Decision Making Using Multi-attributed Alternatives: Raiffa's Contributions in the Context of 21st Century Decision Problems

Fred S. Roberts[1]

DIMACS Center, Rutgers University, Piscataway, NJ 08854 USA

froberts@dimacs.rutgers.edu

Draft of August 4, 2007

Abstract

The pioneering work of Howard Raiffa in [47] introduced methods for analyzing preferences among multi-attributed alternatives. This paper describes the author's interaction with the early stages of that paper. It compares and contrasts the Raiffa point of view with that in the theory of measurement and discusses the key ideas of product structures and lotteries. It then discusses Raiffa's ideas in the context of decision problems in the 21st century.

[1]The author gratefully acknowledges the support of the National Science Foundation under grants SES-0351165 and INT-0339067 to Rutgers University.

1 Introduction

A few years ago, I spent a lot of time thinking about the problem of locating sensors to give early warning of a bioterrorist attack.[2] In the United States, such sensors have been placed in areas around important events, for example near the venues for the 2002 Salt Lake City Olympics or in New York City for the visit of President George Bush on the first anniversary of the September 11, 2001 terrorist attack on the World Trade Center. A program to locate them to protect urban areas is being implemented. However, a systematic way of locating such sensors was (and is) called for. In formulating the problem precisely, it is natural to think of it as a multi-attributed problem. For a given configuration of sensors in an urban area, one needs to consider the probability of detection of a bioterrorist attack, the probability of a false alarm, the speed with which an attack is detected, the geographical area covered, the cost of investing in the sensor configuration, the maintenance cost, etc. Indeed, one might consider some of these factors separately for each of a number of different bioterrorist agents. It seemed natural to think of this as a problem of comparing alternative configurations with a number of attributes. It is hard to remember that in the 1960's, prior to the pioneering work of Howard Raiffa and several of his contemporaries, this kind of analysis would not have occurred to people nor would there have been tools available to have made it. One of the key contributions to the development of this kind of analysis and of tools for making it was contained in the fundamental paper by Raiffa [47], "Preferences for Multi-attributed Alternatives," published as a RAND Corporation Research Memorandum. In this paper, I will describe some of the history of Raiffa's paper (herein to be abbreviated **Raiffa69**) and the context in which it was written, discuss some of the key ideas in the paper, and comment on modern directions/challenges that were not foreseen at the time it was written.

2 A Young Scientist Confronts his First Refereeing Job

Let me start with a personal comment. My first job after graduate school and a postdoctoral fellowship was in the Mathematics Department at the RAND Corporation in Santa Monica, California. RAND was a "think tank" and its Mathematics Department was a lot like an academic department. In particular, publication was encouraged. At RAND, external technical reports such as research memoranda could not be published until they were internally refereed. My first internal refereeing job as a member of the professional staff at RAND was to referee a preliminary version of Raiffa69. Refereeing was considered a very important job of members of the professional staff and, indeed, quantity and quality of refereeing jobs was taken into consideration in semi-annual salary reviews. What made this internal refereeing especially tricky was that it was not blind. So I, as a young scientist fresh out of graduate school, was faced with having to evaluate a paper by a distinguished professor from Harvard and one that I had some serious misgivings about. What was I to do? After, I admit, some serious hemming and hawing, I elected to lay out my reservations. I describe some of these in Section 4. I was delighted (and much relieved) at the response from Raiffa (such responses also being required as part of the vetting process for external RAND reports): "The com-

[2] See [49] for a formulation of a number of analytic challenges arising from problems of bioterrorism.

ments by Roberts are sharp, to the point, and I concur with most of what he says. Indeed, if I had to review my own paper, I would essentially say the same thing about it." Professor Raiffa was gracious about my comments and gave careful, thoughtful responses to them. I have used this example as an object lesson many times over the years, both for my students and my children.

