MANAGEMENT SCIENCE
Vol. 51, No. 3, March 2005, pp. 316–327
ISSN 0025-1909 | EISSN 1526-5501 | 05 | 5103 | 0316
DOI 10.1287/mnsc.1040.0345

The Landscape of Electronic Market Design

G. Anandalingam
The Robert H. Smith School of Business, University of Maryland,
College Park, Maryland 20742-1815, ganand@rhsmith.umd.edu

Robert W. Day
Operations and Information Management, School of Business, University of Connecticut,
Storrs, Connecticut 06269-1041, bob.day@business.uconn.edu

S. Raghavan
The Robert H. Smith School of Business, University of Maryland,
College Park, Maryland 20742-1815, raghavan@umd.edu

This paper presents an introductory survey for this special issue of *Management Science* on electronic markets. We acquaint the reader with some fundamental concepts in the study of electronic market mechanisms, while simultaneously presenting a survey and summary of the essential literature in this area. Along the way, we position each of the papers presented in this special issue within the existing literature, demonstrating the deep impact of these 14 articles on an already broad body of knowledge.

Key words: auctions; electronic markets; economics; management science; information systems
History: Accepted by Hau L. Lee, former editor-in-chief.

1. Overview

Corporations are currently investing billions of dollars in changing the way they buy and sell goods, integrate their supply chains, and interact with their partners and employees. "Electronic markets" (or "eMarkets") are now surfacing as a means of achieving greater efficiency in almost every sphere of economic activity. The emerging scholarly research in this field has focused on the fundamentals of electronic markets: designing market mechanisms and auctions, understanding buyer/seller behavior in network-mediated markets, and the economic modeling of online markets. This special issue of *Management Science* contains state-of-the-art research on various aspects of electronic markets, with a heavy emphasis on eMarkets design. In this introductory piece, we acquaint the reader with some fundamental concepts in the study of electronic market mechanisms, while simultaneously presenting a survey and summary of the essential literature in this area. At the same time, we position each piece of newly presented research within the existing literature, demonstrating the deep impact of the papers in this special issue on an already broad body of knowledge.

Research on eMarkets design relies heavily on theory of auctions as developed in the economics literature through the study of strategic games (Vickrey 1961), and in the operations research literature as constrained optimization since the mid-twentieth century (Friedman 1956). More recently, there has been a great deal of intellectual cross-pollination from management science (MS), information systems (IS), and computer science (CS) in an attempt to develop new electronic markets providing better economic outcomes through centralized computational decision making, costless communication/coordination, and negligibly low search costs.

The computational difficulties presented by these markets are not readily untied from their economic context, making it difficult to apply a purely algorithmic approach that ignores strategic considerations. Conversely, the most general economic allocation problems exhibit a level of computational complexity ($\mathcal{NP}$-hardness) that requires algorithmic expertise, making it difficult to attack these market design problems from a purely economic standpoint. Together, this close interaction between computational and economic considerations gives rise to one of the richest fields of interdisciplinary research, demonstrated by the wide variety of technical contributions in the present issue devoted to the study of eMarkets.

Our introduction to this field of literature is focused heavily on the theory of auctions, which we approach from the perspectives of both applications and theory. At least 10 of the 14 pieces presented in this issue fall neatly into the category of auction theory, while the remaining 4 are closely related to auctions in their discussions of bargaining and coordinated price setting.

This is especially true if we consider a very general definition of an auction:

Definition. An *auction* is a mechanism of information submission, together with rules for assigning items and payments to participants based on this submitted information.

The MS discipline emphasizes the successful implementation of decision-making technology in real-world situations, so we begin with a quick survey of the successful and growing applications of auction theory, keeping this general definition in mind. The most widely recognized electronic auction venues are undoubtedly eBay and similar electronic venues for commerce that may be described as mostly consumer-to-consumer (C2C). The success of these markets demonstrates that the "electronic" services (a Web-based platform) provide added value even when the "market" structure is simple (one-seller, one-item, simple bidding rules). Although several interesting avenues of research receive attention from these C2C settings, the success and future of the MS approach to eMarkets stems from the use of auctions in business-to-business (B2B) commerce and government allocation problems. In these settings the value-added is derived from the ability of the market mechanism to elicit and process information, determining a more favorable outcome through a centralized market. The auctions used in these settings typically have a "combinatorial" component, in which bidders may express their preferences over combinations or bundles of items. This quality promises to extract value from the auction decision mechanism and places the computational challenge for the mechanism at the very frontier of what is possible through computational technique.

