

Teaching Decision Theory in Applied Statistics Courses

by:

Dr. Robert F. Bordley;

GM University; MC 482-D20-B24; Renaissance Center; P.O. Box 100

Detroit, Michigan 48265-1000

robert.bordley@GM.COM

Teaching Decision Theory In Applied Statistics Courses

Abstract

There's been much concern about making the engineering statistics curriculum more relevant to the needs of industry. One proposed solution is to include decision risk analysis in the curriculum. But the current coverage of decision risk analysis in statistics problems is either nonexistent or very introductory. In part, this reflects the fact that decision risk analysis, as currently taught, relies on the complex notion of a utility function.

Recent research in decision theory suggests a way of comprehensively and rigorously discussing decision theory without using utility functions. In this new approach, the decision risk analysis course focuses on making decisions so as to maximize the probability of meeting a target. This allows decision theory to be integrated with reliability theory. As we argue, this course will be just as easy to teach as the conventional introductory treatment of decision theory while being much more comprehensive. In addition, integrating decision theory with reliability theory facilitates its incorporation in curricula that currently exclude decision theory.

Key Words : Decision Theory, Reliability Theory, Introductory Statistics, Goals, Utility Functions

1. Why Decision Theory is Needed

(1.1) Concerns about Statistical Education

Many authors have emphasized the need for improvements in statistical education. As Bisgaard(1991) writes

"There is an implied criticism of the way the statistics profession has approached what has become known as engineering statistics. And that criticism not only concerns teaching style, but also content and organization. I dare say that in the past, we have failed miserably."

Smith(1998) noted that this concern also applies to other courses in the statistics curriculum:

"A radical reform of introductory statistics classes has been advocated by many, often motivated by observations similar to Hogg(1991):`students frequently view statistics as the worst course taken in college.'"

One of the most basic concerns focusses on relevance. As Romero, Ferrer, Capilla, Zunica, Balasch, Serra and Alcover(1995) note:

"The most common criticism of the teaching of statistics in the United States is that it is too academic in focus, excessively theoretically and divorced from the real problems that can appear in the industry and business world."

Barabba(1991) emphasized the importance of addressing these concerns in his presidential address to the American Statistical Association:

"If we are to ensure continued societal support, we as a profession of statisticians must make ourselves more clearly understood and valued by those who use our work. Without this understanding and appreciation, statistics will not be fully used. If we continue on our current course, statistically based information will, at best, be used only to confirm earlier predispositions, and at worst, be done away with in times of scarce resources."

Given such concerns, Snee(1993) concludes:

"Articles in recent issues of the American Statistician highlight the growing feeling that statistical education is in serious trouble and that changes must be made...changing the content of statistical education is needed to help students create value for statistical thinking."

(1.2) Statistical Decision Theory as a Partial Solution

But how should the content of statistical education be changed? Hoerl, Hooper, Jacobs and Lucas(1993) advocate greater focus on managerial and strategic decisionmaking

"Numerous authors have discussed the lack of recognition, respect and effective use of statisticians, particularly in industry. Many have focused improvement efforts at engineering schools, business schools, better selling of statistics or better public relations by the ASA. There appears to be a general consensus that statisticians have typically accepted too narrow a role, being focused more on the statistical methods being employed than the problem addressed. For example, Snee noted the three distinct levels at which a total quality philosophy applies: strategic, managerial and operational. Statisticians, he claimed, tend to ignore the strategic and managerial levels, and concentrate only on the operational, that is, on statistical tools. Very few statistics courses or textbooks teach synthesis of statistical methods into a strategy to solve a class of problems."

In addition, as Hahn and Hoerl(1988) emphasize,

"Today's statistician must strive to get involved in the strategically most significant projects."

