Organization and Decision Theory

edited by
Ira Horowitz

Organization and Decision Theory

edited by
Ira Horowitz
College of Business Administration
University of Florida

Kluwer Academic Publishers
Boston/Dordrecht/London

Recent Economic Thought Series

Warren J. Samuels, Editor
Michigan State University
East Lansing, Michigan, U.S.A.

Previously published books in the series:

Feiwel, G., *Samuelson and Neoclassical Economics*
Wade, L., *Political Economy: Modern Views*
Zimbalist, A., *Comparative Economic Systems: Recent Views*
Darity, W., *Labor Economics: Modern Views*
Jarsulic, M., *Money and Macro Policy*
Samuelson, L., *Microeconomic Theory*
Bromley, D., *Natural Resource Economics: Policy Problems and Contemporary Analysis*
Mirowski, P., *The Reconstruction of Economic Theory*
Field, A., *The Future of Economic History*
Lowry, S., *Pre-Classical Economic Thought*
Officer, L., *International Economics*
Asimakopulos, A., *Theories of Income Distribution*
Earl, P., *Psychological Economics: Development, Tensions, Prospects*
Peterson, W., *Market Power and the Economy*
Mercuro, N., *Law and Economics*

This series is devoted to works that present divergent views on the development, prospects, and tensions within some important research areas of international economic thought. Among the fields covered are macromonetary policy, public finance, labor and political economy. The emphasis of the series is on providing a critical, constructive view of each of these fields, as well as a forum through which leading scholars of international reputation may voice their perspectives on important related issues. Each volume in the series will be self-contained; together these volumes will provide dramatic evidence of the variety of economic thought within the scholarly community.

Contributing Authors

Robert P. Bostrom
College of Business Administration
University of Georgia
Athens, GA 30602

Rex V. Brown
Decision Science Consortium, Inc.
1895 Preston White Drive
Reston, VA 22091

Douglas D. Dankel II
Computer and Information Sciences
University of Florida
Gainesville, FL 32611

James S. Dyer
Graduate School of Business Administration
University of Texas
Austin, TX 78712

Jack Feldman
Department of Psychology
Georgia Institute of Technology
Atlanta, GA 30332

Jack W. Fellers
Graduate School of Business Administration
Indiana University
Bloomington, IN 47401

Ira Horowitz
College of Business Administration
University of Florida
Gainesville, FL 32611

Irving H. LaValle
A. B. Freeman School of Business
Tulane University
New Orleans, LA 70118

Michael K. Lindell
Department of Psychology
Michigan State University
East Lansing, MI 48824

Herbert Moskowitz
Krannert Graduate School of Management
Purdue University
West Lafayette, IN 47907

John W. Payne
Fuqua School of Business
Duke University
Durham, NC 27706

Rakesh K. Sarin
Fuqua School of Business
Duke University
Durham, NC 27706

Howard Thomas
Department of Business Administration
University of Illinois at Urbana-Champaign
Champaign, IL 61820

Contents

1 INTRODUCTION

Ira Horowitz

Depending upon one's perspective, the need to choose among alternatives can be an unwelcome but unavoidable responsibility, an exciting and challenging opportunity, a run-of-the-mill activity that one performs seemingly "without thinking very much about it," or perhaps something in between. Your most recent selections from a restaurant menu, from a set of jobs or job candidates, or from a rent-or-buy or sell-or-lease option, are cases in point. Oftentimes we are involved in group decision processes, such as the choice of a president, wherein one group member's unwelcome responsibility is another's exciting opportunity. Many of us that voted in the presidential elections of both 1956 and 1984, irrespective of political affiliation, experienced both emotions; others just pulled the lever or punched the card without thinking very much about it.

Arriving at either an individual or a group decision can sometimes be a time consuming, torturous, and traumatic process that results in a long-regretted choice that could have been reached right off the bat. On other occasions, the "just let's get it over with and get out of here" solution to a long-festering problem can yield rewards that are reaped for many years to come. One way or another, however, individuals and organizations somehow manage to get the decision-making job done, even if they don't quite understand, and often question, just how this was accomplished.