3 Putting the Paper in Context: The RAND Scene

To understand the context in which this paper was prepared, one has to understand the "scene" at the RAND Corporation in the 1960's. The Mathematics Department had about 35 "mathematical scientists." I use this term loosely, since their expertise ranged over economics, operations research, game theory, linear programming, graph theory, statistics, decision theory and "policy science." Many members of the Department played the role of internal consultants. They were encouraged to do their own foundational research, while being asked to listen to problems faced by researchers throughout the Corporation. Many famous researchers were brought in for the summer or longer periods to visit. While I was there, such luminaries as George Dantzig, Alan Hoffman, Victor Klee, A.W. Tucker, and William T. Tutte, as well as Raiffa, spent time at RAND. They in turn were encouraged to think broadly about "big problems" and "new methods." Indeed, in earlier years, RAND was the birthplace of foundational ideas in mathematical programming, game theory, and other areas of research, not to mention highly speculative and prescient ideas, e.g., about earth-circling satellites [45].

However, some things were changing in the late 1960's. The company, which had been primarily a research arm of the U.S. Air Force, was branching out to work on a variety of problems such as health care, energy use, pollution, and urban problems. These are all areas that engaged me and laid the foundation for many enduring interests. Not the least of these was transportation. Raiffa was brought in under the auspices of the Northeast Corridor Transportation Project, Office of High Speed Ground Transportation, U.S. Department of Transportation. The overarching goal of the project was to develop methods for evaluating alternative transportation proposals. This was a natural application for Raiffa, since the problems dealt with comparing the utilities of complex, multi-attributed alternatives. Prior to Raiffa69 and his predecessor book [46], perhaps the main area of application of utility and decision theory had been to business decisions. Now, a whole new horizon was opening up for these methods. It was natural for Raiffa to develop this further while a consultant to RAND.

4 Seeing Raiffa's Draft

When I first read the draft version of Raiffa69, I was immediately impressed with the no-nonsense, prescriptive point of view: We have to make decisions, so here are some useful techniques that will help us make better ones. At the same time, this somewhat informal attitude seemed to be both the strength and weakness of the paper. Indeed, this seemed in part to be just a "bag of tricks," as I said in my review, though each seemed "amusing" and "meaty," as I also said. They whetted the appetite of the reader and, I presumed, could each be expanded into a sizeable topic for research. I suggested that the paper be expanded to go more in depth into several topics. Indeed, that is just what did happen - though not as I suggested, since RAND felt that the paper was important and should come out quickly. What did happen, of course, was that this foundational paper led to a large amount of exciting and important follow-up work over a long period of time.

I had been heavily influenced in graduate school by the descriptive approach to decision theory. This emphasized axioms that, if satisfied, would allow us to use certain methods for making decisions. Along with the axioms came formally stated theorems that showed that, under the axioms, a specific procedure was specified. I was disappointed that Raiffa didn't take a more axiomatic approach and, even more, that the paper was so informal and he didn't formulate his results/ideas more formally as theorems. In retrospect, of course, this would have detracted greatly from the impact of Raiffa's work, which depended on the use of mathematical language to formulate ideas and approaches, but did not get bogged down in excessive formalism. The published version of Raiffa69 did, however, in important places, lay out conditions that had to hold for the methods described to work.

5 Measurement Theory and Raiffa69

Being exposed to Raiffa's paper at an early stage had a significant effect on my own way of thinking about preference, utility, and decision making. A few years later, when I wrote a book entitled *Measurement Theory, with Applications to Utility, Decisionmaking, and the Social Sciences* [48], the influence of Raiffa69 and his related work in his books [46] and [27] was clearly evident.