But what specifically needs to be done to focus the curriculum on 'a synthesis of statistical methods into a strategy to solve a class of problems'? Barabba(1991), based on decades of practical experience at Kodak, the Census Bureau, General Motors and private consulting, suggests a possible answer:

"Now let me offer our profession an opportunity with potentially greater rewards. The opportunity is to help decision-makers use information to move beyond the realm of 'just seeing the facts that fit' to a higher standard of improved decision quality that leads to wiser decisions. Let me offer a protocol which allows providers of information to meaningfully contribute to quality decisions: a framework that we use at General Motors to enhance the decision process, decision and risk analysis. Through the process, the different functions, perceptions and skills of information users and providers are acknowledged and structured to improve the decision process. The problem is identified, alternatives are analyzed and the best alternative is selected, all with concurrence of the decision-maker and the information providers."

Work by Sharpe & Keelin(1998), Kusnic & Owen(1992), Howard(1988) and Bordley(1998) provide examples of successful applications of decision and risk analysis to high-level corporate problems. As Howard(1988) noted,

" Today it is difficult to find a major corporation that has not employed decision analysis in some form"

Hence incorporating decision and risk analysis into the curriculum might address some of the problems noted in statistical education.

(1.3) But Statistical Decision Theory is not Widely Taught

Barabba recommended decision risk analysis to the statistical community ten years ago. But despite this recommendation and the formation of the American Statistical Association's section on risk analysis, an extensive review of introductory statistics textbooks shows that decision risk analysis is still not widely taught in statistics. One exception is the textbook by Levine, Berenson and Stephan(1998). It's also much more commonly taught in operations research textbooks (Chelst 1998; Anderson et al 1994; Gould, Eppen and Schmidt 1991; Hillier and Lieberman 1995; Ragdale 1995; Winston 1994) . But as the next section will show, even these textbooks provide fairly limited coverage of decision theory.

The purpose of this paper is to help expand the teaching of decision theory in statistics. The next section will examine why decision theory isn't being more widely taught. The third section will propose a solution, based on some recent technical breakthroughs, which integrate decision theory with reliability theory. This section illustrates how this new approach to teaching decision theory might be taught. The fourth section shows how this new approach might be integrated into a statistics course for engineers. The last section discusses the author's experience in applying this approach.

2. Diagnosing the Problem

(2.1) How Decision Theory is Currently Taught

Most textbooks specializing in decision theory generally contain four modules:

1. **PROBLEM-STRUCTURING:** They list the relevant decision alternatives, the uncertain 'states of nature' relevant to the decision problem and the payoff for the different outcomes associated with each decision alternative:
2. **PROBABILISTIC ANALYSIS:** They introduce the criterion of maximizing expected monetary value. After illustrating its use (possibly with a decision tree), they then discuss calculations of the value of perfect and imperfect information using Bayes' Rule.
3. **UTILITY ANALYSIS:** They spend some time showing that expected monetary value ignores individual attitudes toward risk and hence is an unreasonable rule for high-stakes decisions. They then introduce the notion of utility as that probability u making an individual indifferent between receiving that consequence and receiving some best possible consequence with probability u and some worst possible consequence with probability $(1-u)$. They use the utility function to convert the payoff matrix into a matrix of utilities. They then modify the criterion of maximizing expected monetary value by replacing monetary values with utility functions.
4. **MULTIATTRIBUTE UTILITY ANALYSIS:** They also discuss multiattribute utility theory as an extension of utility theory to handle multi-criteria decision making

Introductory textbooks in operations research provide only a very high-level discussion of multiattribute utility theory. Of those statistical textbooks discussing decision theory, most focus on only the first two modules and generally ignore serious treatments of utility.

As Chelst(1998) emphasized, ignoring utility theory (which allows the student to handle risk-preferences) and multiattribute utility theory deprives the student of some of the techniques needed for many successful applications of decision analysis.

(2.2) Why is the Coverage of Decision Theory So Limited?

There are, of course, good reasons why decision isn't theory being taught more extensively in statistics courses. The teacher's time is limited and there are many other statistical topics that also need to be taught. Thus Hogg(1985)'s description of an idealized and ambitious statistics course for engineers includes:

1. Data Collection
2. Basic Probability
3. Basic Statistical Inference
4. Regression Modeling
5. Design of Experiments
6. Reliability

All of these topics are clearly essential. But Hogg's list never even mentions decision theory. There simply isn't room for another topic in his curriculum.