2. Auctions and Governmental Allocation

From the perspective of applications, auction mechanisms first entered the MS literature in the context of government-controlled allocation problems. Throughout this literature, auctions are proposed as a tool of deregulation. Rather than allowing government-controlled resources to be distributed via political decision-making criteria and regulatory control (typically an inefficient procedure), auctions provide a market mechanism to allocate government property fairly. Before such a paradigm shift can take place, researchers must first provide scientific evidence that the auction mechanism will produce favorable outcomes, and that the auction can be implemented smoothly and at a comfortable pace for all participants.

Rassenti et al. (1982) pioneered this line of work with one of the first efforts to confront multi-item auction problems with the computational techniques of integer programming (IP). In this work they propose an auction mechanism for the allocation of airport time-slots and give some of the first experimental verification (using test subjects) that the combinatorial auction paradigm can achieve more efficient outcomes than a noncombinatorial mechanism. Their auctions and corresponding winner-determination IPs were modest in size, reflecting the difficulty and lack of scalability in solving these hard allocation problems.

Ironically, though the airport time-slot application was one of the first proposed for the use of combinatorial auctions, the FAA is only now beginning to consider the use of auctions to control demand for landing/takeoff slots at congested airports. Ball et al. (2005) provide a description of the necessary and desirable features in an airport slot auction; a parallel and quite thorough investigation of the pros and cons of airport-slot auctions in the European market is given by the consulting firm DotEcon-Ltd (2001). Both emphasize the economic benefits of an auction mechanism for the allocation of landing slots at congested airports and advocate a combinatorial auction implementation to achieve the desired results. Most recently, the National Center of Excellence for Aviation Operations Research (NEXTOR) and the FAA have sponsored two workshops (NEXTOR 2002, 2004) exploring the use of auctions for congestion control. Tentative plans include the possibility of an auction for the allocation of airport landing slots at a major U.S. airport by 2007.

Another influential early study rooted in governmental applications is given by the Adaptive User Selection Mechanism (AUSM) of Banks et al. (1989), who proposed to allocate resources such as jobs on a supercomputer or mission time aboard a space station. Like Rassenti et al. (1982), this paper began to stir interest in auction research outside of economics. Though published in an economics journal, this work of Banks et al. suggests that scheduling and logistical problems (typically MS applications) could be handled by adapted auction mechanisms. Further, their approach was progressive in the use of decentralized computation, shifting computational burden away from the central decision-making entity and to the bidding agents themselves. This "agent-based" decision making has been a hot topic in several fields of research and has strongly influenced the more recent auction literature. Kelly and Steinberg (2000), for instance, propose the PAUSE auction for universal telephone service obligations, another governmental allocation problem receiving attention in the auction literature.

Indeed, potential governmental applications for auctions abound, though the one discussed most prominently in the literature seems to be auctions for the allocation of spectrum licenses. The sellers of these licenses (such as the Federal Communication

Commission (FCC) in the United States) have no significant costs to recover for each license granted, so that the value for such items are determined almost exclusively through competition on the demand side. More importantly, several licenses in the same frequency range must be acquired to form a functioning communications network. The value that a particular telecommunications firm places on a given set of licenses is therefore very heavily influenced by which other licenses it has received. This property suggests that a series of single-item auctions would not lead to an efficient outcome in this market and that a more complex auction mechanism should be considered. In the FCC auction-design debate of the early 1990s, it was determined that a "combinatorial" or "package bid" auction should not be considered because the general combinatorial auction winner-determination problem was computationally intractable (i.e., $\mathcal{NP}$-hard, see McAfee 1993). The FCC instead adopted the simultaneous ascending auction (SAA), which allows for simultaneous price discovery without package bidding, foregoing the potential benefits of package bidding due to its computational difficulty.