The theory of measurement was developed by mathematical social scientists (especially mathematical psychologists), philosophers of science, economists, and others, with the goal of putting the way we measure things on a firm mathematical foundation. It started from an attempt to understand what was involved in physical measurement, e.g., of temperature or mass, and moved from there to psychological measurement, e.g., measurement of noise, or more precisely loudness, or of utility. Some of the key figures in this development were S.S. Stevens, Pat Suppes, Duncan Luce, Dave Krantz, Amos Tversky, and Peter Fishburn. Important early papers were due to Stevens [53, 54, 55], and Suppes and Zinnes [57] and an important early book was by Pfanzagl [44]. A critical foundational piece of work here was the book by Krantz, Luce, Suppes, and Tversky [29] (which had two following volumes [31, 56]). Much of the motivation for this work came from psychophysics and was concerned with judgments of loudness, brightness, etc. However, preferences and utility were a major motivation as well. Both the measurement theorists and Raiffa were influenced by the work of such mathematical economists as Gerard Debreu. Indeed, it was Debreu's result about what is sometimes called conjoint measurement that was one of the early attempts to deal with "product structures" for alternatives, something that Raiffa69 built on and that was also built on by Luce and Tukey [32] in the measurement context. One major difference between Raiffa's approach and that of the measurement theorists, however, was that, similar to the emphasis on the descriptive approach to decision theory described in the previous section, the measurement theorists put a great deal of emphasis on an axiomatic foundation for their subject, including preference and utility, while Raiffa was more interested in methods and tools. Raiffa was interested in laying out *prescriptions* for decision makers. The measurement theorists were more concerned with laying out conditions that, if satisfied, allowed one to represent preferences with numbers, interpreted as utilities, which could in turn be used to make decisions. The issue for them then became one of testing those conditions, or axioms, to see if real-world decision makers satisfied them.

My own exposure to the theory of measurement started in graduate school, where I studied under Pat Suppes and then, through him, met Duncan Luce, with whom I later did a postdoc. In graduate school, I also met Dave Krantz and Amos Tversky, and soon after arriving at RAND, I met Peter Fishburn. These individuals heavily influenced my own thinking. However, Raiffa's work, and in particular my exposure to Raiffa69, took me in new directions and exposed me to new applications. My own approach to the theory of measurement was and has been a mixture of emphasis on axioms and practical applications. I very much owe the latter to Raiffa.

6 One Key Idea: Product Structures

A key idea in Raiffa69 is to express a complex alternative as a vector whose i^{th} component x_i is chosen from a set A_i. In the transportation example, A_1 can be the set of possible power sources, A_2 the set of possible vehicle designs, A_3 the set of possible "right of way" designs, etc. Thus, a complex alternative can be thought of as a vector $(x_1, x_2, \ldots, x_n)$ from the Cartesian product $A_1 \times A_2 \times \ldots A_n$, which is called a **product structure**. The sets A_i are often given ahead of time. However, it is often a challenge to find such a product structure. Raiffa69 built on a hierarchical approach to building such a structure that was heavily borrowed from Miller [35] and Mannheim and Hall [33]. Since some product structures have a great many dimensions, i.e., since n can be very large, it is important to develop methods for reducing the dimension, i.e., replacing a product structure with one with smaller n. Raiffa69 discussed methods for doing this by fixing an element $a*$ in one of the sets A_i, say A_1. Then, given a vector $(x_1, x_2, \ldots, x_n)$, one finds an element b in another set A_j, say A_2, that makes the decision maker indifferent between $(x_1, x_2, \ldots, x_n)$ and $(a*, b, x_3 \ldots, x_n)$. This idea, while simple, is still of fundamental importance in decision theory.

In many applications, it is easier to work with product structures where each set A_i is a set of numbers. Raiffa69 concentrated on this special case. The work of Debreu [5] and Luce and Tukey [32] referred to above deals with preferences among vectors from a product structure. If $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ and $\mathbf{y} = (y_1, y_2, \ldots, y_n)$, let $\mathbf{x}R\mathbf{y}$ mean that an alternative corresponding to $\mathbf{x}$ is preferred to an alternative corresponding to $\mathbf{y}$. An ordinal utility function is then a real-valued function F on $A_1 \times A_2 \times \ldots \times A_n$ so that

$$(x_1, x_2, \ldots, x_n)R(y_1, y_2, \ldots, y_n) \leftrightarrow F(x_1, x_2, \ldots, x_n) > F(y_1, y_2, \ldots, y_n). \tag{1}$$