So how can decision theory be incorporated into an already packed curriculum?

First note that the conventional treatment of decision theory begins with a discussion of expected value maximization and then introduces utility as an enhancement of expected value. Note also that the concept of utility---as defined by decision theory---is unique to decision theory and does not appear in any other part of Hogg's curriculum. This suggests that the notion of a utility function is esoteric and hence is harder to teach than other aspects of decision theory.

(2.3) A Possible Solution to these Problems

But recent technical results (Castagnoli and LiCalzi 1996; Bordley and LiCalzi 2000) establish that the traditional decision theory criterion of maximizing expected utility is equivalent to maximizing the probability of exceeding a goal in the presence of background uncertainty. These results indicate that decision theory can be taught, in all of its generality, without using the concept of utility. As the next section will show, this provides a way of teaching decision theory as an extension of reliability theory. Since reliability theory is already part of Hogg's curriculum, teaching decision theory as an extension of reliability theory provides a very simple way of introducing it into an applied statistics course for engineers.

3. Teaching Decision Theory Without Utility Theory

Suppose that the instructor has already familiarized students with basic probability theory and random variables. The remainder of this section now presents an example (in larger font) of a possible lecture using the target-based approach to decision theory:

(3.1) Conceptual Discussion of Individual Goal-Seeking

Suppose your long-run goal is having wealth with a present value of t , ten years from now. Your wealth ten years from now is determined by the sum of the present value of your earnings this year (which we might call x) plus the present value of your earnings in all subsequent years (which we call y). You will have achieved your objective if $x+y$ exceeds t dollars. Let $T=t-y$ define the amount of wealth that must arise from your earnings this year in order for you to reach your goal.

Hence your objective, this year, is to make those decisions whose potential short-term consequences, x , maximize the probability of exceeding T this year. Since you're uncertain about your future earnings in subsequent years, y is uncertain and should be represented as the random variable Y . Thus $T=t-Y$ is also a random variable. Hence getting consequence x gives you a probability $\Pr(x+Y>t)=\Pr(x>T)$ of meeting your goal. If you're offered a gamble, X , giving you a payoff of x this year with probability $\Pr(X=x)$ for various values of x , then the probability of meeting your goal by taking the gamble is $\Pr(X>T)$.

(3.2) Application to a Problem

Consider the following example:

Suppose your only concern is having a wealth of \$1,000,000 by the time you retire (which is ten years from now.) Suppose you also have mineral rights on a piece of land that you believe may have oil underground. You currently only have \$200,000 and it costs \$200,000 to drill for oil. If you drill, there is a 5% chance that you will strike oil. If you strike oil, you

can sell the property for \$5,000,000 this year.

To analyze this problem, suppose first that our wealth ten years from now was exactly equal to whatever we earned from the oil property

1. What's the probability that you'll be able to meet your goals if you don't drill? In this case, it's zero. There's no chance of having a million dollars ten years from now if we only have \$200,000 now.
2. What's the probability that you'll be able to meet your goal if you do drill? In this case, there's a 5% chance of striking oil if we drill. If we do strike oil, we'll have five million dollars which guarantees that we'll achieve our goal. Thus the probability of meeting our target if we drill is five percent.
3. If you're only interested in maximizing the probability of meeting your goal, should you drill? In this case, the answer is clearly yes.

In analyzing this problem, we assumed, quite unrealistically, that our wealth ten years from now was solely equal to the wealth, x , we earned from the oil property.

But there are a lot of other wealth-creating events, which might happen between now, and ten years from now. We might win a million dollars from the lottery five years from now. We might have drastic medical expenses that reduce our wealth by half a million dollars. Thus our wealth ten years from now is really the sum of x and some random adjustments in that current wealth. If we define T to be the difference between our goal and these random adjustments, then maximizing the probability of achieving our goal is the same as maximizing the probability that our near-term payoff exceeds T .