In recent decades, the study of decision making, in particular how individuals and groups actually go about the decision-making process, how they should go about the decision-making process in order to arrive at optimal decisions, and why the decisions of seemingly rational decision makers diverge with such disturbing regularity from what the rational and consistent decision maker ought to do, has taken on all the appearances of a growth industry that has generated employment for, among others, economists, psychologists, statisticians, computer scientists, and mathematicians. This book is concerned with all three of the aforementioned issues—the how we do, the how we should, and the why we don't issues—particularly as they pertain to decision making in organizations.

The book does not attempt to be comprehensive in its coverage, and, in the main, the individual chapters do not start from scratch. It was felt, however, that the final chapter, which deals with the so-called expert systems that are now coming on line, and gets into technical issues that might be somewhat removed from the backgrounds of the majority of readers, would make an appropriate exception and should include more of the basics. Our intention, then, is to provide insights into particular aspects of decision making in organizations, with each of the five authors and the discussants providing his own, perhaps unique perspective, be it that of the psychologist, say, or that of the computer scientist. The common thread that links these papers is their focus on the primary ingredients in any decision problem: notably, the needs to 1) define the problem and determine the decision maker's options, in the process reducing the latter to a manageable few, 2) anticipate the potential outcomes associated with each of the decision alternatives and assess their relative likelihoods of occurrence, and 3) recognize the consequences of any alternative-outcome pair, both for the organization and for the individual making the decision, and translate these consequences into terms that will be meaningful for decision-making purposes. A fourth need arises from the observation that few organizational decisions are made in a vacuum, and many (if not most) involve nonhomogeneous groups, and/or affect other organizations whose possible reactions and (perhaps) conflicting objectives must also be anticipated and incorporated into the decision-making process.

The study of decision making, either descriptively, prescriptively, or at the junction or the separation of the twains, is in its own right an interesting intellectual exercise. And, from a practical standpoint, those of us concerned with becoming better decision makers and with reaching better decisions, might benefit from that exercise, insofar as it calls our attention to 1) some of the things that we do or don't do in the decision-making process; and the foibles that we share with others, and 2) some of the things that we should seriously consider doing if we really want to improve.

But do we really want to improve? A quarter of a century ago, when I still believed in the tooth fairy and was only beginning to harbor doubts about Santa Claus, I was engaged as a consultant by the management of a subsidiary of a major U.S. corporation that, along with the corporation's other subsidiaries, had been asked to use a management science approach in solving a specific well-defined problem that was common to all of the subsidiaries. The results of these individual efforts would then be shared at a conference at which each subsidiary would be represented by its president. Lacking the requisite in-house expertise in management science, the management of one subsidiary asked me for help, and help them I did. I immediately recognized how the problem should be solved, and I identified the appropriate one of my doctoral students to do the dirty work. Knowing whom to call upon and when to do so is one of the hallmarks of the good decision maker who seeks to practice as well as preach. Together, the student and I developed a computer-based decision model that not only provided the answer to one of the subsidiaries' critical and recurring questions, but was user friendly to boot. The management for whom I was working was delighted, the president's conference presentation was exceptionally well received, and the subsidiary's local office was subsequently visited by representatives of other subsidiaries in order that they might learn more of the details of what we had done. I didn't stop patting myself on the back until after the last of several follow-up calls I made to inquire as to whether our model had as yet been implemented, and to ensure that we did in fact have a satisfied customer. What I received repeatedly for answers was a "no" and a "yes," in that order, as well as a bit of an education that dispelled one of the last of my remaining illusions.