This underlying all-or-nothing assumption, that a hard computational problem should either be handled completely or not all, is antithetical to the MS approach and has been challenged in this context. As the seminal work of Rothkopf et al. (1998) points out, many special cases of an $\mathcal{NP}$-hard problem like the combinatorial auction winner-determination problem can be solved efficiently. They suggest several applications in which a combinatorial auction can be quickly implemented, including special cases in which the size or type of bundles that may receive bids are restricted without objection from the bidders. The debate surrounding the FCC's design problem helped introduce auction theory to the MS community as a set of complex decision problems, open to heuristic techniques and algorithms tailored to specific classes of instances. The flurry of research that followed paved the way for the current landscape of auction implementation, in which B2B auctions take place using market-specific structure and IP techniques, as discussed below. The FCC seems to have accepted the possible benefits of such a paradigm shift—that a "partial" or "restricted" combinatorial auction may be successfully implemented using limited package bidding and IP solvers—as evidenced by the proposed Spectrum Auction #31, discussed, for example, by Günlük et al. (2005) in this issue.

Although spectrum auctions and similarly complex markets dominate the decision science–type auction literature, auctions are commonly employed for several other governmental allocation problems that receive less attention for their computational difficulty. The most important examples in this group include treasury bill and electricity auctions, where the auctioned items have fewer distinguishing characteristics and demand can be expressed simply as quantities demanded of a few simple types. Treasury bills in particular are "commodity-like" and do not combine in complicated ways to form valuable bundles, making advanced auction/decision techniques unnecessary in terms of implementation, although the game-theoretic analysis of these markets may not be trivial (see Ausubel and Cramton 2002). It is, however, important that we include treasury auctions among the success stories of governmental auctions, where several "commodity-like" items can be auctioned in a straightforward way.

With spectrum auctions as an example of a market in which items combine to form bundles in a complicated way, potentially benefitting from the use of a combinatorial auction mechanism, and treasury auctions as an example where no such complications arise, making combinatorial machinery unnecessary, deregulated wholesale electricity markets seem to provide a middle ground. On one hand, the items at auction (kilowatt-hours of electrical power supply) are very commodity-like, i.e., divisible and deriving the majority of their value intrinsically, rather than through synergy with other auction items. On the other hand, some synergies do arise in the form of power generator start-up and no-load costs, so that a "pure-commodity" approach may distort a bidding power supplier's ability to communicate true costs of production. This leaves a bidder to either misrepresent her preferences on the side of caution, or to risk dissatisfaction with the results of the auction. The electricity market may therefore be treated with a "tâtonnement" procedure, familiar from the economics literature, modified to handle the nonconvexity imposed by the start-up and no-load costs. Market Design, Inc., an auctions-consulting firm, has successfully implemented such specialized "clock auctions" for the wholesale electricity markets in France and Alberta, Canada, for example. See Cramton (2003) for details on these techniques. (More recently a "clock-proxy" mechanism has been suggested by Ausubel et al. 2005.)

The regional U.S. electricity markets, such as the Pennsylvania–New Jersey–Maryland (PJM) system, take a markedly MS approach to electricity allocation, solving mixed integer programs (MIPs) for daily electricity dispatch. The possibility of a centralized restructuring of the U.S. energy market to allocate and price electricity using sophisticated IP-type auction techniques is currently under investigation by the Federal Energy Regulatory Commission. O'Neill et al. (2005) describe this research. In this issue, Wu and Kleindorfer (2005) present related research on exchanges for goods/services in an environment with both forward contracting and spot markets (as in the wholesale electricity market). Also related to supply chain auctions (discussed below), Wu and Kleindorfer present optimal strategies for buyers and sellers in a model of these markets, prescribing rules for determining the portion of capacity/demand to be negotiated as contracted options, as well as the setting of execution fees and reserve prices on the supply side. In addition, their model provides conditions under which such exchanges are efficient and sustainable as market-clearing platforms for contracting and spot purchases. Related work aimed directly at the electricity market is given by Kleindorfer and Li (2005).

