

Computer Science and Decision Theory

Fred S. Roberts¹

DIMACS Center, Rutgers University, Piscataway, NJ 08854 USA

froberts@dimacs.rutgers.edu

Abstract

This paper reviews applications in computer science that decision theorists have addressed for years, discusses the requirements posed by these applications that place great strain on decision theory/social science methods, and explores applications in the social and decision sciences of newer decision-theoretic methods developed with computer science applications in mind. The paper deals with the relation between computer science and decision-theoretic methods of consensus, with the relation between computer science and game theory and decisions, and with “algorithmic decision theory.”

1 Introduction

Many applications in computer science involve issues and problems that decision theorists have addressed for years, issues of preference, utility, conflict and cooperation, allocation, incentives, consensus, social choice, and measurement. A similar phenomenon is apparent more generally at the interface between computer science and the social sciences. We have begun to see the use of methods developed by decision theorists/social scientists in a variety of computer science applications. The requirements posed by these computer science applications place great strain on the decision theory/social science methods because of the sheer size of the problems addressed, new contexts in which computational power of agents becomes an issue, limitations on information possessed by players, and the sequential nature of repeated applications. Hence, there is a great need to develop a new generation of methods to satisfy these computer science requirements. In turn, these new methods will provide powerful new tools for social scientists in general and decision theorists in particular. This paper describes these new “social-science-based” computer science methodologies and investigates their application to problems of computer science and of the social and decision sciences.

Several major research themes span the developing interface between computer science and the social/decision sciences. They include:

- **Computational Tractability/Intractability.** Limits on what can be efficiently computed are very important in social-science-based computer science applications and applications of computer science methods to the decision sciences/social sciences. There has been considerable work concerning the computational complexity of computing the “winner” or the “consensus” in a voting or social choice situation. This work has shown that many social choice functions are computationally intractable and that there are situations where it is computationally intractable to calculate the winner of an election. In game theory, values, optimal strategies, equilibria, and other solution concepts can be easy, hard, or even impossible to compute. In computer science applications, power indices such as the Shapley-Shubik, Banzhaf, and Coleman indices need to be calculated for huge games, yet in most cases it is computationally intractable to compute them. As we consider larger and larger games in modern economic applications, Nash equilibria, solution concepts in cooperative games, and other game-theoretic solution concepts applied in a variety of settings become harder and harder to compute.

¹The author gratefully acknowledges the support of the National Science Foundation under grants SES-0351165 and INT-0339067 to Rutgers University. Many parts of this paper were originally written in the process of preparing several grant proposals. Numerous people assisted in preparing these proposals and I have tried to acknowledge their contributions and, in many cases, even their wording, in what follows. I hope I have not missed anyone. Special thanks go to those who helped me prepare the relevant proposals by providing wording, references, and editing, in particular Julia Abrahams, Alexander Barg, Denis Bouyssou, Joan Feigenbaum, Lance Fortnow, Eric Friedman, Ron Harstad, Mel Janowitz, Paul Kantor, Brenda Latka, Jon Leland, Michael Littman, Miguel Lobo, David Madigan, Richard McLean, Sunju Park, Aleksandar Pekeć, David Pennock, Michel Regenwetter, Michael Rothkopf, Robert Schapire, Peng Sun, Alexis Tsoukiàs, Vijay Vazirani, Rakesh Vohra, and Rebecca Wright. In addition, Alexis Tsoukiàs gave me extensive comments and suggestions on the entire manuscript. Some of the ideas here were originally described in the report “Challenges for Theoretical Computer Science” [218].

We might ask: Can our understanding of computational complexity help us to build computational tools that will enlarge the set of social/decision science models that can be analyzed and to adapt social-scientific concepts to solve problems of computer science? We might also ask: How can we characterize situations where intractability is a good thing, as for example when we use it to protect privacy or make it difficult to manipulate the outcome of an election?

- **Limitations on Computational Power/Information**². Increasingly, economic decisions have to be made in situations where there are many “actors” each having partial information or where there are limits to computation that prevent actors from obtaining needed information. Limitations include: (1) informational asymmetries (e.g., when agents such as network users have private information that a network designer does not have but requires); (2) computational limits to extracting information that is available (perhaps in large data sets); and (3) cognitive limitations on the part of agents. In classical economics and game theory, participants are assumed to have unbounded computational power, and thus can be assumed to use optimal strategies, so long as such strategies exist. This is clearly an inappropriate assumption in many computer science applications and it is necessary to develop models that assume that participants have limited computational power or information, a concept that is called “bounded rationality.”
- **The Impact on New Methodologies of the Sheer Size of Computer Science Applications.** The sheer size of computer science applications requires new methods/protocols where efficiency is a central focus and where issues of scalability and robustness with respect to errors in data are prominent. Consensus methods developed in the context of voting and decision making have begun to have a variety of applications involving multiple, large databases. Many of the resulting consensus problems are NP-complete in their most general setting, a conclusion that is particularly significant in the context of large databases. This calls for approximate algorithms or heuristic methods. In applications of game theory to computer science problems, we encounter games in which there are huge numbers of players (such as recipients of a multicast or nodes in a communication network). These new kinds of games offer many new and challenging complications.
- **Security, Privacy, Cryptography.** The need for secure communication and privacy sparks the need for cryptographic methods in the analysis and design of voting procedures, games, and decision processes. Such issues arise in areas like electronic commerce and electronic voting. As voting systems become increasingly automated and votes are counted by computer or over the Internet, security and privacy become major issues. The technical issues arising include system security, system availability (designed reliability and redundancy), voter identification/authentication (procedures, approvals, digital signature technologies), and ballot integrity (encryption)³. In game-theoretic models for computational tasks in which many computers have to communicate and cooperate and some of the computers can malfunction, resulting in conflict, we need to be concerned about the extent to which we can carry out computations in this kind of distributed setting in a reliable and secure way. Even if “bad” (failing) players are computationally restricted and so cannot decipher encrypted messages sent by “good” players, “satisfactory computational protocols need to be designed⁴.”
- **Implementation of Auctions**⁵. While dramatic changes in availability of information technology have allowed us to expedite increasingly complex business transactions, for example complex auctions, the ability to implement such transactions often puts a strain on existing computational methods. For example, complex dynamic pricing/auction procedures sometimes put a computational burden on the price-takers/bidders. This opens a whole new set of issues from defining fairness to designing of algorithms that have prespecified levels of complexity/transparency for both bid-takers and bidders. For example, if bidders with more computational power should have no advantage over bidders with less computational power, then we need to design auction procedures that are computationally tractable for winner determination, are transparent, and require minimal computational power to determine minimal bid increases to transform a losing bid into a winning one.

²Thanks to Richard McLean and Ron Harstad for some of the ideas in this paragraph. Others were originally contained in the report [218].

³See [74]. The author thanks Michael Grigoriadis for pointing him to this reference.

⁴See [262] for an expansion of these ideas.

⁵Thanks to Ron Harstad, Aleksandar Pekeć, and Mike Rothkopf for many of the ideas and some of the words on this item.

- **Markets as Information Sources.** For decades, economists have studied an astonishing “side effect” of financial and wagering markets: their ability to serve as highly accurate forecasting devices.⁶ Increasingly, individuals and firms are using markets to learn about their competitors’ preferences and to aggregate and combine information, for example through the generation of prices. These developments can benefit from methods of distributed computing and combinatorial exchange.
- **Dynamic Markets.**⁷ Airlines have routinely used current and historical demand to update their prices and their allocation of seat inventory to different fare classes. This process involves extensive computation, including demand forecasting, modeling of customer behavior, and optimization. Similar problems arise in the pricing of other resource-intensive, time-definite services, such as allocation of bandwidth in telecommunications. Major challenges here lie not in code but in the need for new algorithmic ideas that will allow the use of more sophisticated models of buyer behavior. Ever-more-powerful computers make it increasingly possible to develop sophisticated algorithms that allow actors to adjust prices, bandwidth, etc. in a dynamic way, as well as adjust schedules and allocations in response to changing market conditions or unforeseen events.
- **Computing Utilities**⁸. Utility functions, which have been studied for a long time by economists and psychologists, are increasingly important in considerations of interest to computer scientists. For both groups, there is the need to compute utilities in larger and larger environments and under conditions of distributed information. This is the case for example in computing tasks with massive input data sets as well as tasks in which data corresponds to agent valuations that have to be elicited (such as pricing data like the willingness to buy/pay at a given price). These newer applications call for new models, tools, and algorithms.
- **Learning in Multi-Agent Systems**⁹ Support for decision making is increasingly 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.
- **Learning through Repetition.** Issues of learning, taking advantage of repeated games or auctions or through combining repeated inputs, arise in computer science applications in fascinating ways. As noted above, some tasks may be delegated to software agents and one needs to understand how they should be designed to “learn” user preferences and how to bid and trade. A learning algorithm operates on a set of instances to produce a classifier or classification rule. Voting methods can help to achieve high accuracy in learning algorithms by voting on the predictions of several classifiers. Issues of learning arise relative to the Internet because we can think of agents playing repeated games in a distributed environment where they have very limited a priori information about the other players and the payoff matrix.
- **On-Line Decisions** Decisions increasingly need to be made “on-line,” that is, at the time data becomes available, rather than waiting for the entire input. In both economic and military contexts, large efforts have already been devoted to the development of highly interactive modeling tools to allow decision makers to make on-line decisions. Since much planning is often of a “deliberate” nature, well in advance of a financial development or emergency, deliberate plans need to be changed in an on-line manner in a sequence of decision points. Methods to improve on-line decision making are badly needed, as, more generally, are methods for sequential decision making.

⁶The author thanks Dave Pennock for this observation and pointing him to the literature of “information markets” (see [144, 183, 207, 426, 437]).

⁷The author thanks Jim Dana and Brenda Dietrich for many of the ideas and words about dynamic markets. For more information about this topic, see [275, 276, 283, 405, 406, 438]

⁸Thanks to Aleksandar Pekeč for some of the ideas and words in this bullet.

⁹The author thanks Eric Friedman and Dave Pennock for ideas in this and the next bullet.

- **Decision Making under Uncertainty or in the Presence of Partial Information**¹⁰ Many modern decision making problems have complicated stochastic components involving equipment failure, unforeseen political developments, incomplete situational understanding, etc. Classically, decisions about the design of complex systems have involved oversimplified independence assumptions about component failure. In the economic marketplace or the counter-terror planning venue, system failures can be the result of deliberate hostile actions and are thus highly non-random and non-independent. Uncertainty also enters into decisions since information about competitors is often incomplete or noisy. Input can include explicit mention of uncertainty in the form of probability distributions that need to be aggregated. New models of decision making that account for uncertainty are required, together with new algorithms for making diagnoses and inferences in the face of incomplete or noisy data.
- **Fusing Many Sources of Information** It is complicated enough for a planner or policy maker to fuse many different, conflicting pieces of information from a variety of sources, many unreliable, in advance planning. But in real-time decision making, especially in emergency situations, a huge variety of pieces of information of varying reliability must be efficiently fused in a short time in order to modify previously made plans and determine appropriate tactics. The situation is complicated by the subjective nature of much of the information, including subjective variables such as “readiness,” and even “reliability,” and subjective judgments of probabilities and beliefs from human observers, planners, and intelligence sources. Today, ubiquitous computing networks making use of millions of sensors can promote high speed and high quality decisions. However, we need to determine where to place sensors to have maximum impact on decisions and how to interpret alarms from sensor networks, and algorithmic methods are important for this purpose. Especially needed are ways of fusing uncertain information from subjective judgments of different sources and of combining these judgments with objective but disparate information from automatic devices such as sensors.

This paper is organized into three main sections, one dealing with the relation between computer science and decision-theoretic methods of consensus, a second with the relation between computer science and game theory and decisions, and a third on “algorithmic decision theory.” In some sense, this third section could cover the entire paper, since the first two sections also discuss algorithmic issues in decision making. We point to algorithmic issues in the first two topics that are returned to under the more general theme of algorithmic decision theory.

2 Connections between Computer Science Problems and Decision-theoretic Methods of Consensus

The notion of consensus arises in many decision making applications when we try to combine different opinions to reach one “consensus” opinion. The opinions can be stated in various ways, e.g., as “first choices” among a set of alternatives, as “rankings” (orders) on the alternatives, as scores, as classifications, etc. Here we describe connections between consensus problems and modern issues in computer science.

2.1 Meta-search, Collaborative Filtering¹¹

In computer science, “meta-search” is an example of a consensus problem, involving combining the results of several search engines (see [131]). Cohen, Schapire, and Singer [89] showed how meta-search can be formulated as a problem of consensus. They studied the problem of constructing a new ordering based on a learned probabilistic classifier or regression model, showed that the problem of finding the ordering that best agrees with a learned preference function is NP-complete, and described a simple greedy algorithm guaranteed to find a good approximation. Similar ideas of consensus have been used to address the reduction of spam, through the combination of results from multiple search engines. Dwork, et al. [127, 128] think of a ranking function as spammed by a page in a database of pages if in response to a given query it ranks the page too highly. Dwork, et al. showed that the consensus approach involving Kemeny medians [232, 233, 356, 359],

¹⁰The author thanks Ron Harstad for some of the ideas in this paragraph.

¹¹This section has benefited from the ideas and, in some cases, the words, of Cynthia Dwork and Dave Pennock. Andrew Gelman and D. Sivakumar are also responsible for some of the ideas.

widely used in social science applications, has excellent “spam-resistant” properties. Unfortunately, it is well known that the computation of the Kemeny median is NP-complete [24, 420]. Dwork, et al. developed an approximation to the Kemeny median that preserves most of its desirable properties and is computationally tractable.

The connection between ranking and web search has spawned an increasingly large literature, for example in using personal preference data to improve search outcomes [215], using page importance to guide web crawls [1], or connecting ranking/web search to graph-theoretical optimization problems such as finding the minimum feedback arc set in tournaments [5]. Connections between meta-search and learning are another avenue connecting computer science and the decision sciences/social sciences. See for example [149, 217, 253]. Analogous problems arise in information retrieval, when we try to rank documents according to their probability of relevance to a query.

Consensus methods also arise in collaborative filtering, where we use knowledge about the behavior of multiple users to make recommendations to another user, e.g., combining book or movie ratings to prepare an ordered list of books or movies to recommend [149, 340, 342, 345]. Freund, et al. [149] gave a formal framework for this problem and described a “boosting algorithm” for combining preferences that is very much in the spirit of social choice methods for combining rankings to obtain a preference ranking. Pennock, Horvitz, and Giles [340] observed that under certain axioms, the functional form of a collaborative filtering function is restricted. This constrains the choice of algorithms to those that have this particular form. They showed that under one set of axioms, a “nearest neighbor approach” (used by Freund, et al. [149]) is the only possible approach. Under other axioms, they showed that a weighted average method (used in practice, e.g., in Breese, et al. [67]) is the only possible approach.

Since the publication of the paper [355] by Resnick and Varian, there has been a great deal of interest in the topic known as recommender systems. A recommender system asks you a series of questions about things you liked and didn’t like, compares your answers to others’, and seeks people with similar opinions to yours. Recommender systems improve as more people use them. Recommender systems can be based on collaborative filtering (similarity among decision makers) or content-base (similarities among objects for the same decision maker [19]). From these one moves to data mining and knowledge extraction (what decision makers or consumers have shown to “prefer”, “choose”) and from there to preference elicitation and learning in algorithmic decision theory – topics to be discussed later (see Section 2.2 and Section 4.6). Recommender systems can be seen as an evolution of Decision Support Systems.¹² Of recent interest are “trust-based” recommender systems, where the goal is to take account of declared trust between agents as well as their opinions or reviews. For an axiomatic approach to trust-based recommender systems, see [82].

In meta-search, collaborative filtering, recommender systems, and related applications, what is different from the traditional decision-theoretic problem is that the number of “voters” is small and the number of “alternatives” or “candidates” is large. Moreover, quite often the problem consists of assigning “alternatives” to pre-defined ordered or nominal categories, a problem that is different from the classical ranking problem studied in decision theory. All the above call for new methods and algorithms.

In Section 4.7, we talk about algorithmic decision theory issues involving preference aggregation. The consensus methods used in meta-search, collaborative filtering, recommender systems, etc. are clearly relevant to the issues discussed in that section.

2.2 Large Databases and Inference

Consensus methods also arise in applications involving large databases in anthropology, political science, sociology, and economics. Much of the data in these databases is in the form of sequences and we often seek sequences that occur frequently or are “centrally located” or form a “consensus pattern.” An example arises in social welfare data [245, 246], where an algorithm for finding approximate sequential consensus patterns, ones that appear frequently in a database, is discussed. A similar problem arises in molecular biology, when we seek to choose sequences that occur frequently or are “centrally located” in a database of molecular sequences. For an example of a paper on finding consensus sequences in molecular databases, see [287]. This paper is one example of the emerging field of “bioconsensus” that involves applications of social-science-based consensus methods to problems of biology, many arising from the modern information-theoretic analysis of bioinformatics. Day [102] has compiled a bibliography of many papers in this emerging area. “Bioconsensus”

¹²The author thanks Alexis Tsoukiàs for suggesting the inclusion of these ideas on recommender systems.

has been the subject of several conferences [433, 434]. Several recent books also describe this developing field [103, 213]. New methods in bioconsensus include consensus ideas motivated by the notion of list coloring in graph theory [274] which ties together biological and social-scientific consensus pattern finding. The work on bioconsensus is farther along, in general, than are consensus methods for dealing with social-scientific databases.

Similar problems arise in “homeland security” applications. For instance, some problems of interest to the intelligence community deal with large sets of text messages and attempts to classify them by content. Methods of “fusion” of the results of several different classification algorithms are playing an important “consensus” role in such applications (see [10, 28, 29, 205, 225, 243, 315]).

More generally, the problem of merging or combining databases has been studied within Artificial Intelligence, where decision theory has already introduced several major contributions (see [21, 60, 85, 124, 238, 260, 261, 350]). Similar voting methods can help to achieve high accuracy in learning algorithms by combining the predictions of several classifiers [341, 386]. Decision theoretic approaches have also been used in order to induce classification rules from (partially) inconsistent and/or incomplete data bases [333]. For other examples, see [400, 401, 402, 403]. How can we take advantage of new procedures for combination to build better learning algorithms? This is one major research challenge.¹³

Combining responses to multiple queries is a related issue. Existing database systems have limited ability to restrict the synthesis of critical information. Privileged information can often be extracted through careful analysis of responses to a set of queries, some seemingly innocuous. Merging information from several queries provides the possibility of compromising the privacy of data. This calls for sophisticated access controls to monitor queries so that a user cannot combine queries to extract more information than he or she is allowed to have. Work is needed to discover fundamental laws determining the potential for undesirable inference, for example using functional and multivalued dependencies. Expert systems can be used to automate the process of inferring sensitive data and, by way of response, noise can be used to confuse possible inferences. Another key direction of research involves role-based access control. For a sampling of relevant literature on these topics, see [34, 71, 97, 106, 165, 193, 194, 195, 291, 294, 295, 320, 324, 325, 384, 410, 411].