More generally, if we have a space A and a binary relation R on A, where aRb means that a is preferred to b, then an *ordinal utility function* is a real-valued function F on A so that

$$aRb \leftrightarrow F(a) > F(b).$$

In the case of product structures, Debreu and Luce and Tukey sought functions f_i on A_i so that for the ordinal utility function F,

$$F(x_1, x_2, \ldots, x_n) = f_1(x_1) + f_2(x_2) + \ldots + f_n(x_n), \tag{2}$$

i.e., so that

$$(x_1, x_2, \ldots, x_n)R(y_1, y_2, \ldots, y_n) \leftrightarrow f_1(x_1)+f_2(x_2)+\ldots+f_n(x_n) > f_1(y_1)+f_2(y_2)+\ldots+f_n(y_n). \tag{3}$$

In this case, we can calculate utility over a product structure by calculating individual utilities or scores over each component. They sought conditions – axioms – on the binary relation R so that functions $f_1, f_2, \ldots, f_n$ satisfying Equation 3 exist. Raiffa69, by contrast, assumed we have an ordinal utility function F, and then asked when we could decompose it in such a way as to satisfy Equation 2. Building on the work of Fishburn [10, 11, 12, 13, 14] and Miller [35], he described the use of lotteries to give conditions under which this could be done. (More on lotteries in the next section.) It is typical of Raiffa to prescribe that we be "rational" in our decision making, to assume that we have an ordinal utility function, and to seek out conditions that will allow us to simplify calculation of such a function. Debreu and Luce and Tukey, by contrast, pushed the conditions all the way back to the issue of whether or not our preferences are rational enough to satisfy the ordinal utility model.

More generally, one can ask if a utility function F can be decomposed as

$$F(x_1, x_2, \ldots, x_n) = G[f_1(x_1), f_2(x_2), \ldots, f_n(x_n)].$$

In the special case where G is a polynomial and F satisfies Equation 1, we say that we have *polynomial conjoint measurement*. An early discussion of polynomial conjoint measurement appears in Krantz, et al. [29]. However, axioms giving conditions for polynomial conjoint measurement are not known for many polynomials nor is this idea much used in decision making applications beyond some simple polynomials.

Raiffa69 also considered the possibility that a utility function might not be representable in the additive sense described in Equation 2. He considered more complex decompositions, in particular for $n = 2$:

$$F(x_1, x_2) = f_1(x_1) + f_2(x_2) + g_1(x_1)g_2(x_2). \tag{4}$$

In this case, the third term represents some sort of interaction effect. Raiffa69, Fishburn [15] and Farquhar [9] studied conditions under which one can find this kind of representation. Note that this *quasi additive* representation is not polynomial conjoint measurement. Raiffa's approach to this problem highlighted an approach to decision making not much treated in the literature prior to his work. Not long after the publication of Raiffa69, the quasi additive representation began to be used in many applications, e.g., transportation [6, 7, 25], medical decision making [24], management [20, 59], solid waste disposal [4], air pollution [8], and urban services [26].

7 Another Key Idea: Lotteries

A second key idea discussed in Raiffa69 is the notion of a lottery. Here, one thinks of an alternative x as follows. There are several possible outcomes/consequences $c_1, c_2, \ldots, c_n$ if one chooses the alternative x, and c_i occurs with some probability $p(c_i)$ (probabilities adding to one). Moreover, each consequence c_i has utility $u(c_i)$. Then the *expected utility* of c is given by

$$\sum p(c_i)u(c_i).$$

We anticipate choosing the alternative with the highest expected utility. Note again that the question of axioms under which this "expected utility rule" holds was not the emphasis of Raiffa69. This is in contast to a long line of work in decision theory by many authors, such as L.J. Savage (who, by the way, was the other internal referee of Raiffa69). Raiffa69 asked whether we could find the ordinal utility function u on consequences from knowledge of preferences among lotteries. He described a method that works if one can identify a best consequence K and a worst consequence k. One then seeks for any consequence c a lottery L_p in which one receives K with probability p and k with probability $1 - p$ and so that the decision maker is indifferent between having c for sure and entering the lottery L_p. Let us denote p by $p(c)$. Now one notes that if one assumes that decisions are being made using expected utilities, then $p(c) > p(d) \leftrightarrow u(c) > u(d)$. In other words, if $u(c)$ is an ordinal utility function over the space of consequences, then so is $p(c)$. Note that we do not need to know the utility function u to calculate the utility function p nor do we need to calculate expected utilities. This simple idea is really very clever – typical of the "simple" ideas due to Raiffa that have been so useful in decision theory. Raiffa admitted that the idea is not all that practical. However, Keeney and Raiffa [27] noted that it is really only necessary to calculate $p(c)$ for a few values of c and then to fit a curve to these data points. For a discussion of axioms under which this procedure works, see [48], Theorem 7.1. In contrast to this more formal statement in [48], finding and describing the procedure was a sufficient goal for Raiffa.