Because of our uncertainty about what the future may bring, we are uncertain about the value of T . To reflect this uncertainty, we need to specify the probability distribution of T . This probability distribution needs to be specified based on our limited knowledge of the possible wealth-creating events that might happen. We now consider some possible specifications of the distribution of T . We will show that our solution varies depending upon how we specify the distribution of T . We will then discuss which specification of the distribution of T is best.

Possibly the simplest assumption is to treat T as uniformly distributed. As a result, the probability of achieving the goal given x is $(x/\$5,000,000)$ with $0 < x < \$5,000,000$.

Given this assumption,

- The probability of achieving the goal given we don't drill is $\$200,000/(\$5,000,000)$ or 4%.
- The probability of achieving the goal if we drill and strike oil is one. The probability of achieving the goal if we drill and don't strike oil is zero. Hence the overall probability of achieving the goal if we drill is 5%.

Thus drilling maximizes the probability of achieving the goal.

The uniform distribution presumes that the required amount of wealth you need to earn from the oil well is uniformly distributed between zero and five million. An alternative, more optimistic, assumption is to presume that the amount of wealth we need to earn is exponential, i.e., the probability of achieving the goal given x is exponential and is given by $(1-\exp(-x/R))$ for some constant R

This formula indicates that the amount we need to reach our goal is more likely to be between \$0 and \$2,500,000 than it is to be between \$2,500,000 and \$5,000,000. Hence it suggests that we may reach our goal even if we don't do that well on the oil venture.

To use this formula, we need to assess R . To assess R , suppose there's a 50% chance that you'll meet your goal if your payoff this year is \$2,500,000. Then we can calculate R to equal \$3,600,000. Note that R is the expected value of T , i.e., it is the expected amount required in the first year to ensure that the individual meets the long-run goal.

Given this exponential assumption,

- The probability of achieving the goal given we don't drill is $(1-\exp(-200/3600))=(1-\exp(-0.056))$ which equals 5.4%.
- The probability of achieving the goal if we drill and strike oil is $(1-\exp(-5000/144))=1$. The probability of achieving the goal if we drill and don't strike oil is $(1-\exp(-0/144))=0$. Hence the overall probability of achieving the goal if we drill is 5%.

Thus not drilling maximizes the probability of achieving the goal.

(3.3) Differences in Individual Risk Preferences

Our examples with the uniform and exponential assumptions show how our decision about whether or not to drill depends upon our assumptions about T . This makes it important to think about our long-term goal and about how future events may adjust whatever we earn from the current decision at hand.

In our example, the random variable Y (and thus T) reflects random changes in our wealth that might arise over future years. If Y represents the result of many small changes in wealth, then we might want to assume that Y (or the logarithm of Y) is normally distributed. This is what is generally done in modeling price change on the stockmarket. Specifying this distribution requires a specification of the average percentage rate of increase in our wealth and the variance in this rate of increase.

If we specify a large variance, then we're saying that we know very little about future changes in our wealth. When the variance is very large, our distribution over T will be approximately uniform and we'll tend to take calculated risks. This is referred to as being 'risk-neutral.' In our example, risk-neutrality led us to invest in the oil well.

When we specify a small variance, we're saying that we're very sure about what we need to

earn this year. If 'playing it safe' will give us the return we need this year, we'll play it safe (which is what we did with the exponential distribution). On the other hand, if 'playing it safe' won't give us the return we need, we'll take extreme risks. Hence we're very sensitive to risk when the variance is small.

Thus different assumptions about T imply that individuals have different attitudes toward risk. Conversely suppose that we know an individual's attitude toward risk but don't know whether the individual is explicitly trying to maximize the probability of meeting some long-term goal. Then an important technical result establishes that a 'rational' individual always acts like an individual maximizing the probability of exceeding some uncertain threshold, T . This individual's risk attitude is represented by specifying some distribution over T . (This result also assumes that there is some random variable Z such that $(X-Z)$ and $(T-Z)$ are independent. In other words, if there is correlation between X and T , that correlation can be attributed to a single common uncertainty, Z , which can be subtracted off from X and T .)