Another increasingly important database problem involves preferential queries to databases. Here, one might look, for example, for information about a flight from New York to Paris, but have preferences about the airline, the type of itinerary, the type of ticket, etc. One might also look to combine responses from multiple travel-related websites to preferential queries. For a survey on this topic, see [86]. Both this topic and that of information inference can be viewed as sequential decision making problems, where the next query and/or the access to information about the next query depends on the sequence of queries/responses to that point. We return to the sequential decision making problem in Section 4.1.¹⁴

2.3 Software and Hardware Measurement

A key issue in the theory of measurement is to find ways to combine scores obtained on different criteria. A well-known principle developed in the theory of measurement is that if “ordinal scales” are used, then it is *meaningless* to compare arithmetic mean scores in the sense that the conclusion can change under admissible transformations of scale [357, 358, 363]. Fenton and Pfleeger [135] applied this principle to software measurement, in particular measurement of “understandability,” and observed that it is meaningless to conclude that the arithmetic mean rating of one system is higher than the arithmetic mean rating of a second system if a 5-point scale of understandability is used.

We often approach decision making problems by using several different points of view. For example, in measurement of “quality” of software, we consider such factors as functionality, reliability, efficiency, usability, maintainability, and portability. Methods of multicriteria decision theory [231, 418] can aid in implementing tradeoffs among conflicting ratings and have been used in software measurement with interesting results, as shown for example in [36, 298, 299, 332]. In the case of software measurement, one approach [135] is to weight the criteria, rank alternative software products on each criterion, and then rate alternative A over alternative B if the sum of the weights of all criteria under which A is preferred to B is larger than some threshold. This method cries out for analysis using the theory of meaningfulness in measurement. An alternative procedure is to develop metrics for each of the factors. For instance, early attempts at measuring

¹³Thanks to Alexis Tsoukiàs for ideas and references in this paragraph.

¹⁴Thanks to Alexis Tsoukiàs for suggesting the relevance of the preferential query literature.

portability of software are due to McCall, et al. [280] and Boehm, et al. [38]. Axiomatic and algorithmic approaches to measuring portability, along the lines described for collaborative filtering (Sec. 2.1), also need to be developed.

A widely used procedure in hardware performance measurement is to score a computer system on different benchmarks, normalize the scores relative to the performance of one of the systems, and then average the normalized scores by some averaging procedure. Fleming and Wallace [140] showed that the system that scores highest using this method can depend upon the choice of system used in the normalization, and thus the conclusions are in some sense meaningless. This has led to analysis of the process of merging normalized scores in the papers [2, 360]. The notion of merging normalized scores arises also in the work of Pennock, Horvitz, and Giles [340] on collaborative filtering. They showed that one way to impose a certain “scale invariance” property is to normalize all of the users’ ratings to a common scale before applying the filter. Another way is to design the filter so it depends only on the relative ranking provided and not on the actual rating values. This is the idea behind the collaborative filtering algorithm of Freund, et al. [149]. The notion of merging normalized scores could also be very relevant to the topic of recommender systems discussed in Section 2.1.

2.4 Consensus Computing, Image Processing¹⁵

An old subject in sociology, psychology, and economics involves dynamic models for how individuals change their opinions over time until, hopefully, some consensus is reached. The problems are related to the preference aggregation problems discussed in Section 4.7. Models of the opinion change process range from simple Markov models [104, 148, 356] to dynamic models on graphs closely related to neural nets [121, 173, 349]. These models have begun to be applied in computer science applications in distributed computing [190, 263, 337, 338]. In such applications, the values of processors in a network are updated until all the processors have the same value.

One application of this idea is in noise removal in digital images (see for example [212]). To check if a pixel level represents noise, one compares it with neighboring pixel levels. If the values are beyond a certain threshold, one replaces the value of the given pixel with a mean or median of the values of its neighbors. Related models arise in “distributed consensus” [334] where non-faulty processors are required to reach agreement eventually. Good protocols for how this can be arranged, which include accepting values (“opinions”) of neighboring non-faulty processors through some sort of majority rule process, need to be developed. In [33], a protocol based on the parliamentary procedure known as “cloture” was shown to be very good in terms of a number of important criteria including polynomial computation and communication.

Much of the work on opinion-change models has centered around the question of identifying initial configurations of opinions or opinion reformulation rules that lead to ultimate consensus. It is important to develop good procedures for answering this question in models arising in computer science contexts. Similar methods show promise for modern issues in epidemiology, where we deal with large models involving social networks and the spread of disease (and “opinion” is replaced by having or not having the disease). Here, a key problem is to identify initial configurations of vaccinated individuals so that no matter how many individuals are initially infected with a disease, the ultimate “consensus” will have no one with the disease. Some of the work in this area is motivated by newly emerging diseases and some by diseases deliberately spread by bioterrorists. For some work along these lines, see [121, 122, 189]; see [365] for discussion of the role of such methods in defense against bioterrorism.

2.5 Computational Intractability of Consensus Functions

There has been considerable work concerning the computational complexity of computing the “winner” or the “consensus” in a voting or social choice situation. See for example [23, 24, 25]. This work has shown that many social choice functions are computationally intractable and that there are situations where it is computationally intractable to calculate the winner of an election. Much work is needed to find good algorithms for consensus methods, new efficient approximations in cases where the exact solution cannot be efficiently found, and tractable heuristic methods. (See [127, 128, 279] for some work along these lines.) When

¹⁵The author thanks Juan Garay for the ideas in this section about models of distributed consensus and Mel Janowitz for those about noise removal in digital images.

we have imperfect information about preferences of voters, we might still be able to compute the probability of a particular candidate winning an election. In [191], it is shown that this problem is polynomially solvable in certain situations but $\#P$ -hard in others. In some cases, it is even NP-hard to determine if a given candidate has any chance to be a winner.

Computational intractability can have some positive implications in voting situations. For example, we would like to design voting systems where it is computationally intractable to “manipulate” the outcome of an election by “insincere” voting, e.g., by adding voters, declaring voters ineligible, adding candidates, declaring candidates ineligible, or adjusting the order in which alternatives are considered in a multiple-step voting situation. Early work on these issues is contained in [23]. For more recent work, see for example [129]. Other situations in which we might want computations involved in voting situations to be intractable arise when we are trying to protect the privacy of voters. For early work on this topic, see for example [88]. In turn, new methods developed in these computer science applications should be useful in making it difficult to manipulate the outcomes of electronic voting systems, a topic to which we return in Section 2.8.

2.6 Axiomatic Approaches and Algorithms

The size of modern computer applications requires new algorithms, approximations, etc., since existing algorithms don’t scale well. An intriguing idea is that the axiomatic approach, a standard tool in consensus theory, could help here. The functional form of a consensus function can sometimes constrain the choice of algorithms. See for example [417] and the work of Pennock, Horvitz, and Giles [340] described in Section 2.1. Other relevant work can be found in [51, 52, 55, 56, 57, 139, 182, 284, 285, 361, 362, 414]. We return to this point in our discussion of algorithmic decision theory, specifically in Section 4.1. Axiomatic characterizations of consensus functions can help us in identifying these functional forms without necessarily specifying the exact consensus function, and hence in turn help us to develop algorithms for calculating consensus. Totally new functional forms for consensus functions, derived with computer science applications in mind, can sometimes result from an axiomatic approach (see, e.g., [340] in a computer science framework and [3, 4] in more general contexts).

2.7 Order Relations and Revealed Preferences¹⁶

Establishing a consensus implies being able to compare objects either in order to establish that one is “before” the other or that they are “near.” Concepts developed or used in decision theory, e.g., non-symmetric similarities and special kinds of partial orders such as interval orders or semiorders [137, 347, 412], should be useful here and have already found application in computer science. Applications of ordered sets in computer science (see for instance [388]) include their uses as models for computation, their applications in knowledge representation, text categorization and data mining, and their uses in analyzing crypto-protocols in security and in inductive logic programming. (For examples of such applications, see [435].) Issues of special interest include the way in which preferences are revealed before a consensus is reached (see [439]). (For more on “preference elicitation,” see Section 4.6.) In traditional models of consensus, such preferences are thought of as various kinds of orders that are all revealed at the beginning (perfect information). However, increasingly in computer science applications, especially in situations of economic cooperation and competition using the Internet, we can think of software agents learning about other players’ preferences through repeated or iterated auctions or other procedures. (For more on repeated games, see Section 3.1.) By learning about a competitor’s preferences and, ultimately, its utility function, without seriously intending to win an auction, a firm can then seriously enter a later auction with a significant information advantage. Work of [150, 152, 178, 223, 393] is relevant. Also relevant is the use of repeated games to learn your own true valuation of an object such as in a sequential auction ([32]). Building models for consensus in this new context is another challenge that should be pursued. A related topic building on order relations concerns the issue of solving classic optimization problems when the costs are expressed in ordinal scales (not allowing an additive problem formulation), a problem receiving increasing attention in the planning and scheduling community of computer science (see [105, 272, 273, 344, 343, 364]). Still another topic, returned to in Section 4.4, involves the connection between preferences and artificial intelligence, through the attempt to understand the interplay between reasoning and rationality (see for example [116]).

¹⁶The author thanks Eric Friedman and Alexis Tsoukiàs for ideas used in this section.

2.8 Electronic Voting¹⁷

Electronic voting, for example remotely over the Internet, is being suggested in many contexts as a solution to many problems of modern voting. However, the many requirements it imposes with regard to correctness, anonymity, and availability pose an unusually thorny collection of problems. As voting systems become increasingly automated and votes are counted by computer or over the Internet, security, privacy, and robustness become major issues.

The security risks associated with electronic voting pose major technological challenges for computer scientists [26, 206, 211, 226, 316, 373]. The problems range from the threat of denial-of-service-attacks to the need for careful selection of techniques to enforce private and correct tallying of ballots. Cryptographic methods are needed to make sure that a voter's vote is kept private and that a consensus procedure is robust in the sense that we can be sure that the output is correct and not corrupted or modified. Work on secure multi-party communication and mix-networks is relevant here [30, 80, 81, 133, 155, 171, 209, 210, 330, 379].

Other possible requirements for electronic voting schemes are resistance to vote buying, defenses against malfunctioning software, viruses, audit ability, and the development of user-friendly and universally accessible interfaces. Many research questions arise. How can we prevent penetration attacks that involve the use of a delivery mechanism to transport a malicious payload to the target host (e.g., through a "Trojan horse" or remote control program)? What are the vulnerabilities of the communication path between the voting client (the devices where a voter votes) and the server (where votes are tallied)? Reliability issues are closely related and arise from dealing with random hardware and software failures (as opposed to deliberate, intelligent attack). A key difference between voting and electronic commerce is that in the former, one wants to irreversibly sever the link between the ballot and the voter. Audit trails are a possible way of ensuring this. Another important issue is to design methods for minimizing coercion and fraud, e.g., schemes to allow a voter to vote more than once and only having the last vote count. Further discussion of these topics has taken place at several recent workshops such as the DIMACS Workshop on Electronic Voting – Theory and Practice [436]. See the workshop report [282] for more details.

2.9 Combining Partial Information to Reach Decisions¹⁸

Combining partial information to reach decisions involves consensus methods. Yet, most known results about computational complexity of consensus functions require assumptions about perfect information of individual preferences, and few similar results are known when there is only partial, distributed, or unreliable information or limits to computation. (For some relevant results, however, see [22, 33, 328].) Procedures for merging or aggregating information have been widely studied by the decision science community, but without emphasis on the reliability of communication or the computational complexity of the protocol. The computational characteristics of the agents, especially when they include sensors, are extremely relevant, as is the structure of the communication network. As we study modern issues of merging and aggregating information, we will need to combine algorithmic and decision science approaches, simultaneously taking into account requirements for good group decisions, constraints on the computational power of the agents, and the structure of the communication network.

One major outcome of having only partial information is that there is resulting uncertainty. Decision making under uncertainty is a widely studied and important area of research. New algorithmic issues arising from decision making under uncertainty come from numerous modern applications, for example the recommender systems discussed in Section 2.1. Sometimes in these situations, the traditional tools of decision theory such as numerical (additive) utility functions or (subjective) probability distributions are less useful than algorithmic methods based on symbolic manipulation and "qualitative" methods as widely used in artificial intelligence. For more on the connections between decision theory and artificial intelligence, see Section 4.4. For more on qualitative decision theory, see for example [123, 124] and see the extended discussion and additional references in Section 4.4.

¹⁷The ideas in this section borrow heavily from the ideas and the words of Markus Jakobsson and Ari Juels.

¹⁸The author thanks Juan Garay and Alexis Tsoukiàs for some of the ideas in this section.

3 Connections between Computer Science and Game Theory

Game theory has a long history in the literature of economics, mathematics, and operations research, and in particular in decisionmaking applications. In recent years, computer scientists have found that game-theoretic concepts and methods are very relevant to their problems. Moreover, the increasingly complex games that arise in practical applications such as auctions and the Internet call out for new computer science-based methods. In this section, we review some aspects of the relation between computer science and game theory, with some emphasis on decision making applications.

3.1 Algorithmic Issues in Game Theory and Computer Science¹⁹

Classical work in game theory provides rich mathematical foundations and equilibrium concepts. However, to allow game theory concepts to scale up to large, complex systems in a tractable way, computational and representational insights are required. The rapidly emerging field of computational game theory is addressing such algorithmic issues. In game theory, values, optimal strategies, equilibria, and other solution concepts can be easy, hard, or even impossible to compute [179, 235, 236, 262]. We need to develop new methods for dealing with huge problems. For example, what are methods for computing the solution to a game with a huge number of players or in which the number of players changes constantly? For these huge problems, ordinary polynomial time complexity requirements might not be good enough.

Some specific topics of interest in this field include strategic conflict in matrix games and computation of equilibria, in particular of Nash equilibria in matrix games [228]. A recent result here is that finding Nash equilibria in matrix games with four or more players is complete for the search class PPAD [100]. Graphical models of multiplayer game theory are becoming important [222, 227, 416] and computation of Nash equilibria in graphical games provides an interesting challenge [264]. Other equilibrium concepts that provide computational challenges include correlated equilibria, cooperative equilibria, and evolutionary stable strategies [14, 145, 423]. Other topics of interest are learning in repeated games and regret minimizing algorithms [146, 153, 174] and games with state and their connections to reinforcement learning [50, 61, 229]. Finally, cryptographic methods are needed to make sure that in the course of a game, a player's views, utilities, preferences, strategies are kept private. See [109, 262] for examples of work on the interface between cryptography and game theory.

In computer science applications, power indices of game theory such as the Shapley-Shubik, Banzhaf, and Coleman indices need to be calculated for huge games, yet in most cases it is computationally intractable to compute them. Approximations and other effective heuristics are very much needed. Linial [262] has some interesting results about the game-theoretic concept of influence of coalitions. In some early work, Koller and Megiddo [235, 236] studied the computational complexity of finding max-min strategies for two-person zero-sum games. Grigoriadis and Khachiyan [179] found parallel and sequential randomized algorithms that compute a pair of ϵ -optimal strategies for a matrix game in expected sublinear time. Jain and Vazirani [208] studied approximation algorithms for strategyproof mechanisms in cooperative games arising in multicasting applications.

The theory of repeated games has a long history (see [154, 305]). One recent result [266] is a polynomial algorithm for finding Nash equilibria in repeated games. Of particular interest is the possibility of players adapting their strategy based on past experience. See [393] for work along these lines. Also of interest is the study of efficient algorithms for learning to play repeated games against computationally bounded adversaries [143, 150, 169]. For more on repeated games under bounded computational ability, see Section 3.5. Algorithmic issues that arise in repeated games are related to those arising in sequential decision making – see Section 4.1.

On the Internet, we have large-scale games with a large number of agents, asynchronous play, and an absence of full knowledge about the number of agents one is playing against or the beliefs they possess. The Internet is not the only institution to possess these features nor the first. Markets for traditional goods and services as well as travel networks all possess these features. Research issues involving such large-scale games include algorithmic issues and closely related foundational issues such as how to adapt and extend classical game theoretic models to deal with a large number of players and how to accommodate the absence

¹⁹Many of the ideas and many of the words in this section are due to Michael Kearns. The author also thanks Lance Fortnow and Rakesh Vohra for the material in the paragraph on large-scale games and Nati Linial for pointing him to relevant literature.

of common knowledge, common priors, asynchrony in play and distributed computation. Selected examples of relevant work involve anarchy models [370, 371, 372], robustness in games [224], and convergence issues [288].

3.2 Computational Issues in Auction Design²⁰

Recent advances in information technology and its rapid acceptance by the business community have allowed for expediting of complex business transactions. The most prominent example involves use of auctions in corporate procurement and in government deregulation efforts. Auctions are important and unusually complicated games. In auctions, bidding functions maximizing expected profit can be exceedingly difficult to compute and determining the winner can be easy or extremely hard [185, 368]. When the FCC prepared to accept bids for spectrum rights through auction in 1993, it became clear that existing bidding theory had little to say about such a complicated situation. The simultaneous progressive auction design did not allow bidders to place bids on combinations of licenses since the problem of picking the revenue-maximizing set of bids when bids on any combination are allowed is NP-complete. Yet, when many items with interrelated values are being sold, economic efficiency can be increased by allowing bids on combinations of items. Procedures for auctioning combinations of items have inherent computational problems to overcome. The emergence of these issues raises challenging research problems.

For example, such problems arise in combinatorial auctions in which multiple goods are auctioned and bidders have and wish to express different valuations on which goods complement each other and which goods substitute for each other. (For recent reviews see [99, 336, 419].) Allowing bidders to submit “all-or-nothing” bids for combinations of goods yields NP-complete allocation problems that need to be solved efficiently in designing an auction [368]. Furthermore, bidders face computational and communication problems in combinatorial auctions since they might not be feasibly able to express all possible preferences for all subsets of goods. Rothkopf, Pekeč, and Harstad [368] showed that determining the winner is not NP-complete for many economically interesting kinds of combinations. Park and Rothkopf [331] studied mechanisms that work well in practice even when the combinations lead to an NP-complete problem. Pekeč [335] demonstrated that cooperative game solution concepts such as Shapley value are useful in the design of second price combinatorial auctions. More systematic analysis of the usefulness of such game theory concepts is needed. The applications of methods of combinatorial optimization to combinatorial auctions are increasingly important [9, 156, 158, 368, 419]. Further work on combinatorial auctions with emphasis on partial information, limits to computation, and bidding incentives, is also needed.

Internet auctions are fast becoming an important method for transacting business. In a typical situation, a manufacturer uses an auction site to experiment with quantities supplied in order to learn about demand. While successful bidders might leave the pool of bidders, unsuccessful ones become more informed. These issues call for new models and analysis. See [188] for relevant work on how estimates of asset value and choices of bidding aggressiveness react to the estimates and choices of other bidders. One can study related issues of learning in repeated plays of a game. These arise also in the anticipated scenario under which software agents will act on behalf of humans in electronic marketplaces based on auctions and other forms of trading. Here, algorithmic problems of sequential decision making as discussed in Section 4.1 are very relevant.

In the analysis of auctions, traditional economic considerations of economic efficiency and revenue maximization must be balanced against concern for computational tractability and one should consider the tractability of methods for running the auction, determining winning bids, and informing bidders of minimal supplanting bids. In addition to computability, “transparency” of auction rules [186, 368] is critical to bidders. A framework for dealing with transparency must be developed.

Cryptographic methods are important in auction games too. For instance, we need to preserve the privacy of the participants while maintaining communication and computational efficiency [210, 308].