3. Auctions and B2B Commerce

While the authorities on the government applications motivating much of the early combinatorial auction research have been slow to implement the change to a "package bidding" auction, pioneers of B2B e-commerce quickly recognized the potential benefits of advanced auction techniques and began conducting "combinatorial" auctions as market intermediaries in the mid-1990s. Today, several firms offer auction consulting services, helping a firm conduct an auction for the procurement of transportation lanes, raw materials, and other services. Procurement auctions (with a single buyer and several suppliers or sellers, also called reverse auctions) with at least some combinatorial machinery have been conducted by CombineNet (Levine 2003), IBM (Hohner et al. 2003), Logistics.com (Logistics.com 2001), Net Exchange (Ledyard et al. 2002), and Emptoris (Schneur 2004). These B2B auctions assure savings to the buyer through increased price competition and the ability to constrain product quality, delivery time, and other nonprice attributes of the items provided. Conversely, strategic, logistical, and scaling problems facing suppliers are often alleviated with bundle bidding, which may allow expression of volume discounts, incompatible products/services, and complementary products/services.

In a typical procurement auction, the buyer initiates the auction by inviting quotes from several competing suppliers for a particular set of services (or goods). In a combinatorial procurement auction the buyer often has the additional ability to place restrictions or constraints on the final allocation. These constraints may include, for example, upper and lower bounds on the number of sellers providing services, restrictions that a certain subset of complementary services must be provided by the same seller, and other "rules of the game." These complicating rules specified in the auction setup are prevalent, for example, in transportation or shipping-lane auctions, where both buyers and sellers of shipping services are interested in outcomes that are operationally economical or efficient.

For shippers bidding on a set of transportation paths, it may be important to express how internal constraints affect their contractual decision making (e.g., "I can move shipment A or shipment B but not both. Which one should I bid on?). Also, bidders in a shipping-lane auction frequently want to express which bundles of services they are most willing to provide. For example, certain shipping lanes may be complementary to pre existing contractual obligations of the shipper; a shipper would be especially eager to fill trucks that would otherwise be empty return trips under existing contracts. With a wide range of expressability over both substitutes and complements desirable for both buyers and sellers in the majority of proposed of B2B applications, the optimization approach to winner determination seems to dominate tâtonnement models modified to accept combinatorial information in this context. Indeed, much of the available literature on procurement auctions utilizes an IP approach, which is generally flexible enough to accept a wide range of logical statements of preference.

Sears Logistics held the first large-scale auction for shipping lanes using "combined value" or package bids in the mid-1990s with the software support of Net Exchange (Ledyard et al. 2002). Their encouraging results suggest that satisfaction with the market mechanism and efficiency may be increased on both sides of the market. Sears Logistics (the procurer of shipping lanes in this case) reported savings of $25 million (13%) in their first combined value auction, while shippers were able to eliminate uncertainty and exposure to winning incompatible or incomplete sets of lanes. With this success documented, it is not surprising that several other large corporations decided to introduce the combinatorial auction paradigm for shipping; Logistics.com reports the implementation of procurement auctions on behalf of Walmart Stores, Compaq Computer Co., Staples Inc., The Limited Inc., and Kmart Corporation (Logistics.com 2001). In addition, Elmaghraby and Keskinocak (2003) provide a case study of a successful shipping-lane auction for Home Depot. Providing new research at the frontier of supply chain auctions is the work of Chen et al. (2005) (contained in this issue), which incorporates both transportation costs and production decisions into the auction mechanism, promising more efficient outcomes. As noted earlier, Wu and Kleindorfer (2005) present (in this issue) a related model of B2B supply chain exchanges with integrated forward, option, and spot markets.

The acceptance of a combinatorial auction format using IP decision-making techniques for transportation procurement should not be surprising. Internal supply chain and logistics decisions are often approached from the mathematical programming perspective, making the use of optimization software for community or market decisions a natural transition. Although less apparent, procurement auctions are finding their way into more general settings, though more reports of success stories may be necessary before they are accepted on a large scale. Some notable success stories include those of CombineNet, who report 15% surplus gained by participants in procurement auctions for raw materials (such as coal and steel) and shipping lanes since 2001 (Levine 2003). In addition, IBM has implemented several combinatorial procurement auctions for Mars, Inc., emphasizing that benefits accrue on both sides of the market (implying overall efficiency gains) and that payback on Mars' investment was less than a year (Hohner et al. 2003). In general, experts advocate that a successful procurement auction may be conducted in markets with a relatively small static group of suppliers, and that mechanism design should emphasize efficiency rather than procurer profits to establish favorable long-term relationships.