(3.4) Transition to a Discussion of Reliability Theory

In the oil drilling example, the individual was only concerned with monetary payoffs. But in many actual problems, there are a wide variety of criteria that are important. To approach such problems, we need to draw upon techniques from reliability theory.

At this point, the instructor needs to spend time teaching fault tree analysis and reliability theory. Once the students are reasonably familiar with these techniques, the instructor can then return to the decision problem and start section (3.5).

(3.5) Transition from Reliability Theory to Multiattribute Decision Theory

We now apply these concepts from reliability theory to our decision theory problem.

Suppose the individual has a list of goals that are important in a decision problem. For each goal, the individual needs to ask 'Why is that goal important?' in order to identify that higher-level goal which this goal supports. Continue asking 'Why is that goal important?' until you identify some ultimate goal. Once this goal is identified, begin building a fault tree as follows:

1. Specify achievement of the ultimate goal as the top-event in a fault tree
2. Identify those first-level goals which, if achieved, would directly lead to the achievement of this ultimate goal
3. For each of these first-level goals, identify second-level goals which, if achieved, would directly lead to the achievement of that first-level goal.
4. Continue, using standard fault-tree approaches, to connect lower-level goals to higher-level goals. Stop when the lowest level goals are fairly concrete.
5. For each of your possible consequences x , assess the probability of reaching each of the lowest level goals.
6. Infer the resulting overall probability for meeting the ultimate goal for each of your possible consequences, x .

If the consequences of your decisions are uncertain, then choose that decision which has the highest probability of meeting your long-run goal.

4. A Proposed Curriculum

The previous section showed how the target-based approach allowed us to cover decision theory fairly quickly. This section discusses an even more important benefit of integrating decision theory in the curriculum. The success of decision risk analysis in industry is based on its integration into a structured process for decision making called the decision dialogue process. This process is often advertised in consultant brochures as 'a practical team-based approach to making tough decisions (SDG 1997).'

It involves having the key executives form a 'decision review board' which, because of executive time commitments, meets infrequently. Simultaneously a core team, composed of delegates representing the executives, is formed to work more continuously on the problem. Once formed, the core team:

1. First defines the problem and specifies the problem criteria and then reports back to the decision review board for their comment and approval.
2. Generates possible solutions and meets with the review board for their comment and approval. The core team then develops a model to analyze each alternative. They use the analysis to identify a small number of critical key factors and report the results back to the review board. This model is used to identify needs for further information.
3. Develops a detailed model based on these key factors. The core team then meets with the review board to develop a consensus recommendation and to plan implementation.

Now suppose that Hogg's idealized curriculum and the dialogue decision process are integrated by organizing the curriculum around the following three modules:

Module 1: Defining decision criteria.

In this module, the instructor can focus on specifying the decision criteria in a practical problem. In order to understand the decision criteria, the instructor teaches the material described in our previous section. This will also involve teaching probability theory and reliability theory.

Module 2: Defining the critical factors impacting those criteria

This step involves filtering the myriad variables in the decision problem to identify those few that really matter. Montgomery(1991), Coleman & Montgomery(1993), Box, Hunter & Hunter(1978), Hahn(1977,1984), Natrella(1979), and Lorenzen, Andersen & Andersen(1993) have developed a process for implementing this step based on the design of experiments. Their process presumes knowledge of data collection, statistical analysis and the design of experiments. Hence the module corresponding to this step should teach data collection, statistical analysis and the design of experiments.

Module 3: Analyzing possible solutions

This last step focuses on carefully modeling the problem using only the critical factors. Because a variety of key factors are deliberately omitted, developing this high-level model may involve

constructing a regression model with the omitted variables being implicitly subsumed in the error term. Hence the instructor needs to teach regression analysis during this module.