There are various kinds of transaction costs involved with auctions and auction theory that takes them into account is just beginning to be developed. For example, learning her value is not free to a bidder and could be a major undertaking. Some bidders may decline to incur the cost of preparing a bid since the expected profit from participating in the auction is less than the cost of bid preparation. Bid takers

²⁰This section depends heavily on the ideas and words of Ron Harstad, Aleksandar Pekeč, and Michael Rothkopf. Thanks also go to Markus Jakobsson for his ideas.

also incur costs, for example in preparing documentation on what is to be auctioned and in deciding whom to qualify to be bidders. They need to balance the value of getting additional bidders against the cost of qualifying them and to take into account the possible effect on the behavior of previously qualified bidders. Costs of publicizing auctions and gathering bidders can affect the timing of auctions and grouping of auctions of similar items.²¹ For recent work on transaction costs in auction design, see [187, 239, 247, 250, 353, 428].

3.3 Allocating/Sharing Costs and Revenues²²

Game-theoretic solutions such as the Shapley value and the nucleolus have long been used to allocate costs to different users in shared projects. Examples of such applications in the past include allocating runway fees to different users of an airport, highway fees to different size trucks, costs to different universities sharing library facilities, fair allocation of telephone calling charges among users, and compensation of agents who have to wait in a queue [268, 269, 277, 321, 389, 394, 404, 409, 429]. For cost-sharing problems with submodular costs, Moulin and Shenker [301, 302] characterized the family of cost-sharing mechanisms that are “budgetbalanced” and “group strategyproof” (in the sense that the dominant strategy for a coalition is to reveal the true value of their utility). Among these, the Shapley value minimizes welfare loss. If one requires “efficiency” rather than “budgetbalance,” then a “marginal cost” method is, in a very general setting, the only method satisfying group strategyproofness.

These ideas have found application in multicast transmissions. In unicast routing, each packet sent from a source is delivered to a single receiver. To send the same packet to multiple sites requires the source to send multiple copies of the packet and results in a large waste of bandwidth. In multicast routing, we use a directed tree connecting the source to all receivers and, at branch points, a packet is duplicated as necessary. The bandwidth used by a multicast transmission is not directly attributable to a single receiver and so one has to find a way to distribute the cost among various receivers. Building on the work of Moulin and Shenker, Feigenbaum, Papadimitriou and Shenker [134] studied the computational complexity of the Shapley value and the “marginal cost” cost-sharing mechanisms and showed that under a computational model that considers network complexity, the former does not have a feasible implementation while the latter does. Jain and Vazirani [208] gave a polynomial time computable cost-sharing mechanism that is budgetbalanced and group strategyproof, and moreover, in the multicasting application, computes the cost of the optimum multicast tree within a factor of two of the optimal. Other papers that deal with “noncooperative allocation problems” involving routing, load balancing, and other network issues include [317, 318, 319]. Computational issues involving cost-sharing mechanisms are closely related to computational models of network complexity and much more work needs to be done under different computational models. Other algorithmic issues in decision theory arise from the related notions of aggregation discussed in Section 4.7.

In cost sharing, notions of fairness arise and need to be studied. Specifically, issues of fair division or envy-free division of a commonly owned asset arise frequently in our new information-age economy; fast algorithms are needed for fair or envy-free sharing of revenues in e-commerce. In [159, 180], the authors described fast exact and approximation algorithms for equitable resource sharing and assignment of prices, using combinatorial and continuous optimization methods. In [65], an exponential time algorithm for envy-free division was given. For later approaches, see [132, 346, 367]. Some questions remain. Is the problem NP-hard? Also, the known algorithms are sequential. What about parallel algorithms here? All of us can reveal some of our preferences simultaneously so there is no need to process information in a sequential manner. Also related are cryptographic problems of fair exchange, arising for instance in contract signing [13, 160]. Special problems arise when allocation/sharing involves multi-dimensional outcomes, especially if the outcomes can interact [108, 177, 176]. Finally, related issues arise when demand for goods to be fairly allocated exceeds supply and we must institute “fair rationing.” Relevant papers on this topic deal with probabilistic rationing and fair queueing. See [107, 300, 303, 408].

²¹The ideas in this paragraph are taken from the announcement for the DIMACS workshop on Auctions with Transaction Costs, which was organized by Eric Rasmusen, Michael Rothkopf, and Tuomas Sandholm – see <http://dimacs.rutgers.edu/Workshops/Auctions/announcement.html>.

²²Many of the ideas in this section are due to Aleksandar Pekeć. The author also thanks Eric Friedman, Juan Garay, Michael Grigoriadis, Joan Feigenbaum and Vijay Vazirani for contributing ideas used in this section.

3.4 Game Theory, Computation, and Behavior²³

An exciting recent research direction has focused on the computational and behavioral aspects of game theory. As noted in Sections 3.1 and 3.3, computer scientists have been tackling computational issues arising in traditional game theory and economics and applying game-theoretic analyses to classical computer science problems such as resource allocation and networking. Of interest are improved algorithms for computing equilibria in large games (cf. Section 3.1), using sampling methods, using recent developments in mathematical programming, and finding sophisticated learning algorithms. Rich new representations for complex games now permit the specification of network structures of communication in games, and specialized but natural payoff functions [227, 228].

The computer science community has also been using game theory as a formalism and design principle for many kinds of distributed algorithms [134, 372]. Rather than the traditional view of each node in a distributed protocol or system “behaving” or being “malicious,” these analyses assume that nodes are simply “rational,” for some appropriate notion of self-interest. This approach has addressed problems such as routing, multicast transmission, and load balancing.

One central question of behavioral game theory is: What does it mean for an individual or organization to be rational? Behavioral game theory methods have been used to address this issue, challenging basic notions of rationality [75]. In Section 3.5, we consider a different problem: less than full rationality. Experimental work in behavioral economics has repeatedly shown that human subjects will frequently deviate from traditionally assumed notions of self-interest and greed. (See [186].) Moreover, they seem to do so in ways that are predictable, repeatable, and amenable to mathematical modeling and analysis. Concepts such as altruism, revenge, and envy have been quantified and fruitfully measured. (For some interesting relevant work, see [76, 96, 130, 397].)

One fascinating issue that arises when one tries to characterize rationality is the question of the relationship between rationality and optimization. (See [292] for a discussion.) It would clearly be beneficial for algorithmicists to model rationality as optimization of some appropriate function.

One of the primary difficulties of applications of computational game theory to resource allocation problems has been the specification of appropriate payoff or utility functions. Identifying such payoff or utility functions requires understanding of what the players in a game really care about. In classical game theory, some abstract “currency” is usually a means of measuring payoff and it is usually assumed that players act greedily with respect to such a currency. However, in situations that game theory is being applied to in computer science, where the players of a game are often nodes in a network, newer notions of currency need to be considered, things like negative network latency for routing, execution time for load balancing, or cash compensation.

The revised models of rationality proposed in the behavioral game theory literature present many interesting and challenging computational issues. Computer scientists now have a fairly clear understanding of the algorithmic challenges presented by Nash equilibria and related classical notions, but no deep computational thought has been given to how things change when mathematically formal notions of altruism and fairness are considered.

3.5 Bounded Rationality²⁴

Traditionally, economists and game theorists have assumed that strategic agents are fully rational in the sense that all players can completely reason about the consequences of their actions. However, recent research suggests that human players do not always behave in a way consistent with theoretical predictions. This has raised questions about the postulate of full rationality and led to the formalization of partially or boundedly rational players and games played by such players. (See [223, 376, 391].)

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 play a game in a way consistent with full rationality and also, if players are limited in their computational power, how different equilibrium outcomes can be from the fully rational case. These questions have led some to

²³This section borrows heavily from the ideas and words of Michael Kearns. The author also thanks Colin Camerer for ideas here.

²⁴This section is heavily dependent on the words and ideas of Lance Fortnow and Richard McLean.

examine the computational complexity of finding best responses in games [31, 169, 170, 168, 234, 326]. They have also led some to focus on repeated games played by various types of computing machines with an emphasis on their role in facilitating cooperative behavior. In one branch of this work, bounded rationality is interpreted as bounded recall where a player's strategic options are limited by constraints that are placed on memories of past actions [16, 257, 258]. 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 [15, 309, 310, 311, 312, 313, 314, 329, 374]. Further work that studies strategies implementable by Turing machines may be found in [8] and [286].

Most of the aforementioned work has been carried out by game theorists and, with the exception of a short burst of activity in the mid-1990's [142, 143, 151, 297, 329], there has not been a significant amount of activity in bounded rationality from the computer science point of view (for exceptions see [197, 249]). General research issues in this area include: What makes an appropriate model of bounded rationality? How do models of bounded rationality affect conclusions of standard models? What aspects of human behavior have no compelling model of bounded rationality to explain them? Are there computational models that properly estimate the computational power of bounded players while allowing for an analysis that yields useful results?

Participants in large, complex games arising in computer science applications need to develop optimal strategies in a context of bounded computational power or bounded rationality [375, 376]. 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 [143, 223, 248].) Complex, dynamic pricing/auction procedures sometimes put a computational burden on the price-takers/bidders and we need to design auction procedures in which players with limited computational power can do such things as determine minimal bid increases to transform losing bids to winning ones. (See for example [68].)

3.6 Market Clearing Algorithms²⁵

Markets are important mechanisms for allocating goods, services, tasks, and resources among multiple agents, be they human or software. The market-clearing problem [37] is that of deciding how to allocate the items among the agents. The last few years have witnessed a leap of improvement in market-clearing algorithms both for traditional market designs and entirely new market designs enabled by advanced clearing technology. There are computational implications of different market designs and a variety of corresponding algorithms for clearing markets optimally and approximately. Auctions [381] (one seller, multiple buyers), reverse auctions (one buyer, multiple sellers), and exchanges (multiple buyers and multiple sellers) are examples. Multi-item and multi-unit markets [383] are a major focus and computational implications of different classes of side constraints are of interest. Different bid types studied include price-quantity bids, different shapes of supply/demand curves, and package bids [382]. A recent paper by O'Neill, et al. [322] shows that one can always find equilibrium supporting prices in markets modeled with mixed integer programs if one is willing to pay not just for continuous variables (commodities) but also for nonconvexities such as start-up costs. Once the mixed integer program for the winner determination problem has been solved, the prices require only the solution of a linear program. This partially contradicts what has been thought to have been known about duality gaps and should open up a general discussion of richer kinds of pricing.²⁶

Among research challenges in this area are the development of methods for selective incremental preference elicitation for combinatorial markets [90] and of methods with which the market can be cleared optimally using only a small portion of the agents' preferences as input.

3.7 Streaming Data in Game Theory

In many applications, data must be analyzed as it arrives because the sheer quantity of it makes it infeasible to store in a central database or because of the urgency of the issues involved. In such applications, computations involving the data must be made during an initial scan as the data "stream" by. (See for example [20, 70, 91, 92, 93, 94, 95, 219, 259, 387, 415].) Herzog, Shenker, and Estrin [192] considered the problem of finding

²⁵Thanks go to Tuomas Sandholm for the words and ideas about market-clearing algorithms.

²⁶The author thanks Michael Rothkopf for pointing him to this result and its implications.

a “one-pass” mechanism to implement game-theory-based allocation schemes in multicasting. Similar “one-pass” issues arise in on-line auctions and it is necessary to develop and model bidding strategies in such situations, taking into account the effect of uncertainty, subjective judgments about the value of an item, learning in repeated auctions, etc. [290].

Analysis of data streams from network monitoring infrastructure is used to detect network faults in order to direct timely corrective action. Alarms are set off when scale values measuring disruptions or graphical patterns stray from the “normal.” How do we decide when to set off an alarm? When several measures are above threshold? When only one measure is above threshold but is very high? Measurement theory and decision theory seem relevant to these questions and analogous problems seem relevant to large games in which a central authority might need to interrupt play in case of unauthorized behavior.

4 Algorithmic Decision Theory

Today’s decision makers in fields ranging from engineering to medicine to economics to homeland security have available to them remarkable new technologies, huge amounts of information to help them in reaching good decisions, and the ability to share information at unprecedented speeds and quantities. These tools and resources should lead to better decisions. Yet, the tools bring with them daunting new problems and the need to apply them to problems arising in computer science presents complex new challenges. Among the issues are: The massive amounts of data available are often incomplete or unreliable or distributed and there is great uncertainty in them; interoperating/distributed decision makers and decision-making devices need to be coordinated; many sources of data need to be fused into a good decision. When faced with such issues, there are few highly efficient algorithms available to support decisions. There is a long tradition of algorithmic methods in logistics and planning. But algorithms to automate, speed up and improve real-time decision making are much less common. Algorithms for decision support, especially algorithms that can approximate good analytic solutions, are needed. Such algorithms should improve the ability of decision makers (human or automated) to perform in the face of the new opportunities and challenges that we describe in this section²⁷.

4.1 Sequential Decision Making Models and Algorithms²⁸

With the increasing amount of data faced by decision makers and the increasing speed with which their decisions need to be made, it is often the case that we must make decisions online before having access to all of the relevant data. Sequential decisions are also important in uncertain domains in which decisions impact the environment and therefore the context for future decisions. Sequential decision problems arise in many areas, including communication networks (testing connectivity, paging cellular customers, sequencing tasks, etc.), manufacturing (testing machines, fault diagnosis, routing customer service calls, etc.), artificial intelligence and computer science (optimal derivation strategies in knowledge bases, best-value satisficing search, coding decision tables, etc.), and medicine (diagnosing patients, sequencing treatments, etc.). An illustrative list of references for such applications includes [69, 72, 79, 84, 98, 126, 166, 220, 240, 241, 252, 307, 348, 354, 392, 396, 421, 427]. Sequential decision making is an old subject, but one that has become increasingly important with the need for new models and algorithms as the traditional methods for making decisions sequentially do not scale.

Making decisions that take into account risk attitudes is an important topic in sequential decision analysis. Many important dynamic decision models (in inventory control, revenue management, dynamic pricing in economics, event management in homeland security, etc.) are known to have nicely structured optimal policies under risk-neutral assumptions. Recently, Chen, et al. [83] showed that good structures in optimal policies still exist in risk-averse inventory control and other models. Work is needed to understand the

²⁷The author thanks Michael Littman, David Madigan, and Robert Schapire for ideas about algorithmic decision theory that underlie this introductory paragraph and, indeed, the entire section.

²⁸The author thanks several people for the ideas and many of the specific words in this section. The paragraph on sequential decision making under uncertainty in operations research/management science/economics is mostly due to Miguel Lobo, those on attitudes toward risk and on robust decision making mostly due to Peng Sun, and that on clinical trials to David Madigan. Alexis Tsoukiàs provided ideas and references about qualitative theories of sequential decision making. Endre Boros provided many of the examples of applications of sequential decision making in other disciplines. Madigan and Aleksandar Pekeč provided many of the overall ideas.

computational complexity of obtaining the optimal policies for risk-averse dynamic models and to design good algorithms for this purpose that incorporate the extra dimension of complexity brought in by risk considerations.

Recent research results indicate that, in constructing decision models based on data as it arrives, noisy data has profound effects on the optimal strategy and the value functions derived from dynamic decision models such as Markov decision processes [278]. It is important that we obtain *robust* policies—policies with performance guarantees when the nominal model constructed from data deviates from the unknown true model. While robustness is not a new idea, it remains a major challenge to obtain a robust policy for dynamic decision models. Traditional approaches relying on sensitivity analysis to understand the effects of variations in model assumptions are far from satisfactory in the context of dynamic models with complex structure. A challenge is to develop efficient algorithms or approximations to find robust policies.

There are a number of open problems and computational challenges involving sequential decision making under uncertainty in an operations research/management science/economic setting. In the context of pricing, or more broadly of revenue management, an important application is sequential pricing with uncertain demand. Most academic models in the field require the knowledge of a number of parameters about which, practitioners report, limited knowledge is usually available in practice. There is a need for the development of decision models that explicitly take this uncertainty into account, to introduce robustness and to model the learning process. Optimization-based control policies are a promising approach. One class of policies are heuristics where a measure of informativeness is maximized. Another are approximations of the, otherwise intractable, exact dynamic programming solution. In both cases, convex optimization methods for large-scale problems can be exploited for computational tractability [157, 369]. Recent work on sequential pricing [17, 73, 267] is relevant.

Another application area for sequential decision making involves online text filtering algorithms. These try to identify “interesting” documents from a stream of documents. As each document arrives, the algorithm has to decide whether or not to present the document to an oracle. If the algorithm presents a document to the oracle and if the document is indeed interesting, the oracle rewards the algorithm r “utiles,” or reward units. If the algorithm presents the document to the oracle and it is not interesting, the oracle penalizes the algorithm c utiles. The goal is to maximize expected utiles over either a fixed or infinite horizon. One can look at this as a sequential decision making problem. See for example [147, 221, 225, 432].

Also of interest are clinical trials problems. Conflicting ethical and performance considerations play significant roles in the planning and execution of human clinical trials. Ideally, every aspect of trial design should involve consideration of all foreseeable trial outcomes as well as utilities associated with those outcomes. Recent years have seen some progress towards this goal [77, 87, 398]. Increasingly in clinical trials, decisions as to whether or not to continue the trial are made at various steps. Interim data may provide grounds for terminating a study early if it does not appear to be leading to desired results or to making a new medication available more rapidly. The problem is to balance the financial and ethical considerations of ending a study early against the risk of an incorrect decision. Multiple criteria add considerable complications to analysis of these decision problems. The theory of multi-criteria decision making can help here [51, 58, 66, 418]. Axiomatic approaches to multi-criteria decision making (see [53, 54]) might help here in the spirit of connections between axiomatic approaches and algorithms discussed in Section 2.6. Also possibly of help are the group sequential methods that are designed to deal with the problems caused by repeated analyses of accumulating data and the further multiplicities that arise if there are several endpoints or if three or more treatments are under investigation [216].

A classical approach to sequential decision making involves Markov decision processes. Automated Markov decision processes are of interest, especially in AI applications. See [49] for a discussion of algorithmic approaches that make use of structure in the reward and value functions. In Markov decision processes, actions are stochastic and we measure the satisfaction of agents by numerical utility functions. However, this approach depends on being able to measure transition probabilities and to deal with utilities that are at least “additive.” In many applications, for example in artificial intelligence, uncertainty is often ordinal or qualitative, and similarly for utility. This observation has given rise to qualitative theories of sequential decision making that include logical languages for expressing preferences. Some references on this topic are [44, 125, 377, 378, 407].

4.2 Inspection Problems and Optimal Implementations for Boolean Decision Functions²⁹

An important area of application involving sequential decision making arises in inspection problems. For instance, inspections for weapons of mass destruction or other contraband entering through a port involves analysis of a stream of entities (containers) arriving at the port. A decision maker has to decide how to inspect them, which to subject to further inspection and which to allow to pass through with only minimal levels of inspection. This is a complex sequential decision making problem. One approach to this problem is based on the ideas of Stroud and Saeger [399] and has been applied at Los Alamos National Laboratory in initial formulations of the port-of-entry inspection problem as a problem in decision logic. Entities are thought of as having attributes, each in a number of states. In the simplest case, we seek to classify entities into two categories (“ok” and “suspicious”), there are two states (“present” or “absent”), and the classification can be described as involving a Boolean decision function (BDF). Different binary tree representations for a BDF have different inspection costs associated with them and one looks for an efficient decision tree representation. Modeling sensors used to test for attributes makes the problem more realistic and brings in possible inspection errors and false positives and negatives. Extending the problem to more than two categories and more than two states also makes the problem more realistic. While such problems are in general NP-complete, one can hope for efficient approximations or efficient solutions if one restricts the form of the BDFs under consideration. For a discussion of this problem, see [39, 366].