4. Auction Theory

Several auctions have been studied and implemented for many years, the English auction and its variants being the most familiar. To determine a price that is satisfactory to both buyer and seller, an auctioneer names successively higher prices, and bidders respond with their willingness to accept these prices, or by dropping out of the auction. The auction concludes when only one bidder is willing to pay the current price. This last bidder receives the item at the price that is a single increment higher than the amount that the next highest bidder is willing to pay. This simple auction is both *incentive compatible* (truth telling is the best strategy) and *efficient* (the item goes to the bidder who values it the most).

Other longstanding auction formats for a single item include the Dutch auction, in which the price descends until the first bidder bids and wins the item, and the sealed-bid auction, in which bidders each submit a price for the item in question, with the highest bid winning. (Note that in finance, the term Dutch auction is used differently to refer to a uniform-price auction, as in the context of the recent Google IPO, causing some confusion.) Auction theorists show that under mild assumptions, the same strategic behavior should be expected for the Dutch and the "first-price" sealed-bid auction, in which the highest bidder pays the amount of her bid. The "second-price" sealed-bid auction (a sealed-bid auction together with the condition that the highest bidder wins the item at the price specified by the second-highest bidder) displays incentive compatibility. Under an assumption that a bidder's valuation does not depend on the valuation information revealed by her opponents, the second-price sealed-bid auction produces the same outcome as the English auction (see, for example, Krishna 2002).

These observations on the equivalence of various auction outcomes, together with the strategic analysis showing that a bidder has no ability to benefit from an untruthful strategy in certain simple auctions, provide the "classical" basis of auction theory. There are two major avenues of contemporary auction theory building on these classical foundations. With a behavioral approach, many researchers pursue more accurate models of bidders to explain the differences between empirical data and theoretical predictions. For example, the papers by Carare and Rothkopf (2005), Deltas and Engelbrecht-Wiggans (2005), Ding et al. (2005), and Engelbrecht-Wiggans and Kahn (2005) provide examples of behavioral auction theory within this issue. Alternatively, researchers design new auction formats for simultaneously auctioning multiple items, for example, the papers by Kwasnica et al. (2005) and Kwon et al. (2005), in this issue. The challenge addressed in this line of research is to combine computational feasibility, economic efficiency, and strategic lucidity in a multi-item auction format. Before discussing these research streams in greater depth, we must first make clear some underlying market scenarios and terminology.

4.1. Market Settings for Auctions

Table 1 illustrates the full array of market types for which auctions have been studied. Almost any combination of selections from each column yields a different type of market (the only exception being that the question of differentiation among items has no meaning when only a single item is to be auctioned). Together this enumerates 18 different market settings to which auctions may potentially be applied. Auctions with a single seller are often referred to as *forward* auctions, as this is the most familiar format. An auction with a single buyer and several sellers is often called a *reverse* auction, but is also known in the literature as a *procurement* auction.

Markets with many buyers and sellers have been called both *exchanges* and *double* auctions. Because there is a large amount of literature concerning exchanges (due to their importance in finance), it may be useful to make a distinction between exchanges and double auctions. Based on current business practices, we propose that an exchange may refer to an

Table 1 **Market Settings for Auctions**

Number of participants	Types of items	Number of items	Differentiation
One seller, many buyers One buyer, many sellers Many buyers and sellers	Divisible Discrete	Single-item Multiple-items	Identical items Distinct items

open market in which transactions have the opportunity to occur at any time and require only the agreement of the transactors, as in the stock market. A double auction, on the other hand, should refer to markets in which all transactions are decided by a central information-collecting entity, adhering to our earlier definition of an auction. Although we feel this distinction would eliminate some confusion, there is a great deal of inertia surrounding the use of the word exchange in reference to double auctions, and the reader should be aware that these two terms are often used synonymously.

4.2. Single-Item Auctions

Single-item auctions have been studied in economics as games of incomplete information for more than 40 years (Vickrey 1961), and may be considered well understood. Specific topics of interest include several models for the auction of a discrete item with varying assumptions on the behavior of bidders. These models vary according to the amount of information available to each bidder, whether each bidder knows his own valuation for certain or bases his value somewhat on what others think, and the amount of affiliation or correlation among the values of bidders.