It seems that, notwithstanding management's satisfaction with our efforts, the people responsible for decision making in the area, those managers that would actually be using our model, were at least as satisfied with the way they had been dealing with the problem—the same way in which the company had dealt with the problem for the several decades before we and our model came on the scene—as they were with our approach. The arguments in favor of our model were unassailable: its use only required a technician to update and feed the computer some readily available economic and demographic data that were published annually by the state and the federal governments, and a manager to implement the computer's rapidly generated profit–maximizing recommendations. The arguments against using the model were apparently more compelling: in the first place, it just wasn't much fun to use and as a result it made the manager's job somewhat less interesting; and, in the second place, the company was, quite literally, making more money than they knew what to do with, without our model and without our help, and in the absence of on-the-scene stockholders' surveillance management couldn't see any particular reason to rock the boat.

Was this corporate management unique to the sixties? I very much doubt it. Would this management be unique to the eighties? I doubt that too, but we have indeed come a long way in the last twenty-five years. *Today* how many corporate managers are willing to delegate their decision-making authority to a computer? *Today* how many corporate managers are willing to employ a computer as a consultant that provides expert advice and opinions? The answer to both questions, I suspect, is a growing number of them, and in a growing number of respects. The simple fact of the matter is that present-day organizations are very complex operations whose decisions, throughout the managerial hierarchy, require the processing of ever-increasing volumes of data and information. As a result, it has become increasingly necessary for managers at all levels of the hierarchy to delegate some decision-making authority, be it to a living, breathing subordinate, or to a cruelly efficient and indifferent computer. And, with the spread of the computer and our ability to use it in more sophisticated ways, it is becoming increasingly feasible to employ it as an expert system capable of mirroring the expertise of humans who cannot be in all places at all times, and that has the capacity to process data and "think" more rapidly than can any human expert. In the latter regard, the computer *assists* the decision maker in the decision-making process.

For example, a decade ago most of us made our long-distance travel plans either by consulting an airline guide or individual flight schedules, or else by phoning a travel agent who would consult the guide for us. Today, however, what with a plethora of fares for any given flight on any given day, frequent-flyer plans, and alternative routes to reach the same destination, it is almost de rigueur to phone either an agent or a preferred airline, lay out some guidelines, and then to rely on that source's computers to reduce the total set of options either to those that are feasible, to a manageable few, or even, perhaps, to select both the optimal flight plan and the seat(s) in which one will travel. Indeed, in Paris, at the Jardin de Luxembourg, Monsieur L'Ordinateur—a computer—seemingly delights in telling passersby how best to get from that location to any other place in Paris, given their preferences for bus or train, time considerations, their willingness to walk, and so forth. Monsieur L'Ordinateur will even correct one's (minor) spelling mistakes, or chide the person that attempts to fool "him." Shouldn't present-day executives have available to them the same technological advances to aid *their* corporate decision making as are available to the pedestrians of Paris for solving their somewhat more mundane personal transportation problems?

The need to delegate decision-making authority is especially prevalent where the decisions involved are the repetitive, relatively mundane decisions of the organizations's daily routine. The latter decisions, such as an inventory-stocking decision or a parts-replacement decision, will ordinarily be made through the use of some time-honored rules of thumb or through some more sophisticated mathematical algorithms. In such cases, all that is required to reach the appropriate decision is an understanding of how to implement the rule or the algorithm, and once this understanding is achieved, everybody working with the rule or algorithm, including a properly programmed computer, would reach the same decision. In these cases, then, management can delegate decision-making authority, confident that *its* preferred decision will be reached, even without its hands-on input. Where management has a problem are with those nonrepetitive, nonmundane situations in which it is either forced to delegate decision-making authority, say because of time constraints, or it would prefer to delegate decision-making authority, say to be able to devote more attention to what it judges to be even more interesting or more challenging matters. In the latter situations, too, management would like to be able to delegate decision-making authority with the confidence that the delegate will make *management's* preferred decision, rather than that which he, or she, or, in the case of a computer, it prefers. One of the principal virtues—if not the principal virtue—of the quantitative approach to decision making is that it permits management the latter luxury.