The problem of container inspection at ports of entry has given rise to a number of interesting approaches in recent years. For example, Boros, et al. [40] extended the work of Stroud and Saeger [399] and formulated a large-scale linear programming model yielding optimal strategies for container inspection. This model is based on a polyhedral description of all decision trees in the space of possible container inspection histories. The dimension of this space, while quite large, is an order of magnitude smaller than the number of decision trees. The Stroud-Saeger approach to port-of-entry inspection concentrated on limiting the number of decision trees we must consider by making special assumptions about the underlying Boolean function, i.e., that it is complete and monotone. An alternative approach is to make special assumptions about the topology of the decision tree. An example of this approach can be found in [430]. Anand, et al. [7] did an extensive sensitivity analysis on the results of Stroud and Saeger and found that even with changes in parameters, remarkably few binary decision trees (BDTs) turned out to be optimal. However, they noted that the “combinatorial explosion” of number of possible BDTs, even with restrictions to trees arising from complete, monotone Boolean functions as in Stroud and Saeger, limited the possibility of successful modifications of extensive search algorithms to the case of more types of sensors. Madigan, et al. [270, 271] generalized the notion of complete, monotone Boolean functions and defined notions of complete and monotone binary BDTs. They developed a search algorithm that identifies (locally) optimum (least cost) complete, monotone binary decision trees that is more efficient than the method developed by Stroud and Saeger and allows one to search by using a notion of neighborhood in the space of complete, monotone BDTs. Boros and Kantor [41] used game theory to study how to allocate an available budget for inspection between screening containers and unpacking them entirely. Their methods sought to maximize detection rate while deceiving opponents as to specific tests to which containers will be subjected.

There is a large literature arising in many fields that deals with the problem of finding an optimal binary decision tree (BDT) corresponding to a given Boolean function. The problem appears in AI as part of rule-based systems, where inductive trees are one of the main tools in rule-based learning [289, 306, 352]. In computer science, the problem is really central, since circuit complexity is essentially the depth of a BDT representing a function F [422]. In computer science, the main emphasis in the literature has been on worst case results, which is not necessarily useful for finding a particular representation of a special kind of given function F . Special cases, called “satisficing search” [166], and search strategies in theorem proving [240, 241], are also essentially equivalent to the optimal BDT problem. In reliability theory, many special cases of the problem fall under the rubric of the “diagnosis problem” [27, 35, 69, 72, 98, 214, 380]. It is sometimes called the “quiz show problem” [220] since quiz shows involve decisions about choosing a next category, or whether to stop or continue a process. In the theory of programming/databases, the conversion of decision tables into “if-then-else” structures also involves similar problems [354]. Although in general the problem of finding an optimal BDT is NP-complete [204], several classes of BDFs have been found for which

²⁹Thanks to Endre Boros for the ideas and many of the words in the paragraph on finding optimal binary decision trees.

an efficient solution is possible. This is the case for k -out-of- n systems [27, 163, 164, 181, 392], certain series-parallel systems [12, 43, 293, 380], read-once systems [11, 304], “regular systems” [42], and Horn systems [421].

4.3 Graphical Models for Decision Making³⁰

A growing theme in algorithmic decision theory is the need to find compact representations of beliefs and preferences. In recent years, the need to conserve memory in doing so has become more and more important. In this context, probabilistic graphical models have emerged as a crucial tool for decision making, computer-assisted diagnosis, and multivariate statistical analysis. These models use graphs to represent statistical models, with nodes corresponding to random variables and edges encoding conditional dependencies [251]. Of particular interest are current directions of research concerning decision analytic applications, specifically those involving “influence diagrams,” CP-nets, and GAI networks.

Influence diagrams include nodes corresponding to decisions and also nodes corresponding to utilities [196, 385], and they are general tools for representing and solving problems in Bayesian decision theory. They specify a decision problem on three levels: relational, functional, and numerical. The relational level uses a graph with nodes and edges to specify dependencies between the problem’s variables. The functional level describes the form of the dependencies, be they deterministic or stochastic. The numerical level provides the final details of the decision problem specification. To evaluate an influence diagram is to find an optimal policy and the optimal expected utility. Shachter [385] provided the original algorithm for direct evaluation of influence diagrams. Subsequent algorithmic developments include the decision networks of [431] and the LIMID (Limited Memory Influence Diagram) models of [252].

Most work to date on influence diagrams has focused on ways of manipulating the graph and the probability inputs so as to calculate the overall expected utility of any given decision strategy and, in particular, to identify the optimal strategy that maximizes this utility. More recent developments have focused on applications of influence diagrams to optimal sequential (or “multistage”) decision making [252] and on influence diagrams for causal modeling and inference [101]. In such applications, each decision can draw on the entire prior history of actions and observations. The LIMID algorithms [252], for example, were developed for sequential decision making applications where we only draw upon memory of the recent past (see Section 4.1), thus greatly improving computational performance and greatly extending the range of applications of influence diagrams. Of interest is the development and implementation of both exact and approximate LIMID algorithms. The topic of graphical models of games [222, 227, 416] is related and it is of interest to explore the extension of LIMIDs to multi-agent structures and games.

In many decision making applications, extracting and representing preferences is a key problem. We discuss this problem in Section 4.6. Increasingly, whether because of the need to repeat preference elicitation often as in repeated games (see Section 3.1) or because the decision maker is naive as in collaborative filtering or recommender systems (see Section 2.1), it is important to be able to elicit preferences at least somewhat automatically. Alternatively, one may have to concentrate on qualitative information rather than quantitative information. CP-nets [47] were introduced as a tool to represent preferences succinctly and to provide ways to make inferences about preferences. Several research challenges are central to the study of CP-nets. First, how can one tell whether one outcome in a CP-net “dominates” another. This is in general a difficult problem, indeed NP-complete under certain assumptions [47, 111]. Consistency testing is another challenge with CP-nets. Here, one asks whether or not there is a cycle of dominance, resulting in an outcome that is preferred to itself. In acyclic CP-nets, consistency is not necessarily the case, and testing for consistency can be as hard as testing for dominance [59, 111, 172].

GAI networks [175] are graphical tools designed as tools to aid in the elicitation of utilities. They keep track of dependencies between the different components of utilities in a multi-criteria context and the sequence of questions a decision maker is asked in the elicitation process, and make use of the “generalized additive independence” (GAI) assumption that goes back to Fishburn [136]. Because the GAI approach uses utilities, it can be more powerful than the CP-net approach, but it is more “expensive” since there is a need to elicit the utilities. Many questions about GAI networks remain for new research, including better ways to construct them and how to elicit utilities using them.

³⁰This section heavily uses the ideas and words of David Madigan. Thanks also to Alexis Tsoukiàs for providing ideas concerning CP-nets and GAI networks.

4.4 Decision-Theoretic Artificial Intelligence and Computer-based Reasoning Systems³¹

The complexity of real-world decision making problems has been reflected in recent years in new trends in research in artificial intelligence. All areas of traditional AI research have been enhanced by increasing use of mathematically grounded approaches to optimal decision making. Development of such formal approaches calls for significant contributions from the computer science and mathematics communities, including operations research and statistics.

The connections between computer science and artificial intelligence arise in the analysis of problems of representation, inference, and explanation in a decision-theoretic framework. There is considerable interest in computational foundations of intelligent sensing, with a particular focus on methods for grappling with incompleteness in representations and uncertainty about environments or situations. For instance, Bayesian and decision-theoretic principles are used to understand the construction and operation of intelligent, flexible, computer-based reasoning systems [201]. Theoretical issues arise in reasoning about beliefs and actions under computational constraints [197, 203]. Real-world applications of these ideas include systems for performing diagnosis and for prescribing ideal action for user interfaces [200]; other applications are to operating systems [198], information retrieval [18], machine repair, and communications.

Some of the more fascinating research questions involve “information triage” and alerting that takes human attention into consideration [200]. This involves a mix of computer science and behavioral science, and includes topics such as multitasking and the understanding of interruption and recovery. Other research problems involve methods for guiding computer actions in accordance with the preferences of people [198]. Optimization issues arise in trying to optimize the expected value of computational systems under limited and varying resources. Work along these lines includes ideal “metareasoning” for guiding computation [202], “continual computation” [199], and the construction of bounded-optimal reasoning systems—systems that maximize the expected utility of the people they serve, given the expected costs of reasoning, the problems encountered over time, and assertions about a system’s constitution [197]. Attacking these research challenges will require methods from learning theory and decision making.³²

In Section 2.9, we noted the need to extend decision theory from the formal approaches based on probability distributions and additive utility functions to the use of symbolic approaches not requiring the imposition of further hypotheses to quantify information. This is especially important in AI applications where the qualitative or symbolic is frequently more appropriate than the quantitative or numerical. This has given rise to what is known as “qualitative decision theory” and has developed a rather extensive literature. For some references, see [44, 45, 46, 62, 63, 64, 117, 118, 120, 125, 254, 255, 407, 424].

Another interesting direction of research involves “preferential entailment.” Doyle [112] and Shoham [390] observed that a reasoning system with only simple inferencing capabilities is not able to take into account preferences – a fundamental element of human capability to solve problems. Their suggestion was to enhance inference systems, namely the ones able to perform nonmonotonic reasoning, with a relation ordering the possible interpretations of a logical clause, in order to obtain “preferred” consequences instead of only “true” ones. Such an idea has been explored by several other researchers under different perspectives: [6, 60, 113, 114, 115, 161, 162, 242, 256]. Nevertheless, as Doyle and Wellman have shown [119] (see also [124]), the problem of aggregating such orders leads to an impossibility result as in Arrow’s impossibility theorem. In spite of these “negative results,” they lead to interesting research perspectives such as relaxing the axiomatic framework in order to find a solution (see for instance [425]). Nice “possibility results” have been obtained in the case of fusion of information [243, 244].

4.5 Efficient Instance-based Decision Making³³

In continuous-space control problems, the decision maker must take a series of actions based on a vector of continuous values representing the evolving environmental state. Because the state space itself may be

³¹The author thanks Michael Littman for the ideas and some of the words in this section and Alexis Tsoukiàs for the material and some of the words about qualitative decision theory and preferential entailment. This section, as well as several other sections, has benefited from the very nice survey [413] by Tsoukiàs

³²This paragraph borrows heavily from Eric Horvitz’ web page <http://research.microsoft.com/%7Ehorvitz/>.

³³This section depends heavily on the ideas and words of Michael Littman.

infinite, classical dynamic programming algorithms for decision making do not apply directly. In the machine learning community, researchers have been attracted to instance-based approaches (also known as memory-based, case-based, or episode-based approaches). In *instance-based decision making*, state spaces are constructed from examples of states experienced in interaction with the environment and relevant experience is “retrieved” by finding nearby examples to the current state [296]. Although a number of isolated successful implementations have been described, many significant algorithmic problems remain. The decision making process must be both *effective*—producing behavior that is appropriate to the problem at hand—and *efficient*—executing within available computational bounds.

Creating effective instance-based algorithms is essentially an AI problem; it involves selecting the relevant features of the problem and deciding on an optimization objective that is appropriate for important applications. Researchers have identified algorithms that are effective for challenging robotics problems like robot juggling [296] and navigation [395] as well as artificial driving [141]. Instance-based approaches have also been used in partially observable domains to select, decision-theoretically, actions to gain information [265, 281]. In instance-based learning and decision making, experience is broken down into a set of semi-independent “episodes.” The set of all experienced episodes is stored in a database and treated as an environmental “model” and the model is analyzed to determine reasonable courses of action from each of the experienced states. During a novel episode, the database is consulted to identify “similar” episodes, and this information is used to select an action appropriate to the current situation.

The instance-based approach requires effective algorithms. However, without efficient implementations of these algorithms, instance-based decision making will be limited to working in toy examples. Critical issues along these lines include studying nearest neighbor algorithms [78] to find ways to rapidly match current sensor readings to experience, as well as improved algorithms for similarity computations such as shape matching [323] to apply the instance-based approach to a variety of problem domains such as 3D vision.

There are a number of challenging research issues in the area of efficient instance-based decision making. For example, how can systems use multiple resolution representations in the context of an effective and efficient implementation? The motivation here is to find episodes that match the current episode in “broad strokes” and to use this information to guide additional matching or decision making. Another question is: How should an algorithm best make use of exploration to discover promising actions for the task at hand in an unfamiliar environment? Polynomial time exploration algorithms have been articulated for discrete state spaces [230]. How can these approaches be extended to continuous state spaces, perhaps exploiting geometric constraints? Finally, one can ask: What are appropriate clustering algorithms for efficiently analyzing large collections of experience to reduce computational demands for retrieval without sacrificing effectiveness of decision making? Attacking these problems will require close coordination between experts in machine learning and AI and those working on algorithm design and analysis.

4.6 Computational Approaches to Information Management in Decision Making: Representation and Elicitation³⁴

Computational issues play an increasingly important role in decision making that involves massive, multi-dimensional, diversely held and sometimes imprecise information. A successful approach to these emerging multidisciplinary issues related to decision making in complex environments will require combining state-of-the-art techniques and knowledge from theoretical computer science, decision theory and several social science disciplines such as economics and psychology.

Efficient representation and elicitation of information is a prerequisite to successful decision making. Gelman, et al. [167] make a similar point in a general discussion of decision making, in particular Bayesian decision making, and its drawbacks without well-defined inputs. The issues of preference representation and preference elicitation have been studied for years [136, 231], but the computational aspect of these problems is becoming a focal point because of the need to efficiently deal with massive and complex information. In Section 4.3, we briefly discussed the need to compactly represent preferences and discussed graphical models such as CP-nets and GAI networks designed to aid in preference elicitation. The problems are illustrated by combinatorial auctions, in which the decision maker has to elicit valuations, at least implicitly, from all agents,

³⁴This section depends heavily on the ideas and words of several people. Alexis Tsoukiàs provided most of the material on efficient representation and elicitation of information. Aleksandar Pekeč helped with the overall writeup and Michael Rothkopf provided helpful ideas.

not only for all single items that are being auctioned off, but also for all plausible combinations of the items, only to be faced with the allocation and pricing problem. (See also Section 3.2.) The problem of representing and eliciting valuations for all subsets of an n -element set [110, 138] goes well beyond combinatorial auctions and is of interest in decision making with multi-attribute valuations, in procedures for optimal bundling of goods and services, etc. Even under rather simplified assumptions about relationships among these subset valuations, exponentially many queries are required. However, many applications point to open questions regarding special subinstances of this problem [138]. One challenge is to develop theoretical guidelines for recognizing applicable situations in which efficient representation and elicitation procedures exist. Furthermore, one should study efficient preference representation and elicitation procedures when the structure of objects to be evaluated is more complex than, for example, subset structure. Such situations occur frequently in legal reasoning, dynamic and sequential decision making, automatic decision devices, collaborative filtering, etc. The topic has been addressed by researchers both in artificial intelligence and decision theory. Of relevance are approaches using argumentation theory [351], finding efficient algorithms for conditional preferences [48, 184], developing languages for computationally tractable preference models [44], and discovering ways to learn preferences from holistic examples.

4.7 Computational Approaches to Information Management in Decision Making: Aggregation³⁵

Once the information is efficiently elicited and presented, a central problem in decision analysis becomes that of aggregation; theoretical computer science can be useful in design of appropriate aggregation methods. In many applications, information is represented by multiple attributes of different choice alternatives. This is the case in the central decision analysis problem of aggregating estimates provided by a number of “experts” about a set of parameters (with the parameters often being probabilities). One approach to aggregating numerical estimates builds on the types of measurement theory concepts discussed in Section 2.3 in the context of software and hardware measurement. Since the mid-80s, the so-called “model-based” or “supra-Bayesian” approach has been developed for the problem. A second, related problem involves resolving incoherence in the estimates of different probabilities by the same expert. Research on this problem in recent years has been dominated by heuristic approaches.

In the model-based approach, the “supra-Bayesian decision maker” specifies beliefs about the accuracy of each expert and, given the estimates provided, obtains posteriors for the parameters by Bayesian updating. The models used are usually normal or beta/binomial, or some other conjugate distributions, selected mostly for computational expediency. However, we can easily find examples where these models lead to counter-intuitive results. The normal distribution’s fast-decaying tails make it an inadequate model for a human expert’s accuracy, but the use of other probability models presents computational challenges. A possible alternative is to use a double-sided exponential distribution as the error model. This allows for the computation of maximum a posteriori likelihood estimates with linear programming. LP can also be used to compute confidence intervals, but we are left with the problem of mapping intervals to confidence levels. This requires simulation of the posterior, and the development of efficient Markov-chain Monte Carlo (MCMC) methods for this problem. An interesting approach is to use the maximum-volume ellipsoid inscribed in a polyhedron (which can be computed by determinant maximization with second-order cone constraints) to specify a distribution that results in an efficient MCMC method.

Finally, it is important to design efficient decision making procedures and methods. In many situations, the process of preference representation and elicitation cannot be separated from the information aggregation and trade-off assessments and neither can be separated from the process of selecting the best alternative (making a decision) from the set of available ones. This interrelatedness is accentuated in situations where the information holders might not have incentives to truthfully reveal information, such as in most competitive environments, e.g., pricing and allocation decisions in auctions and markets, It is also accentuated in situations where full elicitation is computationally infeasible, say, when eliciting valuations for all subsets of an n element set for large n . In other cases, it is simply impossible because exact valuations are not available. This direction of research combines theoretical computer science and microeconomic mechanism design and relates to many interesting computational and algorithmic questions that are also important for

³⁵In this section, the first two paragraphs are for the most part originally due to Miguel Lobo. In addition, Aleksandar Pekeč, Mike Rothkopf, and Alexis Tsoukiàs all provided ideas and words on which this section is based.

both business (procurement auctions, business-to-business markets) and government (regulation of electricity markets, allocation of property rights such as those for use of frequencies in telecommunications). There is a lively body of research in this area [317, 318, 327, 368].

5 Concluding Remarks

In recent years, the interplay between computer science and biology has transformed major parts of biology into an information science and led to major scientific breakthroughs such as the sequencing of the human genome. Along the way, this has also led to significant new challenges for and developments in important areas of modern computer science, database search for example. The interplay between computer science and decision theory, and more generally between computer science and the social sciences, is not nearly as far along. Moreover, the problems that arise are spread over a myriad of different disciplines. However, the developments that have already taken place have already achieved a unique momentum of their own, and one can expect more and more exciting outcomes as the partnerships between computer scientists and social scientists expand and mature in the years ahead.