For example, the seeming equivalence in outcome between the English and second-price auctions breaks down when bidders' utility is governed by affiliated signals (see Milgrom and Weber 1982). In this circumstance the English auction has higher expected revenue, because bidders in the English auction learn over the course of the auction and can bid more competitively with better information. In the second-price auction, on the other hand, lack of information about the value of the item induces conservative bidding. The complete result for single-item auctions is that in terms of maximizing auctioneer revenue, the English auction outperforms the second-price auction, which in turn outperforms the first-price auction. (See Krishna 2002 for a good overview of this fundamental material.)

As noted earlier, similar results hold that the Dutch auction may be considered strategically equivalent to the first-price sealed-bid auction. This equivalence is challenged, however, in the recent work of Carare and Rothkopf (2005), presented in this issue, which shows that *slow* Dutch auctions (in which the price drops slowly over the course of a few days in an Internet auction) may produce more revenue than the first-price auction (see also Lucking-Reiley 1999). In the case of Carare and Rothkopf, we see that the simplifying assumptions used in theoretical work may lead to conclusions that no longer hold when the model is extended to describe the real world more accurately. Here, the value of time must be included as a transaction cost in models of Dutch auctions, and failure to capture all costs in an auction model may skew a seller's decision-making criteria when selecting a particular auction.

Other interesting new topics in the world of single-item auctions include explaining the difference in equilibrium behavior between theoretical research and empirical study. Indeed, real-life bidders do not behave as game theory suggests they should, often not understanding the incentive implications of Vickrey-Clarke-Groves (VCG) pricing or the need for proper mitigation of the "winner's curse." In this issue, Deltas and Engelbrecht-Wiggans (2005) and Ding et al. (2005) provide a glimpse into the field of behavioral economics and its intersection with auction theory. The former provides justification for the persistence of strategically naive behavior, while the latter examines the effect of emotional factors (excitement and frustration) on bidder behavior and auction participation. In both cases the auction format is held to the simple case of a single item in order to study bidders' behavior, which may be very complex despite the simple setting.

4.3. Multiple-Item Auctions

The study of multiple-item auctions (or multiunit auctions) has received greater attention than single-item auctions in recent years, and consequently represents a greater portion of the auction literature presented here. Situations proposed for these auctions include forward, reverse, and double-auction settings for identical or distinct items that are either discrete or divisible (i.e., all possibilities). Researchers often assume that results for forward auctions hold for reverse settings, and use whichever terminology is convenient for their proposed area of application. Even though many results do apply generally, whether specific auction designs work better for different numbers and types of market participants is territory that has been explored only on an ad hoc basis.

Combinatorial auctions are multiple-item auctions for which bids may be placed on packages of items, and are often referred to as auctions with *package bidding*. Neither term should be used to refer to multiple-item auctions in general, because many important types of multiple-item auctions need not accept such bids. Indeed, the SAA first adopted by the FCC provides an example of a format that is multiunit but not combinatorial.

Combinatorial auctions may provide more beneficial outcomes when the value of an item received by a bidder is determined heavily by the other items received. Many applications have been proposed, including electricity markets, equities trading, bandwidth auctions, transportation exchanges, pollution right auctions, auctions for airport landing slots, supply chains, and auctions for carrier-of-last-resort responsibilities for universal services. As mentioned above, innovative vendors already host a significant amount of B2B commerce using combinatorial auctions, making the study and development of more specialized combinatorial auctions a lucrative pursuit, despite the computational difficulties to bidders and auctioneers alike.

Certain general developments in the study of multiunit auctions can be attributed to economic reasoning similar to the results surrounding single-item auctions. The behavior of models that extend these ideas to the case of distinct items is not so well understood. Much of the analysis attending multi-item, multiround auctions rely heavily on assumptions of substitutability among the items being auctioned. Complementarity of items is a widespread phenomenon in markets for which multiunit auctions are proposed, and the existence of superadditive prices (among complements) is a strong reason for expecting that multiunit auctions might produce higher revenues than separate single-unit auctions. These strong substitutability assumptions should therefore be treated with some suspicion.