Since this three-step process is focused around making a decision, the instructor might also wish to interject guidelines for effective statistical consulting throughout the course. Ideally the instructor would motivate the course by providing students with a single case study, which is then solved in the process of teaching the three modules.

Thus General Motors --- at which the author works--- has an internal course focused on a recently transferred GM employee who needs to decide how to sell his house. This problem naturally involves multiple criteria, which can be structured using reliability theory. A variety of factors impact these criteria and design of experiments is required in order to identify the truly critical factors. Finally a high-level model is constructed based only on the most critical factors.

Note that the three modules proposed in this course---a probabilistic module, a statistical module and a modeling module---involve teaching all six elements in Hogg's idealized curriculum. Hence the integration of decision theory and reliability theory allows the instructor to teach all of Hogg's curriculum in a course focused around a practical decision process widely used in many large corporations.

5. Actual Experiences

My experience with teaching this proposed new approach is based on:

1. Regular semester (advanced undergraduate) engineering statistics courses at Oakland University
2. Intensive one-week engineering statistics courses at Austrian Technical University. As Birch (1995) noted, the teaching of such courses often involves a diverse audience with one common goal, "to achieve their educational objectives as quickly as possible and to begin using their new skills immediately." Hence it's a different audience than is typically encountered in regular semester classes.
3. In-house continuing education classes to GM employees and executives. This audience is very pragmatic and has little tolerance for concepts that are not obviously applicable.

Teaching these three kinds of classes has given me significant experience on the kinds of demands associated with different teaching formats. It has also allowed me to experiment with different ways of teaching decision theory.

(5.1) Experiences Teaching Decision Theory with Utility Functions

When I initially began teaching, I employed the standard approach to decision theory. As the introduction noted, this approach focuses first on teaching expected value, then on pointing out the limitations of expected value, and then on teaching expected utility. In several of my experiences teaching this traditional approach, a few of my students openly wondered why I was wasting their time teaching expected value if it was a flawed approach.

I would then go on to teach expected utility. But the amount of class time remaining to teach expected

utility was relatively short. While the students did enjoy my discussions of multiattribute utility, there wasn't enough time in this introductory course to make them fully comfortable with the concept. It seemed clear that few, if any of my students, would actually solve any of their own problems by assessing utility functions.

The fact that I was teaching material which the students would almost certainly never apply was obviously of considerable concern. As Romero, Ferrer, Capilla, Zunica, Balasch, Serra and Alcover (1995) wrote:

"What students really learn has little to do with what they do by heart or in a random way or on an examination after days devoted to intensively studying the subject. What they actually learn is what they are able to apply in their jobs ten years later."

From that perspective, then, my students weren't learning much. Of course, one might argue that students would be able to apply these concepts if they also took a follow-on advanced class in utility theory. But like most universities, my college had little interest in holding such a follow-on class.

(5.2) Experiences Teaching Decision Theory using Both Approaches

As a result, I began tentatively teaching decision theory using first the target-based approach and then the standard utility-based approach.

To do this, I would teach decision trees and then focus on the problem of an individual considering several job offers who was concerned with money, social life, and sports activities. I would ask the students to suppose that the individual would only be satisfied if they were satisfied in each of these three areas of life. I would then have them work through the probability of the individual being satisfied with the various job prospects available. The students had little problem understanding this formulation and easily worked through the case study with minimal involvement from me. Indeed, I had several lunchtime experiences helping students solve decision analysis problems from their workplace by making notes on napkins!

Five weeks later, I then had the students go through a related problem using the multiattribute utility approach. In this case, I spend one day teaching multiattribute utility theory and the assessment techniques and the second day applying it to the case study. In this second approach to multiattribute utility theory, I spent considerably more time leading the class through the case. It was clear that the students could not work through the case study without considerable involvement from me. None of the students ever discussed applying this formulation to problems in their workplace.