The idea of using the lottery L_p can be extended to lead to practically-useful methods. Raiffa69 called the lottery L_p the *p-basic reference lottery ticket* or *p-brlt*. He then noted that we can reduce more complicated lotteries to "equivalent" p-brlts. If lotteries R and S are equivalent to p-brlts L_r and L_s respectively, we conclude that R is preferred to S if and only if the number r is greater than the number s. In this way, our preferences among complex decision making problems are reduced to comparison between two numbers – a remarkably simple, yet very clever idea.

One of the most complex applications of Raiffa's ideas on decision making involves applications to medical decision making and public health. I was personally fascinated by the discussion in Raiffa69 of the "value of a life." Here he used the notion of p-brlt to explore questions about public health. He asked questions like: How much would you be willing to give up in resources to reduce the probability of dying from p_1 to p_2. In my book on measurement theory [48], I posed the question this way: Suppose a study has shown that spending x dollars over the next decade can reduce the proportion of people who die from breast cancer from a to b or the proportion of people who die from air pollution from a to *b*. How does society decide whether to spend the x dollars? The question involves tradeoffs between health and resources or between spending on health and spending on other things. Lotteries help in this regard. They lead for example to the conclusion that if it is worth x/N dollars (where N is the size of the population) to reduce deaths from air pollution from 100 people per million to approximately 90 per million, then it should be worth exactly the same amount to reduce the figure from 10 per million to approximately 0 per million. (See [48], Sec. 7.2.6, for details.) This conclusion, if true, should be considered in light of the argument that cutting the probability of death from a particular cause those last few steps down to a very small amount is sometimes proportionately more expensive. Similar comments apply to the bioterrorism example discussed in the Introduction, or more generally to homeland security. We are today investing a large amount of societal resources to prevent acts of terrorism. What would a Raiffa-like analysis say about the societal expenses to cut the probability of a terrorist attack down to very small levels?

8 Raiffa's Ideas in a 21st Century Context

Let me expand on this last point. Raiffa's work, indeed much of decision theory, is heavily involved with the interplay between utility and probability. Yet, in the counter-terrorism domain, we are dealing with unusual situations: low probability, high consequence events. It is widely believed that decision making involving these kinds of events will require new kinds of analysis. The "old standby" methods such as those developed in other contexts might not apply to this kind of situation.

There are other characteristics of decision making problems in the 21st century that surely call for new methods. Raiffa was well aware that dealing with too many alternatives made it difficult to make decisions. We have already mentioned his attention to methods for reducing the number of dimensions in a product structure. However, at the time Raiffa69 was written, could anyone have imagined the massive amounts of data that modern decision theorists are presented with? Dimension reduction remains a major area of research today. However, now we are talking about dimension reduction that is orders of magnitude beyond what Raiffa was dealing with when he reduced dimension one step at a time. The problems of dimension reduction faced by today's decision makers are daunting indeed. (For some modern approaches to dimension reduction, see for example [21, 51, 58, 60].)

The time allowed for making decisions has changed in many ways too. We recently had a workshop on auctions with transactions costs at the DIMACS Center where I work. While auctions on eBay might take from three to ten days, I learned to my amazement that the typical "Dutch auction" at flower markets is completed in an average of four seconds! (See [23], p. 91.) Modern decision makers have to decide whether or not to bid in time that a human being would not be able to handle. Increasingly, decisions are made by software "agents" that are programmed to take into account probabilities and utilities. In competitive situations, such as those arising in the economic marketplace, these agents also "learn" the utilities of their "opponents" by entering many auctions simultaneously, often just for the purpose of learning and thus to enhance their ability to "win" in later "plays" of the "game." Since some decision making tasks may be delegated to software agents, one needs to understand how these agents should be designed to "learn" user preferences and how to bid and trade. There is also a clear need for a theory of decision related to the rapid transaction opportunities provided by the Internet.