References

- [1] Abiteboul, S., Preda, M., and Cobena, G., “Adaptive on-line page importance computation,” *Proc. WWW2003*.
- [2] Aczél, J., “Determining merged relative scores,” *J. Math. Anal. & Appl.*, **150** (1990), 20-40.
- [3] Aczél, J., and Roberts, F.S., “On the possible merging functions,” *Math. Social Sci.*, **17** (1989), 205-243.
- [4] Aczél, J., Roberts, F.S., and Rosenbaum, Z., “On scientific laws without dimensional constants,” *J. Math. Anal. & Appl.*, **119** (1986), 389-416.
- [5] Ailon, N., Charikar, M., and Newman, A., “Aggregating inconsistent information: Ranking and clustering,” *STOC 2005*, 684-693.
- [6] Alchourron, C., Gärdenfors, P., and Makinson, D., “On the logic of theory change: Partial meet contraction and revision functions,” *Journal of Symbolic Logic*, **50** (1985) 510-530.
- [7] Anand, S., Madigan, D., Mammone, R., Pathak, S., and Roberts, F.S., “Experimental analysis of sequential decision making algorithms for port of entry inspection procedures,” in D. Mehrotra, D. Zeng, H. Chen, B. Thuraisingham, and F-X Wang (eds.), *Intelligence and Security Informatics, Proceedings of ISI-2006*, Lecture Notes in Computer Science, no. 3975, Springer-Verlag, New York, 2006.
- [8] Anderlini, L., and Sabourian, H., “Cooperation and effective computability,” *Econometrica*, **63** (1975), 1337-1369.
- [9] Andersson, A., Tenhungen, M., and Ygge, F., “Integer programming for combinatorial auction winner determination,” *Proc. 4th Int. Conference on Multi-agent Systems (ICMAS00)*, 2000, 39-46.
- [10] Angheliescu, A., Boros, E., Fradkin, D., Lewis, D., Menkov, V., Neu, D., Ng, K-B, and Kantor, P., “Prospective data fusion for batch filtering,” preprint, DIMACS Monitoring Message Streams Project, May 19, 2003.
- [11] Angluin, D., Hellerstein, L., and Karpinski, M., “Learning read-once formulas with queries,” *Journal of the Association for Computing Machinery*, **40** (1993), 185-210.
- [12] Arseneau, L., *Optimal Testing Strategies for s,t-Series Parallel Systems*, Master’s Thesis, Combinatorics and Optimization, University of Waterloo, 1996.
- [13] Asokan, N., Shoup, V., and Waidner, M., “Fair exchange of digital signatures,” *Eurocrypt’98*, 591-606.
- [14] Aumann, R.J., “Correlated equilibrium as an expression of Bayesian rationality,” *Econometrica*, **55** (1987), 1-18.
- [15] Aumann, R., “Survey of repeated games,” in *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Zurich, 1981, 11-42.
- [16] Aumann, R., and Sorin, S., “Cooperation and bounded recall,” *Games and Economic Behavior*, **1** (1989), 5-39.
- [17] Aviv, Y., and Pazcal, A., “Pricing of short life-cycle products through active learning,” preprint, Olin School of Business, Washington University, St. Louis, 2002.
- [18] Azari, D., Horvitz, E., Dumais, S., and Brill, E., “A decision making perspective on web question answering,” *Proceedings of the 19th Conference on Uncertainty and Artificial Intelligence*, Morgan Kaufmann, San Francisco, 2003.
- [19] Balabanovic, M., and Shoham, Y., “Fab: Content based, collaborative recommendation,” *Communications of the ACM*, **40** (1997), 66-72.

- [20] Balasubramanian, J.S., Garcia-Fernandez, J.O., Isacoff, D., Spafford, E.H., and Zamboni, D., "An architecture for intrusion detection using autonomous agents," *Proc. 14th IEEE Conf. on Computer Security Applications*, Dec 1998.
- [21] Baral, C., Kraus, S., Minker, J., and Subrahmanian, V.S., "Combining knowledge bases consisting of first order theories," *Computational Intelligence*, **8** (1992), 45-71.
- [22] Bar-Noy, A., Deng, X., Garay, J.A., and Kameda, T., "Optimal amortized distributed consensus," *Information and Computation*, **120** (1995), 93-100.
- [23] Bartholdi, J.J., Tovey, C.A., and Trick, M.A., "The computational difficulty of manipulating an election," *Social Choice and Welfare*, **6** (1989), 227-241.
- [24] Bartholdi, J.J., Tovey, C.A., and Trick, M.A., "Voting schemes for which it can be difficult to tell who won the election," *Social Choice and Welfare*, **6** (1989), 157-165.
- [25] Bartholdi, J.J., Tovey, C.A., and Trick, M.A., "How hard is it to control an election," *Mathematical and Computer Modelling*, **16** (1992), 27-40.
- [26] Benaloh, J., and Yung, M., "Distributing the power of a government to enhance the privacy of the voters," *ACM Symp. on Principles of Distributed Computing*, 1986, 52-62.
- [27] Ben-Dov, Y., "Optimal testing procedures for special structures of coherent systems," *Management Science*, **27** (1981), 1410-1420.
- [28] Benferhat, S., Dubois, D., Kaci, S., and Prade, H., "Bipolar representation and fusion of preferences in the possibilistic logic framework," in *Proceedings of the 8th International Conference on Knowledge Representation and Reasoning, KR'02*, Morgan Kaufman, San Francisco, 2002, 421-432.
- [29] Benferhat, S., Dubois, D., Kaci, S., and Prade, H., "Possibilistic merging and distance-based fusion of propositional information," *Annals of Mathematics and Artificial Intelligence*, **34** (2002), 217-252.
- [30] Ben-OR, M., Goldwasser, S., and Wigderson, A., "Completeness theorems for non-cryptographic fault-tolerant distributed computation," *Proc. 20th Symposium on Theory of Computing*, 1988, 1-10.
- [31] Ben-Porath, E., "The complexity of computing a best response automaton in repeated games with mixed strategies," *Games Econ. Behavior*, **2** (1990), 1-12.
- [32] Bergemann, D., and Valimaki, J., "Dynamic marginal contribution mechanism," Cowles Foundation Discussion Paper No. 1616, Cowles Foundation for Research in Economics, Yale University, July 2007.
- [33] Berman, P., and Garay, J.A., "Cloture votes: $n/4$ -resilient distributed consensus in $t + 1$ rounds," *Mathematical Systems Theory*, **26** (1993), 3-20.
- [34] Bertino, E., Catania, B., Ferrari, E., and Perlasca, P., "A logical framework for reasoning about access control models," *Proc. 6th ACM Symp. on Access Control Models and Technologies*, Chantilly, VA, 2001, 41-52.
- [35] Biasizzo, A., Zuzek, A., and Novak, F., "Enhanced sequential diagnosis," in J. W. Sheppard and W. R. Simpson (eds.), *Research Perspectives and Case Studies in System Test and Diagnosis*, Kluwer Academic Publishers, 1998.
- [36] Blin, M.J., and Tsoukiàs, A., "Multicriteria methodology contribution to software quality evaluations," *Software Quality Journal*, **9** (2001), 113-132.
- [37] Blum, A., Sandholm, T., and Zinkevich, M., "Online algorithms for market clearing," In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, San Francisco, CA, January 2002, 971-980.
- [38] Boehm, B.W., Brown, J.R., Kaspar, J.R., Lipow, M., McLeod, G., and Merritt, M., *Characteristics of Software Quality*, North Holland, Amsterdam, Holland, 1978.

- [39] Boros, E., Elsayed, E., Kantor, P., Roberts, F.S., and Xie, M., "Optimization problems for port-of-entry detection systems," in H. Chen and C.C. Yang (eds.), *Intelligence and Security Informatics: Techniques and Applications*, Springer, to appear.
- [40] Boros, E., Fedzhora, L., Kantor, P.B., Saeger, K., and Stroud, P., "Large scale LP model for finding optimal container inspection strategies," submitted to *Naval Research Logistics Quarterly*. (Preprint at http://rutcor.rutgers.edu/pub/rrr/reports2006/26_2006.pdf.)
- [41] Boros, E., and Kantor, P., "Deceptive detection methods for optimal security with inadequate budgets: The screening power index," preprint, Rutgers University, May 2007.
- [42] Boros, E., and Ünlüyurt, T., "Diagnosing double regular systems," *Annals of Mathematics and Artificial Intelligence*, **26** (1999), 171-191.
- [43] Boros, E., and Ünlüyurt, T., "Sequential testing of series-parallel systems of small depth," in M. Laguna and J. L. González Velarde (eds.), *OR Computing Tools for the New Millennium* (Proc. Conference of the INFORMS Computing Society, Cancun, Mexico, January 5-7, 2000), 2000, 39-74.
- [44] Boutilier, C., "Toward a logic for qualitative decision theory," *Proceedings of the 4th International Conference on Knowledge Representation and Reasoning, KR'94*, Morgan Kaufmann, San Francisco, 1994, 75-86.
- [45] Boutilier, C., "Knowledge representation for stochastic decision processes," in M.J Wooldridge and M. Veloso (eds.), *Artificial Intelligence Today. Recent Trends and Developments*, Springer Verlag, Berlin, 1999, 111-152.
- [46] Boutilier, C., "Decision making under uncertainty: Operations research meets AI (again)," in *Proceedings of the 17th National Conference on Artificial Intelligence, AAAI-2000*, AAAI Press, Menlo Park, 2000, 1145-1150.
- [47] Boutilier, C., Brafman, R.I., Doomshtak, C., Hoos, H., and Poole, D., "Preference-based constrained optimization with CP-nets," *Computational Intelligence*, **20** (2004), 137-157.
- [48] Boutilier, C., Brafman, R.I., and Hoos, H.H., and Poole, D., "Reasoning with condition Ceteris Paribus preference statements," *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, UAI'99*, Morgan Kaufmann, San Francisco, 1999, 71-80.
- [49] Boutilier, C., Dean, T., and Hanks, S., "Decision-theoretic planning: Structural assumptions and computational leverage," *Journal of Artificial Intelligence Research*, **11** 1999, 1-94.
- [50] Boutilier, C., Goldszmidt, M., and Sabata, B., "Continuous value function approximation for sequential bidding policies," *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, 1999.
- [51] Bouyssou, D., "Ranking methods based on valued preference relations: A characterization of the net flow method," *European Journal of Operational Research*, **60** (1992), 61-68.
- [52] Bouyssou, D., "Outranking relations: Do they have special properties?" *Journal of Multi-Criteria Decision Analysis*, **5** (1996), 99-111.
- [53] Bouyssou, D., Marchant, T., Pirlot, M., Perny, P., Tsoukiàs, A., and Vincke, Ph. *Evaluation and Decision Models: A Critical Perspective*, Kluwer Academic, Dordrecht, 2000.
- [54] Bouyssou, D., Marchant, T., Pirlot, M., Perny, P., Tsoukiàs, A., and Vincke, Ph. *Evaluation and Decision Models: Stepping Stones for the Analyst*, Springer Verlag, Berlin, 2006.
- [55] Bouyssou, D., and Perny, P., "Ranking methods for valued preference relations: A characterization of a method based on entering and leaving flows," *European Journal of Operational Research*, **61** (1992), 186-194.

- [56] Bouyssou, D., and Pirlot, M., "A characterization of strict concordance relations," in D. Bouyssou, E. Jacquet-Lagrèze, P. Perny, R. Słowiński, D. Vanderpooten, and Ph. Vincke (eds.), *Aiding Decisions with Multiple Criteria - Essays in Honor of Bernard Roy*, Kluwer Academic, Dordrecht, 2002, 121-146.
- [57] Bouyssou, D., and Pirlot, M., "Non-transitive decomposable conjoint measurement," *Journal of Mathematical Psychology*, **46** (2002), 677-703.
- [58] Bouyssou, D., and Vincke, P., "Ranking alternatives on the basis of preference relations: A progress report with special emphasis on outranking," *J. of Multi-criteria Decision Analysis*, **6** (1997), 77-85.
- [59] Brafman, R., and Dimopoulos, Y., "Extended semantics and optimization algorithms for CP-networks," *Computational Intelligence*, **20** (2004), 219-245.
- [60] Brafman, R.I., and Friedman, N., "On decision-theoretic foundations for defaults," *Artificial Intelligence*, **133** (2001), 1-33.
- [61] Brafman, R.I., and Tennenholz, M., "A near-optimal polynomial time algorithm for learning in certain classes of stochastic games," *Artificial Intelligence*, **121** (2000), 31-47.
- [62] Brafman, R.I., and Tennenholz, M., "On the foundations of qualitative decision theory," in *Proceedings of the 13th National Conference on Artificial Intelligence, AAAI96*, MIT Press, Cambridge, 1996, 1291-1296.
- [63] Brafman, R.I., and Tennenholz, M., "Modeling agents as qualitative decision makers," *Artificial Intelligence*, **94** (1997), 217-268.
- [64] Brafman, R.I., and Tennenholz, M., "An axiomatic treatment of three qualitative decision criteria," *Journal of the ACM*, **47** (2000), 452-482.
- [65] Brams, S.J., and Taylor, A.D., "An envy-free cake division protocol," *Amer. Math. Monthly*, **102** (1995), 9-18.
- [66] Brans, J.P., Vincke, P., and Mareschal, B., "How to select and how to rank projects: The PROMETHEE Method," *European J. of Operational Research*, **24** (1986), 228-238.
- [67] Breese, J., Heckerman, D., and Kadie, C., "Empirical analysis of predictive algorithms for collaborative filtering," *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Madison, WI, July, 1998.
- [68] Brewer, P.J., "Decentralized computation procurement and computational robustness in a smart market," *Economic Theory*, **13** (1999), 41-92.
- [69] Brule, J.D., Johnson, R.A., and Kletsky, E.J., "Diagnosis of equipment failures," *IRE Transactions on Reliability and Quality Control*, **RQC-9** (1960), 23-34.
- [70] Brunken, J.N., Mager, R., and Putzke, R.A., "NEMOS - the network management system for the AT&T long distance network," *Proc. 1989 IEEE Int'l. Conf. on Communications*, **3** (1989), 1193-1197.
- [71] Buczkowski, L.J., "Database inference controller," in C.E. Landwehr (ed.), *Database Security III: Status and Prospects, Results of IFIP WG 11.3 Workshop on Database Security*, North-Holland, 1990.
- [72] Butterworth, R., "Some reliability fault testing models," *Operations Research*, **20** (1972), 335-343.
- [73] Caldentey, R.A., and Bitran, G., "An overview of models for revenue management," *Manufacturing and Service Operations Management*, **5** (2002), 203-229.
- [74] California Department of State Internet Voting Task Force, *A Report on the Feasibility of Internet Voting*, Jan. 2000, <http://www.ss.ca.gov/executive/ivote/finalreport.pdf>.
- [75] Camerer, C.F., *Behavioral Game Theory*, Princeton University Press, 2003.

- [76] Camerer, C.F., Ho, T.H., and Chong, K., “Models of thinking, learning and teaching in games,” *American Econ. Review Papers and Proceedings*, **93** (2003), 192-195.
- [77] Carlin, B.P., Kadane, J.B., and Gelfand, A.E., “Approaches for optimal sequential decision analysis in clinical trials,” *Biometrics*, **54** (1998), 964-975.
- [78] Chakrabarti, A., Chazelle, B., Gum, B., and Lvov, A., “Searching for an approximate nearest neighbor on the Hamming cube: A lower bound,” *Proceedings of the 31st Annual ACM Symposium on the Theory of Computing*, 1999, 305-311.
- [79] Chang, C.L., and Slage, J.R., “An admissible and optimal algorithm for searching and-or graphs,” *Artificial Intelligence*, **2** (1971), 117-128.
- [80] Chaum, D., “Untraceable electronic mail, return addresses, and digital pseudonyms,” *Communications of the ACM*, **24** (1981), 84-88.
- [81] Chaum, D., Crépeau, C., and Damgård, I., “Multiparty unconditionally secure protocols,” *Proc. 20th Symposium on Theory of Computing*, 1988, 11-19.
- [82] Chayes, J., “Trust-based recommendation systems,” DIMACS Workshop on the Boundary between Economic Theory and Computer Science, DIMACS Center, Rutgers University, October 2007.
- [83] Chen, X., Sim, M., Simchi-Levi, D., and Sun, P., “Risk aversion in inventory management,” Fuqua School of Business, Duke University, faculty research paper FRPS05-200, October 2004.
- [84] Chiu, S.Y., Cox, L.A., and Sun, X., “Least-cost failure diagnosis in uncertain reliability systems,” *Reliability Engineering and System Safety*, **54** (1996), 203–216.
- [85] Cholvy, L., and Hunter, T., “Fusion in logic: A brief overview,” in *Proceedings of the 4th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU’97, LNCS 1244*, Springer Verlag, Berlin, 1997, 86-95.
- [86] Chomicki, J., “Preference formulas in relations queries,” *ACM Trans. on Database Systems*, **28** (2003), 1-39.
- [87] Christen, J.A., Mueller, P., Wathen, K., and Wolf, J., “A Bayesian randomized clinical trial: A decision theoretic sequential design,” preprint <http://odin.mdacc.tmc.edu/pm/pap/CMWW03.pdf>, 2003.
- [88] Cohen, J., and Fischer, M., “A robust and verifiable cryptographically secure election scheme,” *Proceedings of 26th Symposium on Foundations of Computer Science*, Portland, OR, October 1985, IEEE, New York, 1985, 372-382.
- [89] Cohen, W.W., Schapire, R.E., and Singer, Y., “Learning to order things,” *J. Artificial Intelligence Research*, **10** (1999), 243-270.
- [90] Conen, W. and Sandholm, T., “Minimal preference elicitation in combinatorial auctions,” *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Workshop on Economic Agents, Models, and Mechanisms*, Seattle, WA, August 6th, 2001.
- [91] Cormode, G., Korn, F., Muthukrishnan, S., and Srivastava, D., “Finding hierarchical heavy hitters in data streams,” *VLDB03* 2003, 464-475.
- [92] Cormode, G., and Muthukrishnan, S., “What’s hot and what’s not: Tracking most frequent items dynamically,” *PODS03*, 2003, 296-306, also in *ACM Trans. Database Systems*, **30** (2005), 249-278.
- [93] Cormode, G., and Muthukrishnan, S., “An improved data stream summary: The Count-Min sketch and its applications,” *LATIN04*, 2004, and *J. of Algorithms*, **55** (2005), 58-75.
- [94] Cormode, G., and Muthukrishnan, S., “What’s new: Finding significant differences in network data streams,” *INFOCOM04*, 2004.

