S. (1992a), "Introduction: What Is Statistics?" in *Perspectives on Contemporary Statistics*, MAA Notes No. 21, eds. D. C. Hoaglin and D. S. Moore, Washington, DC: Mathematical Association of America, pp. 1-17.
- (1992b), "Teaching Statistics as a Respectable Subject," in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 14-25.
- (1993), "The Place of Video in New Styles of Teaching and Learning Statistics," *The American Statistician*, 47, 172-176.
- (1997a), "New Pedagogy and New Content: The Case of Statistics" (with discussion), *International Statistical Review*, 65, 123-165.
- (1997b), "Bayes for Beginners? Some Pedagogical Questions," in *Advances in Statistical Decision Theory and Applications*, eds. S. Panchapakesan and N. Balakrishnan, Boston: Birkhäuser, pp. 3-17.
- (1997c), "Bayes for Beginners? Some Reasons to Hesitate" (with discussion), *The American Statistician*, 51, 254-261.
- Moore, D. S., Cobb, G. W., Garfield, J., and Meeker, W. Q. (1995), "Statistics Education Fin de Siècle," *The American Statistician*, 49, 250-260.
- Moore, T. L., and Roberts, R. A. (1989), "Statistics at Liberal Arts Colleges," *The American Statistician*, 43, 80-85.
- Mosteller, F. (1988), "Broadening the Scope of Statistics and Statistical Education," *The American Statistician*, 42, 93-99.
- Newton, H. J., and Harvill, J. L. (1997), "StatConcepts: A Visual Tour of Statistical Ideas," in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 48-57.
- Noether, G. E. (1992), "An Introductory Statistics Course: The Nonparametric Way," in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 129-138.
- Ottaviani, M. G. (ed.) (1996), *Multimedia and New Educational Environments: Statisticians and Experts Exchange Views*, Rome: Dipartimento di Statistica, Probabilità e Statistiche Applicate, Università di Roma "La Sapienza".
- OzData (1999), "OzData: Australasian Data and Story Library." This archive of data examples is available at <http://www.maths.uq.edu.au/~gks/data/>
- Parr, W. C., and Smith, M. A. (1998), "Developing Case-Based Business Statistics Courses," *The American Statistician*, 52, 330-337.
- Pearl, D. K., Notz, W. I., and Stasny, E. A. (1996), "Finding Examples - The EESSEE Way Out," in *1996 Abstracts: Summaries of Papers Presented at the Joint Statistical Meetings*, Alexandria, VA: American Statistical Association, 168. Further information is available at <http://www.stat.ohio-state.edu/~eeseef/>
- Peck, R., Haugh, L. D., and Goodman, A. (1998), *Statistical Case Studies: A Collaboration Between Academe and Industry*, Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Pollack, S., Fireworker, R., and Borenstein, M. (1995), "Some Resampling Algorithms for the Testing of Hy-

- potheses,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 288-292.
- Roberts, H. V. (1992), “Student-Conducted Projects in Introductory Statistics Courses,” in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 109 - 121.
- Rogers, C. B. (1986), “Design of Experiments for QC Circles,” in *ASQC Quality Congress Transactions—Anaheim*, Milwaukee, WI: American Society for Quality, pp. 354-359.
- Rossman, A. J. (1996), *Workshop Statistics: Discovery with Data*, New York: Springer-Verlag.
- Ruberg, S. J. (1990), “The Statistical Method: The Logic of Science,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 206-211.
- Santy, W., and Einwalter, B. (1997), “Comparison of Classroom Toys for Teaching Experimental Design,” *Quality Engineering*, 10 (1), 75-83.
- SAS Institute Inc. (1999), [SAS Output Delivery System]. For details search the site at <http://www.sas.com/> for the phrase “output delivery system”, including the quotation marks.
- Schau, C., and Mattern, N. (1997a), “Use of Map Techniques in Teaching Applied Statistics Courses,” *The American Statistician*, 51, 171-175.
- (1997b), “Assessing Students’ Connected Understanding of Statistical Relationships,” in *The Assessment Challenge in Statistics Education*, eds. I. Gal and J. B. Garfield, Amsterdam: IOS Press, pp. 91-104.
- Scheaffer, R. L. (1992), “Data, Discernment and Decisions: An Empirical Approach to Introductory Statistics,” in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 69-82.
- Scheaffer, R. L., Gnanadesikan, M., Watkins, A., and Witmer, J. A. (1996), *Activity-Based Statistics: Instructor Resources*, New York: Springer-Verlag.
- Silvers, A., Herrmann, N., Godfrey, K., Roberts, B., and Cerys, D. (1994), “A Prototype Statistical Advisory System for Biomedical Researchers,” *Computational Statistics and Data Analysis*, 18, 341-369.
- Simon, J. L. (1993), *Resampling: The “New Statistics”*. Available at <http://www.resample.com/text.html>
- (1994), “What Some Puzzling Problems Teach About the Theory of Simulation and the Use of Resampling,” *The American Statistician*, 48, 290-293.
- Simon, J. L., and Bruce, P. (1991), “Resampling: A Tool for Everyday Statistical Work,” *Chance*, 4 (1), 22-32.
- Smith, A. F. M., and Gelfand, A. E. (1992), “Bayesian Statistics Without Tears: A Sampling-Resampling Perspective,” *The American Statistician*, 46, 84-88.
- Snee, R. D. (1993), “What’s Missing in Statistical Education?” *The American Statistician*, 47, 149-154.
- Snell, J. L. (1999), “Chance Database.” Available at <http://www.dartmouth.edu/~chance/>
- Snell, J. L., and Finn, J. (1992), ‘A Course Called “Chance”,’ *Chance*, 5 (3-4), 12-16.
- StatLib (1999). This archive of datasets and other statistical information is available at <http://lib.stat.cmu.edu/>
- Sylwester, D. L., and Mee, R. W. (1992), “Student Projects: An Important Element in The Beginning Statistics Course,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 137-141.
- Tukey, J. W. (1977), *Exploratory Data Analysis*, Reading, MA: Addison-Wesley.
- Velleman, P. F. (1998), *ActivStats Version 2.0* (computer software), Reading MA: Addison Wesley Longman.
- Velleman, P. F., and Hoaglin, D. C. (1992), “Data Analysis,” in *Perspectives on Contemporary Statistics*, MAA Notes No. 21, eds. D. C. Hoaglin and D. S. Moore, Washington, DC: Mathematical Association of America, pp. 19-39.
- Velleman, P. F., Hutcheson, M. C., Meyer, M. M., and Walker, J. H. (1996), “DASL, the Data and Story Library: An Online Resource for Teaching Statistics,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 45-49. The library is at <http://lib.stat.cmu.edu/DASL/>
- Velleman, P. F., and Moore, D. S. (1996), “Multimedia for Teaching Statistics: Promises and Pitfalls,” *The American Statistician*, 50, 217-225.
- Watkins, A., Burrill, G., Landwehr, J. M., and Scheaffer, R. L. (1992), “Remedial Statistics?: The Implications for Colleges of the Changing Secondary School Curriculum,” in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 45-55.
- Wild, C. J. (1995), “Continuous Improvement of Teaching: A Case Study in a Large Statistics Course,” *International Statistical Review*, 63, 49-68.
- Willemain, T. R. (1994), “Bootstrap on a Shoestring: Resampling using Spreadsheets,” *The American Statistician*, 48, 40-42.
- Willett, J. B., and Singer, J. D. (1992), “Providing a Statistical ‘Model’: Teaching Applied Statistics using Real-World Data,” in *Statistics for the Twenty-First Century*, MAA Notes No. 26, eds. F. Gordon and S. Gordon, Washington, DC: Mathematical Association of America, pp. 83-98.
- Wonnacott, T. H. (1992), “More Foolproof Teaching Using Resampling,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 273-274.
- Zahn, D. A. (1992), “Student Projects in a Large Lecture Introductory Business Statistics Course,” in *American*