Today, support for decision making is frequently provided by computing systems of various kinds. Such systems increasingly consist of numerous computing devices that must learn about themselves and their environment in situations of dynamic change. Learning has traditionally played an important role in decision support and in particular machine learning by a single agent in a relatively static environment has been well studied. Theoretical computer science has played an important role in the development of learning theory. However, today's complex decision making problems involve numerous agents that must continually adapt to their environment and each other. There are very few theoretical results to guide us in understanding learning by multi-agent systems.[3]

As these examples illustrate, computer science has clearly come to decision theory. Generally speaking, computational power of a decision theorist needs to be considered in advising him or her or it on how best to make decisions. It is safe to say that this was not even on Howard Raiffa's radar screen when Raiffa69 was written. For a survey article on the interface between computer science and decision theory, see [50]. A special issue of *Annals of Operations Research* on this topic is scheduled to appear some time in 2007 or 2008. If one thinks of decision making in economic or other social situations as computation in the formal sense of theoretical computer science, then one is led to some notion of bounded computational power as a formal expression of bounded rationality. One can then ask what computational power is required in order to make a decision. Some ideas along these lines are in the literature of computer science and game theory where the issue to to examine the computational complexity of finding best responses in games [3, 17, 18, 19, 28, 42]. A larger literature models bounded rationality in terms of finite automata, e.g., where the strategies of players are limited to those that are implementable by finite state automata [2, 36, 37, 38, 39, 40, 41, 43, 52]. Further work that studies strategies implementable by Turing machines may be found in [1] and [34]. In general, there is need to combine algorithmic and decision-theoretic approaches, taking into account constraints on the computational power of the agents. For some examples of work in this area, see [16, 22, 30].[4]

These modern challenges for decision makers call for the development of new ideas, new concepts, and new methods. In short, we need another Raiffa69.

[3] The author thanks Eric Friedman and Dave Pennock for these ideas.

[4] The author thanks Lance Fortnow and Richard McLean for some of these ideas.

References

[1] Anderlini, L., and Sabourian, H., "Cooperation and effective computability," *Econometrica*, **63** (1975), 1337-1369.

[2] Aumann, R., "Survey of repeated games," in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Zurich, 1981, 11-42.

[3] Ben-Porath, E., "The complexity of computing a best response automaton in repeated games with mixed strategies," *Games Econ. Behavior*, **2** (1990), 1-12.

[4] Collins, J.P., *The Development of a Solid Waste Environmental Evaluation Procedure*, Doctoral Dissertation, Department of Civil Engineering, University of Michigan, Ann Arbor, MI, 1974.

[5] Debreu, G., "Topological methods in cardinal utility theory," in K.J. Arrow, S. Karlin, and P. Suppes (eds.), *Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, CA, 1960, 16-26.

[6] de Neufville, R., and Keeney, R.L., "Use of decision analysis in airport development for Mexico City," in A.W. Drake, R.L. Keeney, and P.M. Morse (eds.), *Analysis of Public Systems*, MIT Press, Cambridge, MA, 1972.

[7] de Neufville, R., and Keeney, R.L., "Multiattribute preference analysis for transportation system evaluation," *Transportation Research*, **7** (1973), 63-76.

[8] Ellis, H.M., *the Application of Decision Analysis to the Problem of Choosing an Air Pollution Control System for New York City*, Doctoral Dissertation, Graduate School of Business Administration, Harvard University, Cambridge, MA, 1970.

[9] Farquhar, P.H., "A survey of multiattribute utility theory and applications," in M. Starr and M. Zeleny (eds.), *TIMS/North Holland Studies in the Management Sciences: Multiple Criteria Decisionmaking*, **6** (1977), 256-271.