- [95] Cormode, G., and Muthukrishnan, S., "Summarizing and mining skewed data streams," *SIAM Conf on Data Mining (SDM)*, 2005.
- [96] Costa Gomes, M., Crawford, V., and Broseta, B., "Experimental studies of strategic sophistication and cognition in normal-form games," *Econometrica*, **69** (2001), 1193-1235.
- [97] Cox, L.H., "Modeling and controlling user inference," in C.E. Landwehr (ed.), *Database Security: Status and Prospects*, North-Holland, 1988.
- [98] Cox, L.A., Jr., and Qiu, Y., "Optimal inspection and repair of renewable coherent systems with independent components and constant failure rates," *Naval Research Logistics*, **41** (1994), 771-788.
- [99] Cramton, P., Shoham, Y., and Steinberg, R. (eds.), *Combinatorial Auctions*, MIT Press, 2006, to appear.
- [100] Daskalakis, C., Goldberg, P.W., and Papadimitriou, C.H., "The complexity of computing a Nash equilibrium," Electronic Colloquium on Computational Complexity Report TR05-115, October 2005.
- [101] Dawid, A.P., "Influence diagrams for causal modelling and inference," *International Statistical Review*, **70** (2002), 161-189.
- [102] Day, W.H.E., "The sequence analysis and comparison bibliography," <http://edfu.lis.uiuc.edu/~class/sequence/>.
- [103] Day, W.H.E., and McMorris, F.R., *Axiomatic Consensus Theory in Group Choice and Biomathematics*, SIAM Publications, Philadelphia, PA, 2003.
- [104] Degroot, M.H., "Teaching a consensus," *J. Amer. Stat. Assoc.*, **69** (1974), 167-182.
- [105] Della Croce, F., Paschos, V.T., and Tsoukiàs, A., "An improved general procedure for lexicographic bottleneck problems," *Operations Research letters*, **24** (1999), 187-194.
- [106] Delugach, H., and Hinke, T.H., "Using conceptual graphs to represent database inference security analysis," *Jour. Computing and Info. Tech.*, **2** (1994), 291-307.
- [107] Demers, A., Keshav, S., and Shenker, S., "Analysis and simulation of a fair queueing algorithm," *Symposium Proceedings on Communications Architectures and Protocols*, ACM Press, 1989, 1-12.
- [108] De Waegenaere, A., and Wakker, P.P., "Nonmonotonic Choquet integrals," *Journal of Mathematical Economics*, **36** (2001), 45-60.
- [109] Dodis, Y., Halevi, S., and Rabin, T., "A cryptographic solution to a game theoretic problem," *Advances in Cryptology – Crypto2000 Proceedings*, Lecture Notes in Computer Science, **80** (2000), 112-130.
- [110] Doignon, J.-P., Pekeč, A., and Regenwetter, M., "The repeated insertion model for rankings: Missing link between two subset choice models," *Psychometrika*, **69** (2004), 33-54.
- [111] "CP-nets: Reasoning and consistency testing," *Proceedings of KR-2002*, Morgan Kaufmann, San Francisco, 2002, 121-132.
- [112] Doyle, J., "Reasoned assumptions and Pareto optimality," in *Proceedings of the 9th International Joint Conference on Artificial Intelligence, IJCAI85*, Morgan Kaufmann, San Francisco, 1985, 87-90.
- [113] Doyle, J., "Constructive belief and rational representation," *Computational Intelligence* **5** (1989), 1-11.
- [114] Doyle, J., "Rationality and its roles in reasoning," in *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI'90)*, MIT Press, Cambridge, 1990, 1093-1100.
- [115] Doyle, J., "Reasoned assumptions and rational psychology," *Fundamenta Informaticae*, **20** (1994), 35-73.
- [116] Doyle, J., "Prospects for preferences," *Computational Intelligence*, **20** (2004), 111-136, 2004.

- [117] Doyle, J., Shoham, Y., and Wellman, M.P., “A logic of relative desire,” in *Methodologies for Intelligent Systems, 6th International Symposium, ISMIS 91*, Springer-Verlag, Berlin, 1991, 16-31.
- [118] Doyle, J., and Thomason, R.H., “Background to qualitative decision theory,” *AI Magazine*, **20** (1999), 55-68.
- [119] Doyle, J., and Wellman, M.P., “Impediments to universal preference-based default theories,” *Artificial Intelligence*, **49** (1991), 97-128.
- [120] Doyle, J., and Wellman, M.P., “Representing preferences as ceteris paribus comparatives,” in *Decision-Theoretic Planning: Papers from the 1994 Spring AAAI Symposium*, AAAI Press, Menlo Park, California, 1994, 69-75.
- [121] Dreyer, P., *Applications and Variations of Domination in Graphs*, Ph.D. Thesis, Department of Mathematics, Rutgers University, Sept. 2000.
- [122] Dreyer, P.A. Jr., and Roberts, F.S., “Irreversible k -threshold processes: Graph-theoretical threshold models of the spread of disease and of opinion,” preprint, DIMACS Center, Rutgers University, September 2007.
- [123] Dubois, D., Fargier, H., Perny, P., and Prade, H. “On the limitation of ordinal approaches to decision making,” in *Proceedings of the 8th International Conference on Knowledge Representation and Reasoning, KR02*, Morgan Kauffman, San Francisco, 2002, 133-144.
- [124] Dubois, D., Fargier, H., Perny, P., and Prade H., “Qualitative decision theory: From Savage’s axioms to non-monotonic reasoning,” *Journal of the ACM*, **49** (2002), 455-495.
- [125] Dubois, D. and Prade, H., “Possibility theory as a basis for qualitative decision theory,” in *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI95*, Morgan Kaufmann, San Francisco, 1995, 1924-1930.
- [126] Duffuaa, S.O., and Raouf, A., “An optimal sequence in multicharacteristics inspection,” *Journal of Optimization Theory and Applications*, **67** (1990), 79-87.
- [127] Dwork, C., Kumar, R., Naor, M., and Sivakumar, D., “Rank aggregation, spam resistance, and social choice,” preprint, 2000, theory.stanford.edu/muri/reports/1999-2000/cynthia2.ps.
- [128] Dwork, C., Kumar, R., Naor, M., and Sivakumar, D., “Rank aggregation methods for the Web,” *Proc. WWW10*, 2001.
- [129] Elkind, E., and Lipmaa, H., “Hybrid voting protocols and hardness of manipulation,” in X. Deng and D-Z. Du (eds.), *The 16th Annual International Symposium on Algorithms and Computation, ISAAC 2005*, Lecture Notes in Computer Science, Springer-Verlag, Volume **3827**, 206-215.
- [130] Erev, I., and Barron, G., “On adaptation, maximization and reinforcement learning among cognitive strategies,” *Psychol Rev.*, **112** (2005), 912-931.
- [131] Etzioni, O., Hanks, S., Jiang, T., Karp, R.M., Madani, O., and Waarts, O., “Efficient information gathering on the internet,” *37th Annual Symposium on Foundations of Computer Science*, 1996.
- [132] Even, S., and Paz, A., “A note on cake cutting,” *Discr. Appl. Math.*, **7** (1984), 285-296.
- [133] Feigenbaum, J., Ishai, Y., Malkin, T., Nissim, K., Strauss, M.J., and Wright, R.N., “Secure multiparty computation of approximations,” *Proceedings 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, Lecture Notes in Computer Science, **2076** (2001), 927-938.
- [134] Feigenbaum, J., Papadimitriou, C.H., and Shenker, S., “Sharing the cost of multicast transmissions,” *J. of Computer and System Sciences*, **63** (2001), 21-41. (See also *Proceedings 32nd ACM Symposium on Theory of Computing*, 2000.)
- [135] Fenton, N.A., and Pfleeger, S.L., *Software Metrics*, 2nd Ed., PWS Publishing Co., Boston, MA, 1997.

- [136] Fishburn, P.C., *Utility Theory for Decision Making*, John Wiley and Sons, New York, 1970.
- [137] Fishburn, P.C., *Interval Orders and Interval Graphs*, Wiley-Interscience, New York, 1985.
- [138] Fishburn, P.C., Pekeč, A., and Reeds, J.A., "Subset comparisons for additive linear orders," *Mathematics of Operations Research*, **27** (2002), 227-243.
- [139] Fishburn, P.C., and Roberts, F.S., "Axioms for unique subjective probability on finite sets," *Jour. of Math. Psychol.*, **33** (1989), 117-130.
- [140] Fleming, P.J., and Wallace, J.J., "How not to lie with statistics: The correct way to summarize benchmark results," *Comm. of ACM*, **29** (1986), 218-221.
- [141] Forbes, J., and Andre, D., "Real-time reinforcement learning in continuous domains," *AAAI Spring Symposium on Real-time Autonomous Systems*, 2000.
- [142] Fortnow, L., and Kimmel, P., "Beating a finite automaton in the big match," *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, Morgan Kaufmann, San Francisco, 1998, 225-234.
- [143] 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.
- [144] Forsythe, R., Rietz, T.A., and Ross, T.W., "Wishes, expectations, and actions: A survey on price formation in election stock markets," *J. of Economic Behavior and Organization*, **39** (1999), 83-110.
- [145] Foster, D., and Vohra, R., "Calibrated learning and correlated equilibrium," *Games and Economic Behavior*, **21** (1997), 40-55.
- [146] Foster, D., and Vohra, R., "Regret in the on-line decision problem," *Games and Economic Behavior*, **29** (1999), 7-35.
- [147] Fradkin, D., and Littman, M., "Exploration approaches to adaptive filtering," DIMACS Technical Report 2005-01, DIMACS Center, Rutgers University, Piscataway, NJ, 2005.
- [148] French, J.R.P., Jr., "A formal theory of social power," *Psychol. Rev.*, **63** (1956), 181-194.
- [149] Freund, Y., Iyer, R., Schapire, R.E., and Singer, Y., "An efficient boosting algorithm for combining preferences," *J. of Machine Learning Research*, **4** (2003), 933-969.
- [150] Freund, Y., Kearns, M., and Mansour, Y., "Efficient algorithms for learning to play repeated games against computationally bounded adversaries," *Proceedings 36th Symposium on Foundations of Computer Science*, November 1995.
- [151] Freund, Y., Kearns, M., Mansour, Y., Ron, D., Rubinfeld, R., and Schapire, R., "Efficient algorithms for learning to play repeated games against computationally bounded adversaries," *Proceedings of the Thirty Seventh Annual Symposium on Foundations of Computer Science*, 1996, 332-341.
- [152] Friedman, E.J., and Shenker, S., "Learning and implementation on the Internet," Department of Economics, Rutgers University, Working Paper Series, 1998.
- [153] Fudenberg, D., and Levine, D., *The Theory of Learning in Games*, MIT Press, 1999.
- [154] Fudenberg, D., and Tirole, J., *Game Theory*, MIT Press, 1991.
- [155] Fujioka, A., Okamoto, T., and Ohta, K., "A practical secret voting scheme for large scale elections," *Auscrypt '92*, 244-251.
- [156] Fujishima, Y., Leyton-Brown, K., and Shoham, Y., "Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches," *Proc. of Int. Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, 1999, 548-553.

- [157] Gallego, G., and van Ryzin, G., “Optimal dynamic pricing of inventories with stochastic demand over finite horizons,” *Management Sci.*, **40** (1993), 999-1020.
- [158] Gallien, J., and Wein, L.M., “Design and analysis of a smart market for industrial procurement,” Working paper, Sloan School of Management, MIT, 2000.
- [159] Gallo, G., Grigoriadis, M.D., and Tarjan, R.E., “A fast parametric maximum flow algorithm and applications,” *SIAM J. Computing*, **18** (1989), 30-55.
- [160] Garay, J., Jakobsson, M., and MacKenzie, P., “Abuse-free optimistic contract signing,” *CRYPTO-99*, 449-466.
- [161] Gärdenfors, P., *Knowledge in Flux*, MIT Press, Cambridge, 1988.
- [162] Gärdenfors, P.G., and Makinson, D., “Nonmonotonic inference based on expectations,” *Artificial Intelligence* **65**, 197-245.
- [163] Garey, M.R., “Simple binary identification problems,” *IEEE Transactions on Computers*, **21** (1972), 588-590.
- [164] Garey, M.R., “Optimal task sequencing with precedence constraints,” *Discrete Mathematics*, **4** (1973), 37-56.
- [165] Garvey, T.D., and Lunt, T.F., “Cover stories for database security,” in C.E. Landwehr and S. Jajodia (eds.), *Database Security V: Status and Prospects, Results of the 5th IFIP WG 11.3 Working Group Conference on Database Security*, North-Holland, 1991, 363-380.
- [166] Geiger, D., and Barnett, J.A., “Optimal satisficing tree searches,” in *AAAI91*, Morgan Kaufmann, Anaheim, CA, 1991, 441-445.
- [167] Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B., *Bayesian Data Analysis*, second edition, Chapman and Hall, 2003.
- [168] Gilboa, I., Kalai, E., and Zemel, E., “The complexity of eliminating dominated strategies,” *Math. Oper. Res.*, **18** (1993), 553-563.
- [169] Gilboa, I., and Samet, D., “Bounded versus unbounded rationality: The tyranny of the weak,” *Games and Economic Behavior*, **1** (1989), 213-221.
- [170] Gilboa, I., and Zemel, E., “Nash and correlated equilibria: Some complexity considerations,” *Games and Econ. Behavior*, **1** (1989), 80-93.
- [171] Goldreich, O., Micali, S., and Wigderson, A., “How to play any mental game – A completeness theorem for protocols for honest majority,” *Proc. 19th Symposium on Theory of Computing*, 1987, 218-229.
- [172] Goldsmith, J., Lang, J., Truszczyński, M., and Wilson, N., “The computational complexity of dominance and consistency in CP-nets,” *IJCAI '05*.
- [173] Goles, E., and Martinez, S., *Statistical Physics, Automata Networks, and Dynamical Systems*, Kluwer Academic Publishers, Dordrecht, Netherlands, 1992.
- [174] Gordon, G., Greenwald, A., Marks, C., and Zinkevich, M., “No-regret learning in convex games,” Technical Report CS-07-10, Dept. of Computer Science, Brown University, October 2007.
- [175] Gonzalez, C., and Perny, P., “GAI networks for utility elicitation,” *Proceedings of KR-2004*, Morgan Kaufmann, San Francisco, 2004, 224-234.
- [176] Grabisch, M., and Labreuche, C., “Fuzzy measures and integrals in MCDA,” in J. Figueira, S. Greco, and M. Ehrgott (eds.) *Multiple Criteria Decision Analysis: State of the Art Surveys*, Springer Verlag, Boston, Dordrecht, London, 2005, 563-608.

- [177] Grabisch, M., Murofushi, T., Sugeno, M., and Kacprzyk, J., *Fuzzy Measures and Integrals. Theory and Applications*, Physica Verlag, Berlin, 2000.
- [178] Greenwald, A., Friedman, E.J., and Shenker, S., "Learning in network contexts: Experimental results from simulations," *Games and Economic Behavior*, **35** (2001), 80-123.
- [179] Grigoriadis, M.D., and Khachiyan, L.G., "A sublinear-time randomized approximation algorithm for matrix games," *Operations Research Letters*, **18** (1995), 53-58.
- [180] Grigoriadis, M.D., and Khachiyan, L.G., "Coordination complexity of parallel price-directive decomposition," *Math. Oper. Res.*, **21** (1996), 321-340.
- [181] Halpern, J.Y., "Fault testing for a k-out-of-n system," *Operations Research*, **22** (1974), 1267-1271.
- [182] Hansen, P., and Roberts, F.S., "An impossibility result in axiomatic location theory," *Math. of Operations Research*, **21** (1996), 195-208.
- [183] Hanson, R., "Combinatorial information market design," *Information Systems Frontiers*, **5** (2003), 105-119.
- [184] Hansson, S.O., "What is ceteris paribus preference?," *Journal of Philosophical Logic*, **425** (1996), 307-332.
- [185] Harstad, R.M., "Alternative common-value auction procedures: Revenue comparisons with free entry," *J. of Political Economy*, **98** (1990), 421-429.
- [186] Harstad, R.M., "Dominant strategy adoption, efficiency, and bidder's experience with pricing rules," *Experimental Economics*, **3** (1990), 261-280.
- [187] Harstad, R., "Endogenous competition alters the structure of optimal auctions," DIMACS Workshop on Auctions with Transaction Costs, DIMACS Center, Rutgers University, March 2007.
- [188] Harstad, R.M., Kagel, J.H., and Levin, D., "Equilibrium bid functions for auctions with an uncertain number of bidders," *Economics Letters*, **33** (1990), 35-40.
- [189] Hartke, S., *Graph-Theoretic Models of Spread and Competition*, Ph.D. Thesis, Department of Mathematics, Rutgers University, October 2004.
- [190] Hassin, Y., and Peleg, D., "Distributed probabilistic polling and applications to proportionate agreement," *Proc. 26th Int. Colloq. on Automata, Languages and Programming ICALP*, 1999, 402-411.
- [191] Hazon, N., Aumann, Y., Kraus, S., and Wooldridge, M., "Evaluation of election outcomes under uncertainty," DIMACS Workshop on the Boundary between Economic Theory and Computer Science, DIMACS Center, Rutgers University, October 2007.
- [192] Herzog, S., Shenker, S., and Estrin, D., "Sharing the "cost" of multicast trees: An axiomatic analysis," *IEEE/ACM Trans. on Networking*, **5** (1997), 847-860.
- [193] Hinke, T.H., "Inference aggregation detection in database management systems," *IEEE Symp. on Security and Privacy*, Oakland, CA, 1988.
- [194] Hinke, T.H., "Database inference engine design approach," in C.E. Landwehr (ed.), *Database Security II: Status and Prospects*, IFIP WG 11.3 Workshop on Database Security, North-Holland, 1988.
- [195] Hinke, T.H., Delugach, H. S., and Wolf, R.P., "Protecting databases from inference attacks," *Computers and Security*, **16** (1997), 687-708.
- [196] Howard, R.A. and Matheson, J.E., "Influence diagrams," in R.A. Howard and J.E. Matheson (eds.), *Readings in the Principles and Applications of Decision Analysis*, Strategic Decisions Group, Menlo Park, California, 1984.

- [197] Horvitz, E.J., "Reasoning about beliefs and actions under computational resource constraints," *Proceedings of the 1987 Workshop on Uncertainty in Artificial Intelligence*, 1987.
- [198] Horvitz, E.J., "Continual computation policies for utility-directed prefetching," *Proceedings of the 7th ACM Conference on Information and Knowledge Management*, ACM Press, New York, 1998, 175-184.
- [199] Horvitz, E.J., "Principles and applications of continual computation," *Artificial Intelligence Journal*, **126** (2001), 159-196.
- [200] Horvitz, E., and Barry, M., "Display of information for time-critical decision making," *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, 1995, 296-305.
- [201] Horvitz, E., Breese, J., and Henrion, M., "Decision theory in expert systems and artificial intelligence," *J. of Approximate Reasoning, Special Issue on Uncertainty in Artificial Intelligence*, **2** (1988), 247-302.
- [202] Horvitz, E., Cooper, G., and Heckerman, D., "Reflection and action under scarce resources: Theoretical principles and empirical study," *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, MI, Morgan Kaufmann, San Mateo, CA, August 1989, 1121-1127.
- [203] Horvitz, E., Kadie, C.M., Paek, T.E., and Hovel, D., "Models of attention in computing and communications: From principles to applications," *Communications of the ACM*, **46** (2003), 52-59.
- [204] Hyafil, L., and Rivest, R.L., "Constructing optimal binary decision trees is NP-complete," *Information Processing Letters*, **5** (1976), 15-17.
- [205] Ibraev, U., Ng, K-B., and Kantor, P., "Exploration of a geometric model of data fusion," *Proc. 2002 Conference of the American Society for Information Science and Technology*, 2002.
- [206] Internet Policy Institute, *Report of the National Workshop on Internet Voting: Issues and Research Agenda*, technical report, March 2001.
- [207] Jackwerth, J.C., and Rubinstein, M., "Recovering probability distributions from options prices," *J. of Finance*, **51** (1996), 1611-1631.
- [208] Jain, K., and Vazirani, V., "Applications of approximation algorithms to cooperative game theory," *Proceedings of Symposium on Theory of Computing*, 2000.
- [209] Jakobsson, M., "Flash mixing," *PODC '99*, 83-89.
- [210] Jakobsson, M., and Juels, A., "Mix and match: Secure function evaluation via ciphertexts," *Asiacrypt '00*, 162-177.
- [211] Jakobsson, M., Juels, A., and Rivest, R.L., "Making mix nets robust for electronic voting by randomized partial checking," *USENIX Security Symposium*, 2002, 339-353.
- [212] Janowitz, M., "Ordinal filtering in digital image processing," in E. Wegman and D. DePriest (eds.), *Statistical Image Processing and Graphics*, Marcel Dekker, 1986, 25-41.
- [213] Janowitz, M., LaPointe, F-J., McMorris, F.R., Mirkin, B., and Roberts, F.S. (eds.), *Bioconsensus*, DIMACS Series, Volume 61, American Mathematical Society, Providence, RI, 2003.
- [214] Jedrzejowicz, P. "Minimizing the average cost of testing coherent systems: Complexity and approximate algorithms," *IEEE Transactions On Reliability*, **R-32** (1983), 66-70.
- [215] Jeh, G., and Widom, J., "Scaling personalized Web search," Stanford University Technical Report, 2002.
- [216] Jennison, C., and Turnbull, B.W., *Group Sequential Methods with Applications to Clinical Trials*, Chapman and Hall, 1999.