In particular, the decision-theoretic approach to decision making seeks to quantify management's preferences and judgments in a manner that will permit these to be incorporated as inputs, along with the relevant environmental data, into an information processing mechanism that is designed to yield as its outputs a numerical ranking, or management's preference ordering, of the decision alternatives. That is, once the decision maker's preferences and judgments have been elicited and quantified, the processing mechanism produces from the problem-specific data a set of numbers that serve as signals of the decision maker's preferred alternatives. The processing mechanism, then, is an impersonal "black box" that decision makers can use directly to their own advantage when making decisions, or, alternatively, that they can rely on indirectly to signal their preferences when they have delegated decision-making authority to others. In either event, the black box has the appealing quality that it doesn't make mistakes; it always signals the delegator's preferred decision. This is not to say that after the fact one might not regret having taken the decision that was signalled as being optimal. In a world of uncertainty there are no guarantees. Rather, reliance on the black box permits us to suffer our regrets with the often-comforting knowledge that, given the uncertainties and imperfect knowledge of the world in which we live, we did just about as well as we could have done at the time, and under the same set of circumstances we would do the same thing again with, we would hope, a more rewarding spin of the wheel and result.

The original mechanism for the black box, which is associated with the names von Neumann and Morgenstern, produces as its decision-ranking signals a set of numbers called expected utilities, a somewhat unfortunate appellation that can conjure up a variety of meanings. In this context, the word *expected* means a mathematical expectation, or a statistical weighted average of a set of numbers, with each number's probability of occurrence being its assigned weight. What is being averaged are the *utilities*, which are arbitrary numbers assigned to reflect an individual's ordinal ranking of an entity in his or her preference ordering, *and* that individual's attitude towards risk.

For example, suppose apples are the fruit that I most prefer, and pears are the fruit that I least prefer. The mechanism permits me to arbitrarily assign apples the number 8 and pears the number 4. The specific numbers that I choose to assign are unimportant, although some numbers—in particular, a one and a zero—will usually be more convenient to work with than others. The only condition imposed upon me is that I assign a higher number to an apple than to a pear, because I prefer apples to pears. Under the latter circumstance, once the rest of the world sees my numerical assignments it can immediately discover that I prefer apples to pears. What the rest of the world does not learn is the extent to which I prefer apples to pears, nor, indeed, whether I actually detest all fruit, but apples somewhat less than pears. That is, in and of themselves the 8 and the 4 convey nothing about the strength of my preferences.

Suppose, now, that placed before me are two barrels, one of which is loaded with apples and pears, and the other of which contains nothing but peaches. Suppose, too, that I am going to be given the opportunity to reach into the barrel of my choice one time, but blindfolded, and I will receive for my efforts and cooperative spirit the first piece of fruit that I touch. From which barrel should I choose to draw? By my previously expressed attitude towards fruit I must feel at least as kindly to peaches as I do to pears, but I cannot prefer a peach to an apple. Therefore, the answer to my decision problem would seem to hinge on the relative number of apples and pears in the barrel that contains both, because I will without question get a peach if I should reach into the barrel that contains only peaches.