[10] Fishburn, P.C., "Additivity in utility theory with denumerable produce sets," *Econometrica*, **34** (1964), 500-503.

[11] Fishburn, P.C., "Independence in utility theory with whole product sets," *Operations Research*, **13** (1965), 28-45.

[12] Fishburn, P.C. "Additive utilities with finite sets: Applications in the management sciences," *Naval Research Logistics Quarterly*, **14** (1967), 1-13.

[13] Fishburn, P.C., "Additive utilities with incomplete product sets: Application to priorities and assignments," *Operations Research*, **15** (1967), 537-542.

[14] Fishburn, P.C., "Interdependence and additivity in multivariate, unidimensional expected utility theory," *International Economic Review*, **8** (1967), 335-342.

[15] Fishburn, P.C., "Von Neumann-Morgenstern utility functions on two attributes," *Operations Research*, **22** (1974), 35-45.

[16] Fortnow, L., and Whang, D., "Optimality and domination in repeated games with bounded players," *Proceedings 26th Symposium on the Theory of Computing*, ACM, New York, 1994, 741-749.

[17] Gilboa, I., Kalai, E., and Zemel, E., "The complexity of eliminating dominated strategies," *Math. Oper. Res.*, **18** (1993), 553-563.

[18] Gilboa, I., and Samet, D., "Bounded versus unbounded rationality: The tyranny of the weak," *Games and Economic Behavior*, **1** (1989), 213-221.

[19] Gilboa, I., and Zemel, E., "Nash and correlated equilibria: Some complexity considerations," *Games and Econ. Behavior*, **1** (1989), 80-93.

[20] Huber, G.P., "Multi-attribute utility models: A review of field and field-like studies," *Management Science*, **20** (1974), 1393-1402.

[21] Indyk, P., "Dimensionality reduction techniques for proximity problems," *Proc. 9th SODA*, 2000, 371-378.

[22] Kalai, E., "Bounded rationality and strategic complexity in repeated games," in T. Ichiishi, A. Neyman, and Y. Tauman (eds.), *Game Theory and Applications*, Academic Press, New York, 1990, 131-157.

[23] Kambill, A., and van Heck, E., *Making Markets: How Firms can Design and Profit from Online Auctions and Exchanges*, Harvard Business School Press, Boston, MA, 2002.

[24] Keeney, R.L., "An illustrated procedure for assessing multiattributed utility functions," *Sloan Management Review* (1972), 38-50.

[25] Keeney, R.L., "A decision analysis with multiple objectives: The Mexico City airport," *Bell J. Econ. Management Sci.*, **4** (1973), 101-117.

[26] Keeney, R.L., "A utility function for response times of engines and ladders to fires," *Urban Analysis*, **1** (1973), 209-222.

[27] Keeney, R.L., and Raiffa, H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York, 1976.

[28] Knuth, D.C., Papadimitriou, C., and Tsitsiklis, N., "A note on strategy elimination in bimatrix games," *Oper. Res. Letters*, **7** (1988), 103-108.

[29] Krantz, D.H., Luce, R.D., Suppes, P., and Tversky, A., *Foundations of Measurement*, Vol. I, Academic Press, New York, 1971.

[30] Larson, K., and Sandholm, T., "Deliberation in equilibrium: Bargaining in computationally complex problems," *Proc. 17th National Conference on Artificial Intelligence*, AAAI Press, 2000, 48-55.

[31] Luce, R.D., Krantz, D.H., Suppes, P., and Tversky, A., *Foundations of Measurement*, Vol. III, Academic Press, New York, 1990.

[32] Luce, R.D., and Tukey, J.W., "Simultaneous conjoint measurement: A new type of fundamental measurement," *J. Math. Psychol.*, **1** (1964), 1-27.

[33] Mannheim, M.L., and Hall, F.L., "Abstract representation of goals," Paper P-67-24, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, MA, January 1968.

[34] Megiddo, N., and Wigderson, A., "On play by means of computing machines," in J. Halpern (ed.), *Reasoning about Knowledge*, Kaufman Publishers, Los Altos, 1986, 259-274.