- [217] Joachims, T., "Optimizing search engines using clickthrough data," *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*, ACM, 2002.
- [218] Johnson, D., Papadimitriou, C., Wigderson, A., and Yannakakis, M., "Challenges for theoretical computer science," report, June 2000, <http://www.research.att.com/~dsj/nsflist.html>.
- [219] Johnson, T., Cormode, G., Korn, F., Muthukrishnan, S., Spatscheck, O., and Srivastava, D., "Holistic UDAFs at streaming speeds," *SIGMOD04*, 2004.
- [220] Kadane, J.B., "Quiz show problems," *Journal of Mathematical Analysis and Applications*, **27** (1969), 609-623.
- [221] Kaelbling, L.P., Littman, M.L., and Cassandra, A.R., "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, **101** (1998), 99-134.
- [222] Kakade, S., Kearns, M., Langford, J., and Ortiz, L., "Correlated equilibria in graphical games," *Proceedings of the 4th ACM Conference on Electronic Commerce*, ACM Press, San Diego, 2003, 42-47.
- [223] 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.
- [224] Kalai, E., "Large robust games," *Econometrica*, **72** (2004), 1631-1665.
- [225] Kantor, P.B., and Roberts, F.S., "Monitoring message streams: Algorithmic methods for automatic processing of messages," *J. of Intelligence Community Research and Development*, Feb. 2, 2007; see <http://www.datalaundering.com/download/fredpoul.pdf>.
- [226] Kazue, S., and Killian, J., "Receipt-free mix-type voting scheme: A practical solution to the implementation of a voting booth," *Advances in Cryptology-EUROCRYPT'95*, Lecture Notes in Computer Science, **921** (1995), 393-403.
- [227] Kearns, M., Littman, M., and Singh, S., "Graphical models for game theory," *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2001, 253-260.
- [228] Kearns, M., and Mansour, Y., "Efficient Nash computation in large population games with bounded influence," *Proceedings of UAI 2002*, 259-266.
- [229] Kearns, M., Mansour, Y., and Singh, S., "Fast planning in stochastic games," *Proceedings of UAI*, 2000, 309-316.
- [230] Kearns, M.J., and Singh, S.P., "Near-optimal reinforcement learning in polynomial time," *Machine Learning*, **49** (2002), 209-232.
- [231] Keeney, R., and Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York, 1976.
- [232] Kemeny, J.G., "Mathematics without numbers," *Daedalus*, **88** (1959), 575-591.
- [233] Kemeny, J.G., and Snell, J.L., *Mathematical Models in the Social Sciences*, Blaisdell, New York, 1962. Reprinted by MIT Press, Cambridge, MA, 1972.
- [234] Knuth, D.C., Papadimitriou, C., and Tsitsiklis, N., "A note on strategy elimination in bimatrix games," *Oper. Res. Letters*, **7** (1988), 103-108.
- [235] Koller, D., and Megiddo, N., "The complexity of two-person zero-sum games in extensive form," *Games and Economic Behavior*, **4** (1992), 528-552.
- [236] Koller, D., and Megiddo, N., "Finding mixed strategies with small supports in extensive form games," *International J. of Game Theory*, **25** (1996), 73-92.
- [237] Konieczny, S., and Perez, R.P., "Merging information under constraints: A logical framework," *Journal of Logic and Computation*, **12** (2002), 772-808.

- [238] Konieczny, S., and Perez, R.P., “Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links to social choice theory,” *European Journal of Operational Research*, **160** (2005), 785-802.
- [239] Koppius, O., “The need for speed: Non-price considerations in auction design at the Dutch flower auctions,” DIMACS Workshop on Auctions with Transaction Costs, DIMACS Center, Rutgers University, March 2007.
- [240] Kowalski, R., “Search strategies for theorem proving,” in B. Meltzer and D. Mitchie (eds.) *Machine Intelligence*, **5**, Edinburgh University Press, Edinburgh, 1969, 181-201.
- [241] Kowalski, R., “And-or graphs, theorem proving graphs and bi-directional search,” in B. Meltzer and D. Mitchie (eds.) *Machine Intelligence*, **7**, Edinburgh University Press, Edinburgh, 1972, 167-194.
- [242] Kraus, S., Lehmann, D., and Magidor, M., “Nonmonotonic reasoning, preferential models and cumulative logics,” *Artificial Intelligence*, **44** (1990), 167-207.
- [243] Konieczny, S., and Perez, R.P., “Merging information under constraints: a logical framework,” *Journal of Logic and Computation*, **12** (2002), 772-808.
- [244] Konieczny, S., and Perez, R.P., “Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory,” *European Journal of Operational Research*, **160** (2005), 785-802.
- [245] Kum, H-C., Pei, J., Wang, W., and Duncan, D., “ApproxMAP: Approximate mining of consensus sequential patterns,” *Proc. of the 3rd SIAM International Conference on Data Mining*, San Francisco, CA, 2003.
- [246] Kum, H-C., Duncan, D., Flair, K., and Wang, W., “Social welfare program administration and evaluation and policy analysis using knowledge discovery and data mining (KDD) on administrative data,” *Proceedings of the NSF National Conference on Digital Government Research (D.G.O)*, 2003, 39-44.
- [247] Larson, K., “Reducing costly information acquisition in auctions,” *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, ACM, 2006, 1167 - 1174.
- [248] 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.
- [249] Larson, K., and Sandholm, T., “Bargaining with limited computation: Deliberation equilibrium,” *Artificial Intelligence*, **132** (2001), 183-217.
- [250] Larson, K., and Sandholm, T., “Costly valuation computation in auctions,” *TARK*, 2001, 169-182.
- [251] Lauritzen, S.L., *Graphical Models*, Clarendon Press, Oxford, UK, 1996.
- [252] Lauritzen, S.L. and Nilsson, D., “Representing and solving decision problems with limited information,” *Management Science*, **47** (2001), 1238-1251.
- [253] Lebanon, G., and Lafferty, J., “Cranking: Combining rankings using conditional probability models on permutations,” *Machine Learning: Proc. Nineteenth International Conference*, 2002.
- [254] Lehmann, D., “Generalized qualitative probability: Savage revisited,” in *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence, UAI96*, Morgan Kaufmann, San Francisco, 1996, 381-388.
- [255] Lehmann, D., “Expected qualitative utility maximization,” *Games and Economic Behavior*, **35** (2001), 54-79.
- [256] Lehmann, D., “Nonmonotonic logics and semantics,” *Journal of Logic and Computation*, **11** (2001), 229-256.

- [257] Lehrer, E., "Repeated games with stationary bounded recall strategies," *J. Econ. Theory*, **46** (1988), 130-144.
- [258] Lehrer, E., "Finitely many players with bounded recall in infinitely repeated games," *Games and Economic Behavior*, **7** (1994), 390-405.
- [259] Li, C.-S., and Ramaswami, R., "Automatic fault detection, isolation, and recovery in transparent all-optical networks," *Journal of Lightwave Technology*, **15** (1997), 1784-1793.
- [260] Lin, J., "Integration of weighted knowledge bases," *Artificial Intelligence*, **83** (1986), 363-378.
- [261] Lin, J., and Mendelzon, A.O., "Merging databases under constraints," *International Journal of Cooperative Information Systems*, **7** (1998), 55-76.
- [262] Linial, N., "Game-theoretic aspects of computing," in R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory with Economic Applications*, **II**, chapter 38, (1994), 1340-1395.
- [263] Linial, N., Peleg, D., Rabinovitch, Y., and Saks, M., "Sphere packing and local majorities in graphs," *Proc. 2nd ISTCS*, 1993, 141-149.
- [264] Littman, M., Kearns, M., and Singh, S., "An efficient algorithm for singly connected games," *Neural Information Processing Systems*, 2002.
- [265] Littman, M.L., Ravi, N., Fenson, E., and Howard, R., "An instance-based state representation for network repair," in *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI)*, 2004, 287-292.
- [266] Littman, M.L., and Stone, P., "A polynomial-time (Nash) equilibrium algorithm for repeated games," *Decision Support Systems*, **39** (2005), 55-66.
- [267] Lobo, M., and Boyd, S., "Pricing and learning with uncertain demand," preprint, Fuqua School of Business, Duke University, 2003.
- [268] Loehman, E., and Whinston, A., "An axiomatic approach to cost allocation for public investment," *Public Finance Q.*, **2** (1974), 236-251.
- [269] Lucas, W.F., "Applications of cooperative games to equitable allocations," *Proc. of Symposia in Applied Mathematics*, American Mathematical Society, 1981, 19-36.
- [270] Madigan, D. Mittal, S., and Roberts, F.S. "Sequential decision making algorithms for port of entry inspection: Overcoming computational challenges," to appear in *Proceedings of the International Conference on Intelligence and Security Informatics*, 2007.
- [271] Madigan, D. Mittal, S., and Roberts, F.S., "Efficient sequential decision making algorithms for container inspection operations," preprint, DIMACS Center, Rutgers University, 2007.
- [272] Mahadev, N.V.R., Pekeč, A., and Roberts, F.S., "Effect of change of scale on optimality in a scheduling model with priorities and earliness/tardiness penalties," *Mathematical and Computer Modelling*, **25** (1997), 9-22.
- [273] Mahadev, N.V.R., Pekeč, A., and Roberts, F.S., "On the meaningfulness of optimal solutions to scheduling problems: Can an optimal solution be non-optimal?," *Operations Research*, **46** **supp.** (1998), S120-S134.
- [274] Mahadev, N.V.R., and Roberts, F.S., "Consensus list colorings of graphs and physical mapping of DNA," in M. Janowitz, F.-J. Lapointe, F.R. McMorris, B. Mirkin, and F.S. Roberts (eds.), *Bioconsensus*, DIMACS Series, American Mathematical Society, Providence, RI, **61** (2003), 83-95.
- [275] Mahajan, S. and van Ryzin, G.J., "Inventory competition under dynamic consumer substitution," *Operations Research*, **49** (2001), 646-657.

- [276] Mahajan, S. and van Ryzin, G.J., "Stocking retail assortments under dynamic consumer substitution," *Operations Research*, **49** (2001), 334-351.
- [277] Maniquet, F., "A characterization of the Shapley value in queueing problems," *Journal of Economic Theory*, **109** (2003), 90-103.
- [278] Mannor, S., Simister, D.I., Sun, P., and Tsitsiklis, J.N., "Biases and variances in value function estimates," Fuqua School of Business, Duke University, Faculty Research paper FRPS05-199, August 2004. Preliminary version "Bias and variance in value function estimation," *Proceedings of the 21st International Conference on Machine Learning*, Banff, Canada, 2004.
- [279] Marcotorchino, P. and Michaud, P., "Heuristic approach of the similarity aggregation problem," *Methods of Oper. Res.*, **43** (1981), 395-404.
- [280] McCall, J.A., Richards, P.K., and Walters, G.F., "Factors in software quality," RADC TR-77-369, 1977. Also in US Rome Air Development Center Reports NTIS AD/A-049 014, 015, 055, **I**, **II**, **III**, 1977.
- [281] McCallum, R.A., "Instance-based utile distinctions for reinforcement learning with hidden state," *Proceedings of the 12th International Conference on Machine Learning*, Morgan Kaufmann, San Francisco, 1995, 387-395.
- [282] McGaley, M., "Report on DIMACS Workshop on Electronic Voting – Theory and Practice," <http://dimacs.rutgers.edu/Workshops/Voting/e-voting-final.pdf>, December 2004.
- [283] McGill, J. and van Ryzin, G.J., "Revenue management: Research overview and prospects," *Transportation Science*, **33** (1999), 233-256.
- [284] McMorris, F.R., Mulder, M., and Roberts, F.S., "The median procedure on median graphs," *Discrete Applied Math.*, **84** (1998), 165-181.
- [285] McMorris, F.R., Roberts, F.S., and Wang, C., "The center function on trees," *Networks*, **38** (2001), 84-87.
- [286] 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.
- [287] Mirkin, B., and Roberts, F.S., "Consensus functions and patterns in molecular sequences," *Bull. Math. Bio.*, **55** (1993), 695-713.
- [288] Mirrokni, V.S., and Vetta, A., "Convergence issues in competitive games," *RANDOM-APPROX 2004*.
- [289] Miyakawa, M., "Criteria for selecting a variable in the construction of efficient trees," *IEEE Transactions on Computers*, **38** (1989), 130-141.
- [290] Mizuta, H., and Steiglitz, K., "Agent-based simulation of dynamic online auctions," *Winter Simulation Conference*, Orlando, FL, Dec. 2000.
- [291] Moffet, J.D., and Lupu, E.C., "The uses of role hierarchies in access control," *Proc. 4th ACM Workshop on Role-based Access Control*, Fairfax, VA, 1999, 153-160.
- [292] Mongin, P., "Does optimization imply rationality?," *Synthese* **124** (2000), 73-111.
- [293] Monma, C.L., and Sidney, J.B., "Optimal sequencing via modular decomposition: Characterization of sequencing functions," *Mathematics of Operations Research*, **4** (1979), 215-224.
- [294] Morgenstern, M., "Security and inference in multilevel database and knowledge-base systems," *Proc. SIGMOD ACM Special Interest Group on Management of Data*, ACM, 1987.
- [295] Morgenstern, M., "Controlling logical inference in multilevel database systems," *IEEE Symp. on Security and Privacy*, Oakland, CA, 1988.

- [296] Moore, A.W., Atkeson, C.G. and Schaal, S., "Memory-based learning for control," Technical Report CMU-RI-TR-95-18, CMU Robotics Institute, 1995.
- [297] Mor, Y., and Rosenschein, J., "Time and the prisoner's dilemma," *Proceedings of the First International Conference on Multi-agent Systems ICMAS-95*, AAAI Press, Menlo Park, CA, 1995, 276-292.
- [298] Morisio, M., Stamelos, I., and Tsoukiàs, A., "A new method to evaluate software artifacts against predefined profiles," in *Proceedings of the SEKE-02 conference*, ACM Press, New York, 2002, 811-818.
- [299] Morisio, M., and Tsoukiàs, A., "IUSWARE: A formal methodology for software evaluation and selection," *IEEE Proceedings on Software Engineering*, **144** (1997), 162-174.
- [300] Moulin, H., "Equal or proportional division of a surplus, and other methods," *International J. of Game Theory*, **16** (1987), 161-186.
- [301] Moulin, H., "Incremental cost sharing: Characterization by coalition strategy-proofness," *Social Choice and Welfare*, **16** (1999), 279-320.
- [302] Moulin, H., and Shenker, S., "Strategyproof sharing of submodular costs: Budget balance versus efficiency," *Economic Theory*, **18** (2001), 511-533.
- [303] Moulin, H., and Stong, R., "Fair queuing and other probabilistic allocation methods," *Mathematics of Operations Research*, **27** (2002), 1-30.
- [304] Mundici, D., "Functions computed by monotone boolean formulas with no repeated variables," *Theoretical Computer Science*, **66** (1989), 113-114.
- [305] Myerson, R., *Game Theory*, Harvard University Press, Cambridge, MA, 1991.
- [306] Natarajan, K.S., "Optimizing depth-first search of AND-OR trees," Technical report, IBM T.J. Watson Research Center, Yorktown Heights, NY, 1986.
- [307] Nilsson, N.J., *Problem-solving Methods in Artificial Intelligence*, McGraw-Hill, New York, 1971.
- [308] Naor, M., Pinkas, B., and Sumner, R., "Privacy preserving auctions and mechanism design," *Proceedings 1st ACM Conference on Electronic Commerce, Denver, Colorado*, November 1999.
- [309] Neyman, A., "Bounded complexity justifies cooperation in finitely repeated prisoner's dilemma," *Economics Letters*, **19** (1985), 227-229.
- [310] 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.
- [311] Neyman, A., "Finitely repeated games with finite automata," *Math. Op. Res.*, **23** (1999), 513-552.
- [312] Neyman, A., and Okada, D., "Strategic entropy and complexity in repeated games," *Games and Econ. Behavior*, **29** (1999), 191-223.
- [313] Neyman, A., and Okada, D., "Repeated games with bounded entropy," *Games and Econ. Behavior*, **30** (2000), 228-247.
- [314] Neyman, A., and Okada, D., "Two person repeated games with finite automata," *Int. J. Game Theory*, **29** (2000), 309-325.
- [315] Ng, K-B., *An Investigation of the Conditions for Effective Data Fusion in Information Retrieval*, PhD Thesis, Rutgers University, New Brunswick, NJ, 1998.
- [316] Niemi, V., and Renvall, A., "How to prevent buying of votes in computer elections," *Advances in Cryptology - ASIACRYPT'94*, Lectures Notes in Computer Science, **917** (1994), 164-170.
- [317] Nisan, N., "Algorithms for selfish agents," *Proc. 16th Symposium on Theoretical Aspects of Computer Science STACS99*, Lecture Notes in Computer Science, Springer-Verlag, Berlin, **1563** (1999), 1-17.