The fact that my students could apply the target-based approach with minimal involvement from me was especially appealing and indicated that they would, in fact, leave the class with a technique which they could use.

(5.3) Experiences Teaching Decision Theory without Utility Functions

This emphasis on leaving students with techniques they can immediately use was even more important in my internal classes at General Motors. There were many GM project teams who wished to apply decision risk analysis. These teams were often assigned decision analysis facilitators and were also asked to take a short course in decision risk analysis prior to actually using decision risk analysis to make the decision. (My colleagues and I referred to this as 'Just-In-Time Education'.) We initially tried teaching such utility-based concepts as 'willingness to pay' which focuses on an individual's

hypothetical willingness to pay for improvements along various attributes. Thus in one application, we had individuals assess their willingness to pay for a one percent improvement in fuel economy, their willingness to pay for a one percent improvement in quality, etc. We did, in fact, complete many projects using these kinds of notions.

But we continued to encounter resistance from individuals accustomed to setting targets and then focusing efforts to meet those targets. Thus in GM's Powertrain Planning groups, individuals preferred to use future envisioning exercises to set long-run corporate fuel economy targets. They would then evaluate various decisions in light of their contribution toward meeting those targets. This target-based approach to decision-making was quite entrenched in the organization.

We did initially spend considerable time trying to convince people to abandon target-based approaches for utility-based approaches. But since there are no degrees offered in internal GM university classes, the professor's only control over the students is their belief that what is being taught will help them with their work. As a result, we eventually stopped trying to replace the target-oriented approach and focused more on making that target-oriented way of thinking more valid.

(5.4) My Interpretation of these Experiences

Were my experiences unique? A review of the literature on organizational decision making shows that this kind of target-based approach to decision making is common to most bureaucratic organizations. Indeed Herb Simon earned the Nobel Prize, in part, for his theory of bounded rationality which emphasized that individual decision making was based on looking for alternatives that would meet prespecified targets. As Simon(1978) noted,

"research in information processing psychology provides conclusive evidence that the decision-making process in problem situations conforms closely to the models of bounded rationality."

The prevalence of target-seeking behavior in organizations reinforces my belief that the target-oriented approach to teaching decision risk analysis is much more natural for those who teach applied statistics to engineers.

Conclusions

This paper presented an alternate way to teach decision theory. This approach has several potential advantages:

1. Formally teaching the concept of utility requires a discussion of reference lotteries and the principle of maximizing expected utility. In the proposed approach, we eliminate this entire discussion and replace it with the idea of maximizing the probability of achieving a goal. This notion of maximizing the probability of achieving a goal is familiar to many statistics students.
2. Multiattribute utility theory is currently taught via tools and techniques specific to decision theory. In the proposed approach, it is taught as an application of reliability theory, which most students must learn anyway.

This approach to decision theory draws heavily upon existing elements in the statistical curriculum with minimal introduction of non-statistical concepts. Hence it may be easier to

incorporate into a statistics class.

But there is one further advantage of this approach:

3. It allows the instructor to organize all of the elements in Hogg's idealized curriculum as part of a single coherent process for making decisions. This can enable the instructor to give a unifying theme to the survey course in applied statistics.

By doing so, it also shows, consistent with Barabba's vision, how decision theory can be blended with other statistical tools in a way that's extremely helpful to high-level decision making.