[35] Miller, J.R., "Assessing alternate transportation systems," Memorandum RM-5865-DOT, RAND Corporation, Santa Monica, CA, 1969.

[36] Neyman, A., "Bounded complexity justifies cooperation in finitely repeated prisoner's dilemma," *Economics Letters*, **19** (1985), 227-229.

[37] Neyman, A., "Cooperation, repetition, and automata," in S. Hart and A. Mas-Colell (eds.), *Cooperation: Game Theoretic Approaches*, NATO ASI Series F, **155** (1997), 233-255.

[38] Neyman, A., "Finitely repeated games with finite automata," *Math. Op. Res.*, **23** (1999), 513-552.

[39] Neyman, A., and Okada, D., "Strategic entropy and complexity in repeated games," *Games and Econ. Behavior*, **29** (1999), 191-223.

[40] Neyman, A., and Okada, D., "Repeated games with bounded entropy," *Games and Econ. Behavior*, **30** (2000), 228-247.

[41] Neyman, A., and Okada, D., "Two person repeated games with finite automata," *Int. J. Game Theory*, **29** (2000), 309-325.

[42] Papadimitriou, C., "On players with a bounded number of states," *Games and Econ. Behavior*, **4** (1992), 122-131.

[43] Papadimitriou, C., and Yannakakis, M., "On complexity as bounded rationality," *STOC-94*, 1994, 726-733.

[44] Pfanzagl, J., *Theory of Measurement*, Wiley, New York, 1968.

[45] Project RAND, "Preliminary design of an experimental world- circling spaceship," RAND Corporation, Santa Monica, CA, May 1946.

[46] Raiffa, H. *Decision Analysis*, Addison-Wesley, Reading, MA, 1968.

[47] Raiffa, H., "Preferences for multi-attributed alternatives," Memorandum RM-5868-DOT, RAND Corporation, Santa Monica, CA, 1969.

[48] Roberts, F.S. *Measurement Theory, with Applications to Utility, Decisionmaking, and the Social Sciences,* Addison-Wesley, Reading, MA, 1979.

[49] Roberts, F.S. "Challenges for discrete mathematics and theoretical computer science in the defense against bioterrorism," in C. Castillo-Chavez and H.T. Banks (eds.), *Mathematical and Modeling Approaches in Homeland Security*, SIAM Frontiers in Applied Mathematics Series, 2003, 1-34.

[50] Roberts, F.S., "Computer science and decision theory," *Annals of Operations Research*, to appear.

[51] Roweis, S.T., and Saul, L.K., "Nonlinear dimensionality reduction by locally linear embedding," *Science*, **290** (2000), 2323-2326.

[52] Rubinstein, A. "Finite automata play the repeated Prisoner's Dilemma," *J. Econ. Theory*, **39** (1986), 83-96.

[53] Stevens, S.S., "On the theory of scales of measurement," *Science*, **103** (1946), 677-680.

[54] Stevens, S.S., "Mathematics, measurement, and psychophysics," in S.S. Stevens (ed.), *Handbook of Experimental Psychology*, Wiley, New York, 1951, 1-49.

[55] Stevens, S.S., "Measurement, psychophysics, and utility," in C.W. Churchman and P. Ratoosh (eds.), *Measurement: Definitions and Theories*, Wiley, New York, 1959, 18-63.

[56] Suppes, P, Krantz, D.H., Luce, R.D., and Tversky, A., *Foundations of Measurement*, Vol. II, Academic Press, New York, 1989.

[57] Suppes, P., and Zinnes, J., "Basic measurement theory" in R.D. Luce, R.R. Bush, and E. Galanter (eds.), *Handbook of Mathematical Psychology*, Vol. 1, Wiley, New York, 1963, 1-76.

[58] Tenenbaum, J.B., De Silva, V., and Langford, J.C., "A global geometric framework for nonlinear dimensionality reduction," *Science*, **290** (2000), 2319-2323.

[59] Ting, H.M., "Aggregation of attributes for multiattributed utility assessment," Tech. Rep. 66, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, August 1971.

[60] Witten, I. H., and Frank, E., *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*, Morgan Kaufman, 1999.