The way that von Neumann and Morgenstern would approach the problem is not to ask me from which barrel I prefer to draw, but rather to ask me how many apples and how many pears have to be in the barrel that contains both in order to make me indifferent as to the barrel into which I am going to be reaching. If my answer is "ten apples and ten pears," then von Neumann and Morgenstern would infer that the number 6 should be assigned to reflect my preference for peaches. The reasoning behind this assignment is that I have indicated that I am indifferent between the certainty of a peach and having a fifty-fifty chance of an apple or a pear, or what is equivalent a fifty-fifty chance of their associated "utilities" of 8 and 4, which average out to be a 6. Von Neumann and Morgenstern's system *demands*, then, that the certainty of a peach be assigned a "utility" of 6. Had I instead responded "seventy-five apples and twenty-five pears," the corresponding probabilities of ¾ for an apple or an 8, and ¼ for a pear or a 4, would have resulted in a peach being associated with the number $(3/4)8 + (1/4)4 = 7$. The latter number, which is my assigned "utility" of a peach, immediately tells the outside observer—*any* outside observer who understands the system—that I prefer apples to peaches to pears. Given the latter assignment, if I could draw from yet a third barrel that contains thirty apples, twenty pears, and fifty peaches, the expected utility of being blindfolded and being allowed to draw one piece of fruit from that barrel would be $(.3)8 + (.5)7 + (.2)4 = 6.7$, since, for example, there is a .5 probability of drawing a peach and its associated utility of 7. What von Neumann and Morgenstern showed is that as a rational decision maker, I would prefer to reach into the barrel that contained all peaches than to reach into this third barrel, and that this preference has been *signalled* by the fact that I would assign a higher expected utility (7) to the all-peach barrel than I would to the three-fruit barrel (6.7). Moreover, now that I have presented for all the world to see an ordinal numerical ranking of my preferences, irrespective of the combination of apples, peaches, and pears in the next barrel that I will into existence, I can delegate my decision-making authority to anybody or any computer that has been instructed in how the system works, and that person or computer can immediately rank my preferences between the various combinations, with the preferred one signalled as that which maximizes my expected utility.

There is to be sure a trick in all of this. In order to show that the expected utility measure accurately reflects the decision maker's ordinal ranking of alternatives, von Neumann and Morgenstern had to make a series of assumptions as to what constitutes rationality in decision making. Their assumptions were summarized in a series of axioms or rules for rational behavior. In the previous example, these axioms would state, among other things, that I know my own mind and that I can say whether I prefer apples to peaches, say, and that once I assert my preference for apples to peaches and go on to assert a preference for peaches to pears, I would not further assert a preference for pears to apples. Consistency may be "the hobgoblin of little minds," but for most of us it is a desideratum of a decision-making framework. The axioms would also state that I derive no particular pleasure from reaching into barrels; my only concern, at least in the previous example, is with the piece of fruit that I am going to enjoy at the completion of the decision-making exercise.

If one wants to be rational in the von Neumann-Morgenstern sense, and if one would prefer that his or her decisions always satisfy the von Neumann-Morgenstern axioms of rational behavior, then one should always compute and take the decision that maximizes expected utility, because that decision has been signalled as the preferred decision—preferred under *their* set of rules. Put otherwise, the ball game is over and we can all go home. The reason that the game continues and that we can't all go home is that decision makers do not necessarily obey the rules, even when they are aware of the rules, understand and subscribe to the rules, and would like their decision making to be coldly rational and adhere to the rules; nor, if the whole truth be known, do they necessarily want to obey the rules—always, if indeed ever. That is, decision makers persist in acting as individuals, with all of the individual's warts and blemishes. And so the game plays on, with theorists developing variations on the von Neumann-Morgenstern theme, which commonly consist of a modified set of axioms of rational behavior, modified with the goals of more accurately reflecting human foibles and peccadilloes, and with behaviorists providing the grist for the theorists' mills by studying decision makers at work, observing and describing the latters' penchants for violating one or more of the latest proffered set of axioms of rational behavior, and speculating as to the further modifications that will be necessary to create a more accessible, acceptable, and applicable framework for real-world decision making, one principal element of which is the quantification of individual and group preferences.

A second critical element is the quantification of judgment, which implies a need to improve our understanding of how individuals and groups process information in a world of uncertainty, and how these judgments are revised as new information is received. Suppose, to use a variation of a well-known example, that the barrel containing the apples and pears is large enough to hold literally billions of pieces of fruit, and has been filled to the brim with equal amounts of each of the two varieties. What are the chances that one random draw from the barrel will produce an apple? It's the rare person, even one that has had only a passing acquaintance with probability, that fails to agree that the only reasonable response is 1/2. That is, each of us would correctly translate the expression "the chances that" into probabilistic terms and would quantify our judgments into the same probability statement of "one half." The follow-up question is: What are the chances that two successive random draws from the barrel will produce two apples as its prize? Again, virtually everybody can translate this question correctly, and most people will go on to answer "one out of four," with 1/4 being the objectively correct answer provided the barrel is big enough and does indeed contain billions upon billions of apples and pears.