Statistical Association Proceedings of the Section on Statistical Education, pp. 147- 154.

(version of February 2, 1999)

——— (1994), “A Brain-Friendly First Day of Class,” in *American Statistical Association Proceedings of the Section on Statistical Education*, pp. 53-62

CONTENTS

1. INTRODUCTION..... 1

2. COURSE GOALS..... 1

 2.1 The Value of Emphasizing Goals..... 1

 2.2 A Definition of “Empirical Research” 1

 2.3 Topic-Based Goals Have a Serious Drawback 2

 2.4 Recommended Goals 2

3. FIVE CONCEPTS 2

 Concept 1: Entities 2

 Concept 2: Properties of Entities (Variables) 3

 Concept 3: A Goal of Empirical Research: To Predict and Control the Values of
 Variables..... 3

 Concept 4: Relationships Between Variables as a Key to Prediction and Control..... 4

 Concept 5: Statistical Techniques for Studying Relationships Between Variables..... 5

 Summing Up 5

4. EVALUATION OF THE EPR APPROACH 5

 4.1 The Concepts of the Approach Are Easy to Understand..... 5

 4.2 The Approach Provides a Broad and Deep Foundation for Discussion of Statistical
 Concepts 6

 4.3 The Approach Unifies the Main Statistical Methods 6

 4.4 The Approach Links Well With Scientific Explanation 7

 4.5 The Concepts of the Approach Are Fundamental..... 7

 4.6 The Fundamental Concepts Should Be Taught First..... 7

 4.7 The Approach Suggests the Role of Statistics 7

 4.8 The Approach Gives Students a Lasting Appreciation of Statistics 7

 4.9 The Concepts Are Old But the Approach Is New 7

 4.10 Compatibility With Other Approaches 7

 4.11 Responses to Criticisms of the EPR Approach 8

5. TESTS OF THE EPR APPROACH..... 9

6. IMPLEMENTING THE EPR APPROACH..... 9

 6.1 Motivating Students on the First Day..... 9

 6.2 What Topics Should Follow the Five Concepts?..... 9

 6.3 An Easy Implementation 9

 6.4 Implementation in Phases 10

 6.5 Practical Examples 11

 6.6 Instantiation and Generalization 12

 6.7 Feedback Systems..... 12

 6.8 The Use of Mathematics in the Introductory Statistics Course 12

 6.9 Univariate Distributions..... 13

 6.10 Implementation With Computer Support 13

7. SUMMARY 13

APPENDIX A: THE PRIORITY OF THE CONCEPT OF ‘ENTITY’ 13

APPENDIX B: DO RESEARCH PROJECTS STUDY RELATIONSHIPS?..... 14

 B.1 Physics or Chemistry Experiments 14

 B.2 Parameter Estimation 14

 B.3 Interval Estimation..... 14

 B.4 Statistical Tests 15

APPENDIX C: SUPPORT FOR CLAIMS IN SUBSECTION 4.3 15

APPENDIX D: FUTURE SYSTEMS FOR STUDYING RELATIONSHIPS BETWEEN
 VARIABLES 17

REFERENCES 18