REFERENCE LIST

1. Anderson, D.; Sweeney, D. and T. Williams (1994). *An Introduction to Management Science*. West Publishing, St. Paul, Minnesota.
2. Barabba, V. (1991). "Through a Glass Less Darkly." *Journal of the American Statistical Association*. 3,86,413, pg.1-8.
3. Bisgard, S. (1991). "Teaching Statistics to Engineers." *The American Statistician*. 45,4.
4. Birch, J. (1995). "Ten Suggestions for Effectively Teaching Short Courses to Heterogeneous Groups." *American Statistician*, 49,2,190-195.
5. Bordley, R. and M. LiCalzi (2000). "Decision Analysis Using Targets instead of Utility Functions." *Decisions in Economics and Finance*. 23,53-74.
6. Bordley, R. (1998). "R&D Project Generation versus R&D Project Selection." *IEEE Transactions in Engineering Mgt*. November.
7. Box, E., W. Hunter & J.S. Hunter (1978). *Statistics for Experimenters*. John Wiley & Sons. New York, 1978.
8. Castagnoli, E. and M. Li Calzi (1996). "Expected Utility without Utility." *Theory and Decision*. 41,281-301
9. Chelst, Kenneth (1998). "Can't See the Forest Because of the Decision Trees: A Critique of Decision Analysis in Survey Texts." *Interfaces*. 28:2,80-98.
10. Coleman, D. and D. Montgomery (1993). "A Systematic Approach to Planning for a Designed Industrial Experiment." *Technometrics*. 35,1,1-12.
11. Gould, F., Eppen, G. and C. Schmidt (1991) *Introductory Management Science*. Prentice Hall, Englewood Cliffs, New Jersey.
12. Hahn, G. (1977) "Some Things Engineers Should Know About Experimental Design." *Journal of Quality Technology*. 9, 13-20.
13. Hahn, G. (1984). "Experimental Design in a Complex World." *Technometrics*. 26.19-31.
14. Hahn, G. and R. Hoerl (1998). "Key Challenges for Statisticians in Business and Industry." *Technometrics*. 40,4, 195-200.
15. Hoerl, R., J. Hooper, P. Jacobs and J. Lucas (1993). "Skills for Statisticians to survive and prosper in the emerging quality environment." *American Statistician*. 47,4, 280-291.
16. Hillier, F. and G. Lieberman (1991). *Introduction to Operations Research*. McGraw-Hill, New York.
17. Hogg, R. (1985). "Statistical Education for Engineers: An Initial Task Force Report." *The American Statistician*, 39, 168-175.
18. Howard, R. (1988) "Decision Analysis: Practice & Promise." *Management Science*. 34,5,679-695.
19. Keeney, R and H. Raiffa. (1976). *Decision Analysis with Multiple Objectives* (New York: Wiley.)
20. Kusnic, M & D. Owen. (1992). "The Unifying Vision Process: Value Beyond Traditional Analyses"

- in Multiple Decision Maker Environments." *Interfaces*.22,6.
21. Levine, D., M. Berenson and D. Stephan.(1998). *Statistics for Managers*. Prentice Hall, New Jersey.
 22. Lorenzen & Anderson, T and V. Anderson.(1993). *Design of Experiments:A No-Name Approach*. Marcel Dekker, New York,1993.
 23. Montgomery, D.(1991). *Design and Analysis of Experiments*. New York. John Wiley.
 24. Natrella, M.(1979) "Design and Analysis of Experiments" in *Quality Control Handbook*. ed. J.M.Juran. New York, McGraw-Hill,pg.27-35.
 25. Ragsdale, C.(1995). *Spreadsheet Modeling and Decision Analysis*. Course Technology Inc., Cambridge, Massachusetts.
 26. Romero,R., A. Ferrer, C.Capilla, L. Zunica, S. Balasch, V.Serra and R.Alcover.(1995). "Teaching Statistics to Engineers; An Innovative Pedagogical Experience." *Journal of Statistical Education*. 3,1.
 27. Sharpe, P. & T. Keelin(1998). "How SmithKline Beecham Makes Better Resource-Allocation Decisions." *Harvard Business Review*. March,45-57
 28. Simon, Herbert(1978). "Rationality as Process and as Product of Thought." *American Economic Review*. 68,2.
 29. Smith, G..(1998) "Learning Statistics by Doing Statistics." *Journal of Statistical Education*. 6, 3.
 30. Snee, R. (1993)."What's Missing in Statistical Education." *The American Statistician*. 47,149-154.
 31. Winston, W. (1994). *Operations Research:Applications and Algorithms*.Wadsworth Publishing, Belmont, California.