Now, however, suppose one is given a bit of information: notably, two pieces of fruit have been drawn at random from the barrel, and a surreptitious peek reveals that one of these pieces of fruit is an apple. What are the chances that the other is also an apple? My personal and rather extensive experience is that most people to whom I pose the question quantify their judgments into the probabilistic responses of either 1/2 or 1/4, and a substantial proportion remain annoyed with me even after hearing my simple explanation as to why the *only* reasonable answer is 1/3! The explanation is that when two pieces of fruit are to be drawn at random from the barrel, there are four equally likely possibilities: either both pieces of fruit will be pears, both will be apples, or the first piece will be an apple and the second a pear, or the first piece will be a pear and the second an apple. Given that my surreptitious peek revealed that at least one of the pieces of fruit in the drawn pair was in fact an apple, the first of the four combinations that was possible *prior* to the drawing could not have been drawn. Therefore, *posterior* to the drawing only one of the three equally likely possibilities satisfies the condition of providing two apples, so that the probability that the second piece of fruit is also an apple is one out of three, or 1/3.

The answer of "one third" can also be obtained through the use of Bayes' Rule, a systematic procedure for revising one's quantified judgments—probability statements—in the light of additional information. In the previous example, one would initially make the judgment that the chances of drawing two apples in successive random draws would be one out of four, and this judgment would have to be revised in light of the information that at least one of the draws has yielded an apple, a sample outcome that on average will occur three out of every four times that two pieces of fruit are drawn from the barrel. Bayes' Rule is an important element in the quantification of judgment because, experimental evidence has revealed, in the absence of this systematic data processor, individuals tend not to appropriately use the additional information that they are given. The latter failure compounds the potential problems that can arise when we attempt to quantify our judgments prior to receiving this additional information. The latter problems include the difficulty encountered in assigning probabilities to either rare outcomes or those that are considered to be virtual certainties, as well as the potential for biasing one's probability assignments, either intentionally or subconsciously. The salesperson that assigns a low probability of making a sale in order to look good to the boss, when in reality the sale had already been consummated, is a case in point; the salesperson that assigns a low probability of making a sale so as not to feel too badly when rebuffed is another. Yet, the salesperson's judgments, as inputs into a decision process, can be extremely influential in management's decision making. Thus, still another problem is how to get the individuals whose probability judgments will be used in the decision-making process to reveal their true beliefs, irrespective of how good or bad these people are as probability assessors—good or bad as indicated by the extent to which, over the long haul, their quantified judgments seem to bear some relationship to what has actually taken place.

As divine as forgiving may be, the business world has an overwhelming tendency to punish error, and to err is indeed human in a world fraught with uncertainty. Of course there are people that always land on their feet no matter how ridiculous their decisions seem to the rest of us to be, and no matter how cavalier are the decision processes that have been employed to arrive at those seemingly ridiculous decisions. In a world of uncertainty, where a monkey left alone in a room with a word processor and given enough time would eventually produce the Bible, the occasional person that always lucks out is a veritable certainty. More to the point, however, is Damon Runyon's astute observation that "The race is not always to the swift, nor the battle to the strong—but that's the way to bet." The business world demands that its decision makers be both strong and swift, and the cynics notwithstanding, by and large they are. The general problem facing the decision sciences is to provide both the tools and training to help decision makers throughout the organizational hierarchy to achieve their aspirations, both in the short term and the long, given their beliefs, their judgments, and the constraints under which they operate. In the subsequent chapters our authors and commentators will explore where we've been, where we are, where we're headed, and what sort of problems we will be facing down the road, in our efforts to make the decision makers of the future both stronger and swifter.