- [318] Nisan, N., and Ronen, A., "Algorithmic mechanism design," *Games and Economic Behavior*, **35** (2001), 166-196.
- [319] Nisan, N., and Ronen, A., "Computationally feasible VCG mechanisms," *Games 2000*.
- [320] Nyanchama, M., and Osborn, S., "The role graph model," *Proc. 1st ACM Workshop on Role-based Access Control*, Gaithersburg, MD, 1995, 12.
- [321] Okada, N., Hashimoto, T., and Young, P., "Cost allocation in water resources development," *Water Resources Research*, **18** (1982), 361-373.
- [322] O'Neill, R.P., Sotkiewicz, P.M., Hobbs, B.F., Rothkopf, M.H., and Stewart, W.R., "Efficient market-clearing prices in markets with nonconvexities," *European J. of Operations Research*, **164** (2005), 269-285.
- [323] Osada, R., Funkhouser, R., Chazelle, B., and Dobkin, D., "Matching 3D models with shape distributions," *Shape Modeling International*, 2002, 154-166.
- [324] Osborn, S., and Guo, Y., "Modeling users in role-based access control," *Proc. 5th ACM Workshop on Role-based Access Control*, Berlin, 2000, 31-37.
- [325] Osborn, S., Sandhu, R., and Munawer, Q., "Configuring role-based access control to enforce mandatory and discretionary access control policies," *ACM Trans. on Information and System Security TISSEC*, **3** (2000), 85-106.
- [326] Papadimitriou, C., "On players with a bounded number of states," *Games and Econ. Behavior*, **4** (1992), 122-131.
- [327] Papadimitriou, C.H., "Algorithms, games, and the Internet," *Proceedings of Symposium on Theory of Computing STOC*, 2001.
- [328] Papadimitriou, C.H., and Yannakakis, M., "On the value of information in distributed decision-making," *PODC 91*.
- [329] Papadimitriou, C., and Yannakakis, M., "On complexity as bounded rationality," *STOC-94*, 1994, 726-733.
- [330] Park, C, Itoh, K., and Kurosawa, K., "All/nothing election scheme and anonymous channel," *Eurocrypt '93*, 248-259.
- [331] Park, S., and Rothkopf, M., "Auctions with bidder-determined allowable combinations," *European Journal of Operational Research*, **161** (2005), 399-415.
- [332] Paschetta, E., and Tsoukiàs, A., "A real world mcda application: Evaluating software," *Journal of Multi-Criteria Decision Analysis*, **9** (2000), 205-226.
- [333] Pawlak, Z., *Rough Sets - Theoretical Aspects of Reasoning about Data*, Kluwer Academic, Dordrecht, 1991.
- [334] Pease, M., Shostak, R., and Lamport, L., "Reaching agreement in the presence of faults," *JACM*, **27** (1980), 228-234.
- [335] Pekeč, A., "How to design a second-price combinational auction," *INFORMS 2000*.
- [336] Pekeč, A., and Rothkopf, M.H., "Combinatorial auction design," *Management Science*, **49** (2003), 1485-1503.
- [337] Peleg, D., "Local majority voting, small coalitions and controlling monopolies in graphs: A review," *Proc. 3rd Colloq. on Structural Information and Communication Complexity*, 1996, 170-179.
- [338] Peleg, D., "Size bounds for dynamic monopolies," *Discrete Appl. Math.*, **86** (1998), 263-273.

- [339] Peleg, G., and Wool, A., "The availability of quorum systems," *Information and Computation*, **123** (1995), 210-223.
- [340] Pennock, D.M., Horvitz, E., and Giles, C.L., "Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering," *Proc. 17th National Conference on Artificial Intelligence AAAI-2000*, AAAI Press, Menlo Park, CA, 2000, 729-734.
- [341] Pennock, D.M., Maynard-Reid, P., Giles, C.L., and Horvitz, E., "A normative examination of ensemble learning algorithms," *Proc. 17th International Conference on Machine Learning (ICML-2000)*, Morgan-Kaufmann, 2000, 735-742.
- [342] Perny, P., "Multicriteria filtering methods based on concordance/non-discordance principles," *Annals of Operations Research*, **80** (1998), 137-167.
- [343] Perny, P., and Spanjaard, O., "Preference-based search in state space graphs," *Proceedings of the AAAI'2002 Conference*, Edmonton, Canada, July 2002, 751-756.
- [344] Perny, P., and Spanjaard, O., "A preference-based approach to spanning trees and shortest paths problems," *European Journal of Operational Research*, **162** (2005), 584-601.
- [345] Perny, P., and Zucker, J.D., "Preference-based search and machine learning for collaborative filtering: The 'film-conseil' recommender system," *Information, Interaction, Intelligence*, **1** (2001), 9-48.
- [346] Pikhurko, O., "On envy-free cake division," *Amer. Math. Monthly*, **107** (2000), 736-738.
- [347] Pirlot, M., and Vincke, P., *Semiororders: Properties, Representations, Applications*, Kluwer Academic Publishers, Dordrecht, 1997.
- [348] Pohl, I., "Bi-directional search," in B. Meltzer and D. Mitchie (eds.) *Machine Intelligence*, **6**, Edinburgh University Press, Edinburgh, 1971, 127-140.
- [349] Poljak, S., and Sura, M., "On periodical behavior in society with symmetric influences," *Combinatorica*, **3** (1983), 119-121.
- [350] Poole, D., "Decision-theoretic defaults," in *Proceedings of the Ninth Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, Morgan Kaufmann, San Francisco, 1992, 190-197.
- [351] Prakken, H., and Vreeswijk, G., "Logics for defeasible argumentation," in D. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic*, Vol 4, Kluwer Academic, Dordrecht, 2002, 219-318.
- [352] Quinlan, J.R., "Induction of decision trees," *Machine Learning*, **1** (1986), 81-106.
- [353] Rasmusen, E., "Strategic implications of uncertainty over one's own private value in auctions," *Advances in Theoretical Economics*, **6** (2006), 1261.
- [354] Reinwald, L.T., and Soland, R.M., "Conversion of limited-entry decision tables to optimal computer programs I: Minimum average processing time," *Journal of the Association for Computing Machinery*, **13** (1966), 339-358.
- [355] Resnick, P., and Varian, H.R., "Recommender systems," *Communications of the ACM*, **40** (1997), 56-58.
- [356] Roberts, F.S., *Discrete Mathematical Models, with Applications to Social, Biological, and Environmental Problems*, Prentice-Hall, Englewood Cliffs, NJ, 1976.
- [357] Roberts, F.S., *Measurement Theory, with Applications to Decisionmaking, Utility, and the Social Sciences*, Addison-Wesley, Reading, MA, 1979.
- [358] Roberts, F.S., "Applications of the theory of meaningfulness to psychology," *J. Math. Psychol.*, **29** (1985), 311-332.

- [359] Roberts, F.S. (ed.), *Applications of Combinatorics and Graph Theory in the Biological and Social Sciences*, IMA Volumes in Mathematics and its Applications, Springer-Verlag, New York, 1989.
- [360] Roberts, F.S., "Merging relative scores," *J. Math. Anal. and Appl.*, **147** (1990), 30-52.
- [361] Roberts, F.S., "Characterizations of the plurality function," *Math. Soc. Sci.*, **21** (1991), 101-127.
- [362] Roberts, F.S., "On the indicator function of the plurality function," *Math. Soc. Sci.*, **22** (1991), 163-174.
- [363] Roberts, F.S., "Limitations on conclusions using scales of measurement," in A. Barnett, S.M., Pollock, and M.H. Rothkopf (eds.), *Operations Research and the Public Sector*, Elsevier, Amsterdam, 1994, 621-671.
- [364] Roberts, F.S., "A functional equation that arises in problems of scheduling with priorities and lateness/earliness penalties," *Math. and Computer Modelling*, **21** (1995), 77-83.
- [365] 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, SIAM, Philadelphia, PA, 2003, 1-34.
- [366] Roberts, F.S., "Decision support algorithms for port-of-entry inspection," *Working Together: Research and Development Partnerships in Homeland Security, Proceedings of DHS/IEEE Conference*, Boston, 2005.
- [367] Robertson, J.M., and Webb, A., "Near exact and envy-free cake division," *Ars Combin.*, **45** (1997), 97-108.
- [368] Rothkopf, M.H., Pekeć, A., and Harstad, R.M., "Computationally manageable combinatorial auctions," *Management Sci.*, **44** (1998), 1131-1147.
- [369] Rothschild, M., "A two-armed bandit theory of market pricing," *J. of Economic Theory*, **9** (1974), 185-202.
- [370] Roughgarden, T. "The price of anarchy is independent of the network topology," *Journal of Computer and System Sciences*, **67** (2003), 341-364.
- [371] Roughgarden, T., *Selfish Routing and the Price of Anarchy*, MIT Press, May 2005.
- [372] Roughgarden, T., and Tardós, E., "How bad is selfish routing?," *Proceedings of IEEE Foundations of Computer Science*, 2000.
- [373] Rubin, A., "Security considerations for remote electronic voting," *Comm. of the ACM*, **45** (2002), 39-44.
- [374] Rubinstein, A. "Finite automata play the repeated Prisoner's Dilemma," *J. Econ. Theory*, **39** (1986), 83-96.
- [375] Rubinstein, A., "The electronic mail game: A game with almost common knowledge," *Amer. Economic Review*, **79** (1989), 385-391.
- [376] Rubinstein, A., *Modeling Bounded Rationality*, MIT Press, Cambridge, MA, 1998.
- [377] Sabbadin, R., "Possibilistic Markov decision processes," *Engineering Applications of Artificial Intelligence* **14** (2001), 287-300.
- [378] Sabbadin, R., Fargier, H., and Lang, J., "Towards qualitative approaches to multi-stage decision making," *International Journal of Approximate Reasoning*, **19** (1998), 441-471.
- [379] Sako, K., and Kilian, J., "Receipt-free mix-type voting scheme," *Eurocrypt '95*, 393-403.

- [380] Salloum, S., and Breuer, M.A., "An optimum testing algorithm for some symmetric coherent systems," *Journal of Mathematical Analysis and Applications*, **101** (1984), 170-194.
- [381] Sandholm, T., "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, **135** (2002), 1-54.
- [382] Sandholm, T. and Suri, S. "Side constraints and non-price attributes in markets," *International Joint Conference on Artificial Intelligence (IJCAI), Proceedings of the Workshop on Distributed Constraint Reasoning*, Seattle, WA, August 4th, 2001.
- [383] Sandholm, T. and Suri, S., "Optimal clearing of supply/demand curves." In *Proceedings of the 13th Annual International Symposium on Algorithms and Computation (ISAAC)*, Vancouver, Canada. 2002.
- [384] Sandhu, R., and Munawer, Q., "How to do discretionary access control using roles," *Proc. 3rd ACM Workshop on Role-based Access Control*, Fairfax, VA, 1998, 47-54.
- [385] Shachter, R.D., "Evaluating influence diagrams," *Operations Research*, **34** (1986), 871-882.
- [386] Schapire, R.E., Freund, Y., Bartlett, P., and Lee, W.S., "Boosting the margin: A new explanation for the effectiveness of voting methods," *The Annals of Statistics*, **26** (1998), 1651-1686.
- [387] Schuba, C.L., Krsul, I., Kuhn, M.G., Spafford, E.G., Sundaram, A., and Zamboni, D., "Analysis of a denial of service attack on TCP," *Proc. 1997 IEEE Symp. on Security and Privacy*, 208-233.
- [388] Scott, D., "Some ordered sets in computer science," in I. Rival (ed.), *Ordered Sets*, D. Reidel, Dordrecht, 1982, 677-718.
- [389] Sharkey, W., "Network models in economics," in M.O. Ball, et al. (eds.), *Network Routing*, Handbook in Operations Research and Management Science, North-Holland, Amsterdam, **8** (1995).
- [390] Shoham, Y., "Nonmonotonic logic: Meaning and utility," in *Proceedings of the 10th International Joint Conference in Artificial Intelligence, IJCAI87*, Morgan Kaufmann, San Francisco, 1987, 388-393.
- [391] Simon, H.A., *Models of Bounded Rationality*, Vol. 2, MIT Press, 1982.
- [392] Simon, H.A., and Kadane, J.B., "Optimal problem-solving search: All-or-none solutions," *Artificial Intelligence*, **6** (1975), 235-247.
- [393] Singh, S., Kearns, M., and Mansour, Y., "Nash convergence of gradient dynamics in general-sum games," *UAI 2000*.
- [394] Shubik, M., "Incentives, decentralized controls, the assignment of joint costs and internal pricing," *Management Sci.*, **8** (1962), 325-343.
- [395] Smart, W.D., and Kaelbling, L.P., "Practical reinforcement learning in continuous spaces," *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, 2000, 903-910.
- [396] Smith, D.E., "Controlling backward inference," *Artificial Intelligence*, **39** (1989), 145-208.
- [397] Stahl, D., and Wilson, P., "On players' models of other players – theory and experimental evidence," *Games and Economic Behavior*, **10** (1995), 218-254.
- [398] Stallard, N., Thall, P.F., and Whitehead, J., "Decision theoretic designs for phase II clinical trials with multiple outcomes," *Biometrics*, **55** (1999), 971-977.
- [399] Stroud, P.D. and Saeger, K.J., "Enumeration of increasing boolean expressions and alternative digraph implementations for diagnostic applications," in H. Chu, J. Ferrer, T. Nguyen, and Y. Yu (eds), *Proceedings Volume IV, Computer, Communication and Control Technologies: I*, International Institute of Informatics and Systematics, Orlando, FL, 2003, 328-333.
- [400] Stefanowski, J., and Tsoukiàs, A., "Incomplete information tables and rough classification," *Computational Intelligence*, **17** (2001), 454-466.

- [401] Stefanowski, J., and Tsoukiàs, A., "Valued tolerance and decision rules," in W. Ziarko and Y. Yao (eds.), *Rough Sets and Current Trends in Computing*, Springer Verlag, LNAI 2005, Berlin, 2001, 212-219.
- [402] Stefanowski, J., and Tsoukiàs, A., "Induction of decision rules and classification in the valued tolerance approach," in J.J. Alpigini, J.F. Peters, A. Skowron, and N. Zhong (eds.) *Proceedings of the RSCTC 2002 Conference*, LNAI 2475, Springer Verlag, Berlin, 2002, 271-278.
- [403] Stefanowski, J., and Vanderpooten, D., "Induction of decision rules in classification and discovery-oriented perspectives," *International Journal on Intelligent Systems*, **16** (2001), 13-27.
- [404] Straffin, P.D., and Heaney, J.P., "Game theory and the Tennessee Valley Authority," *Int. J. of Game Theory*, **10** (1981), 35-43.
- [405] Talluri, K. and van Ryzin, G.J., "An analysis of bid-price controls for network revenue management," *Management Science*, **44** (1998), 1577-1593.
- [406] Talluri, K. and van Ryzin, G.J., "A randomized linear programming method for computing network bid prices," *Transportation Science*, **33** (1999), 207-216.
- [407] Tan, S.W., and Pearl, J., "Qualitative decision theory," in *Proceedings of the 12th National Conference on Artificial Intelligence, AAAI94*, MIT Press, Cambridge, 1994, 928-933.
- [408] Tasnádi, A., "On probabilistic rationing methods," *Mathematical Social Sciences*, **44** (2002), 211-221.
- [409] Thomas, A., *A Behavioral Analysis of Joint Cost Allocation and Transfer Pricing*, Arthur Andersen and Company Lecture Series, Stipes Publishing Company, 1980.
- [410] Thuraisingham, B., "The use of conceptual structures for handling the inference problem," in C.E. Landwehr and S. Jajodia (eds.), *Database Security V: Status and Prospects, Results of the 5th IFIP WG 11.3 Working Group Conference on Database Security*, North-Holland, 1991, 333-362.
- [411] Thuraisingham, B., Ford, W., Collins, M., and O'Keefe, J., "Design and implementation of a database inference controller," *Data and Knowledge Engineering*, **11** (1993), 271-298.
- [412] Trotter, W.T., *Combinatorics and Partially Ordered Sets: Dimension Theory*, The Johns Hopkins University Press, Baltimore, MD, 1992.
- [413] Tsoukiàs, A., "From decision theory to decision aiding methodology," DIMACS Technical Report 2003-21, DIMACS Center, Rutgers University, Piscataway, NJ, 2003.
- [414] Tsoukiàs, A., and Vincke, Ph., "A characterization of PQI interval orders," *Discrete Applied Mathematics*, **127** (2003), 387-397.
- [415] Tzafestas, S.G., and Dalianis, P.J., "Fault diagnosis in complex systems using artificial neural networks," *Proc. 3rd IEEE Conf. on Control Applications*, **2** (1994), 877-882.
- [416] Vickrey, D., and Koller, D., "Multi-agent algorithms for solving graphical games," *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2002.
- [417] Vincke, P., "Exploitation of a crisp relation in a ranking problem," *Theory and Decision*, **32** (1992), 221-240.
- [418] Vincke, P., *Multicriteria Decision-Aid*, Wiley, New York, 1992.
- [419] DeVries, S., and Vohra, R.V., "Combinatorial auctions: A survey," *INFORMS Journal of Computing*, **15** (2003), 284-309.
- [420] Wakabayashi, Y., *Aggregation of Binary Relations: Algorithmic and Polyhedral Investigations*, Thesis, Augsburg, 1986.

- [421] Wang, J., and Vande Vate, J., "Question asking strategies for Horn clause systems," *AMAI*, **1** (1990), 359-370.
- [422] I. Wegener. *The Complexity of Boolean Functions*, Wiley-Teubner, 1987.
- [423] Weibull, J., *Evolutionary Game Theory*, MIT Press, 1995.
- [424] Wellman, M.P., and Doyle, J., "Preferential semantics for goals," in *Proceedings of the 9th National Conference on Artificial Intelligence, AAAI91*, AAAI Press, Menlo Park, 1991, 698-703.
- [425] Weymark, J.A., "Arrow's theorem with social quasi-orderings," *Public Choice*, **42** (1984), 235-246.
- [426] Williams, L. V., "Information efficiency in betting markets: A survey," *Bulletin of Economic Research*, **51** (1999), 1-30.
- [427] Yener, A., and Rose, C., "Highly mobile users and paging: Optimal polling strategies," *IEEE Transactions on Vehicular Technology*, **47** (1998), 1251-1257.
- [428] Yilankaya, O, and Celik, G., "Optimal auctions with simultaneous and costly participation," Micro Theory Working Papers celik-05-05-09-03-55-40, 2005.
- [429] Young, H.P., *Cost Allocation: Methods, Principles, Applications*, North-Holland, Amsterdam, 1985.
- [430] Zhang, H., Schroepfer, C. and Elsayed E. A., "Sensor thresholds in port-of-entry inspection systems," *Proceedings of the 12th ISSAT International Conference on Reliability and Quality in Design*, Chicago, Illinois, USA, August 3-5, 2006, 172-176.
- [431] Zhang, N.L., Qi, R., and Poole, D., "A computational theory of decision networks," *International Journal of Approximate Reasoning*, **11** (1994), 83-158.
- [432] Zhang, Y., Xu, W., and Callan, J., "Exploration and exploitation in adaptive filtering based on Bayesian active learning," *Proceedings of the Twentieth International Conference on Machine Learning (ICML-2003)*, 2003.
- [433] DIMACS Workshop and Working Group Meeting on Bioconsensus, October 2000, <http://dimacs.rutgers.edu/Workshops/Bioconsensus/>.
- [434] DIMACS Tutorial and Workshop on Bioconsensus II, October 2001, <http://dimacs.rutgers.edu/Workshops/BioconII/>.
- [435] DIMACS Workshop on Applications of Lattices and Ordered Sets to Computer Science, July 2003, <http://dimacs.rutgers.edu/Workshops/Lattices/>.
- [436] DIMACS Workshop on Electronic Voting – Theory and Practice, May 2004, <http://dimacs.rutgers.edu/Workshops/Voting/>.
- [437] DIMACS Workshop on Markets as Predictive Devices (Information Markets), <http://dimacs.rutgers.edu/Workshops/Markets/>.
- [438] DIMACS Workshop on Yield Management and Dynamic Pricing, <http://dimacs.rutgers.edu/Workshops/Yield/>
- [439] Constructive Preference Elicitation for Multiple Criteria Decision Aid (MCDA), Decision Aiding Research Area, LAMSADE, <http://www.lamsade.dauphine.fr/mcda/siteTheme/p4.html>