It can be shown that the "Walrasian" approach may not converge without these very special utility functions displaying the "gross substitutes property" (Scarf 1960). The reason for this difficulty becomes clear when we consider that the problem of finding a revenue-maximizing allocation of items in the general combinatorial auction is $\mathcal{NP}$-hard. The substitutability conditions on multiunit auctions in the economics literature have been proven to be equivalent to a submodular restriction on each bidder's indirect utility function (Ausubel and Milgrom 2002), assuring that various iterative algorithms will quickly converge to global solutions. This suggests that the difficulty of the general combinatorial auction problem is the optimization of a nonconvex objective function, with complementary packages causing local peaks in utility.

The $\mathcal{NP}$-hardness of the winner-determination problem opens the door for an algorithmic approach to combinatorial auctions. As with any $\mathcal{NP}$-hard problem, few have hope of finding a general procedure that can be guaranteed to work rapidly for all situations. Instead, computer scientists and MS researchers focus on identifying situations for which winner determination is easy (as in Rothkopf et al. 1998), specializing in solution techniques to various market structures (as in Günlük et al. 2005, presented in this issue), or developing approximation algorithms for the general case (Zurel and Nisan 2001, Anandalingam et al. 2002). The computationally minded literature often seeks to take advantage of special structure and special techniques to design a multiunit auction tailored to a specific market.

Table 2 Auction Design Components

Component	Objectives	Typical choices
Winner determination	Maximization of revenue or efficiency	Provisional winners with stopping rules Solve an IP formulation
Payment determination	Incentive compatibility Revenue maximization	Pay-as-bid Uniform-price Vickrey-Clarke-Groves
Information flow	Privacy preservation Minimize cost of elicitation	Sealed-bid price offers Dynamic price w/demand reporting
Bid language	Minimize the exponential Bundles problem	All-or-nothing package Bids XOR-of-OR

To better understand the developments that are being made, we introduce the categories of auction design components, tabulated in Table 2. Although every auction design must define rules for each design component, in many auctions the choices may not be unusual. For example, determining the winner and payment in a first-price, sealed-bid, single-item auction is a trivial matter, and the information flow and bid language are as simple as possible. Still, it is precisely the variations on these four auction attributes that fuel the growing literature on the design of multiunit distinct-item auctions. As such, Table 2 provides an outline for the remainder of this section.

4.3.1. Winner Determination. Variations on the winner-determination problem are the most common innovations in the algorithmic approach, using different IP formulations and different algorithms for solving these various formulations. This is both because the computational complexity lies in winner determination, and because most innovations in the other components dictate at least some adjustment of the winner-determination process. Within this issue, Günlük et al. (2005) and Sandholm et al. (2005) demonstrate significant advances in the development of new computational methods tailored to solving the winner-determination problem of a specific auction. Other work on efficient algorithms for the winner-determination problem is provided by Kwasnica et al. (2005 in this issue), Anandalingam et al. (2002), and Fujishima et al. (1999). Additionally, de Vries et al. (2003) explore the connections between auction winner determination and the theory of optimization, recognizing several ascending auction formats as either primal-dual or subgradient optimization.

Winner determination is impeded not only by the $\mathcal{NP}$-hardness of the general problem, but also by exponential growth of input relative to the number of items N. This may be referred to as the *exponential bundles problem*: There are $2^N - 1$ nontrivial packages in an auction of N items, too many for a human to consider explicitly for an auction of more than a few items. Because the general winner-determination problem (see de Vries and Vohra 2003) allows too many subsets to be specified for even a modestly large number of items, some methods suggested in the literature employ an auction in which bids are only accepted on a certain set of permissible packages of items, necessarily much smaller than the set of all subsets of the N items. Rothkopf et al. (1998) provide several such *restricted-subset* combinatorial auctions for which polynomial-time algorithms exist. Although these methods work well for very specific market structures, there are two problems limiting their applicability.

First, the auctioneer decides in advance which packages may be bid on, which may be restrictive when the bidders are firms competing for raw materials or governmental licenses that can be utilized in different combinations according to varying technologies and market strategies. In this case, a mechanism in which the bidders themselves choose the packages may be more desirable. Secondly, if a bidder does not submit a bid on a particular package, how much should she pay for that package if it is awarded to her? If the auction dictates that this package may not be awarded to her, the mechanism may sacrifice efficiency; it considers a package of nonzero value to have zero value and may thus miss an optimal allocation. If instead the mechanism relaxes the single-bundle-per-bidder constraints (i.e., a bidder may receive a bundle together with other bundles at an additive cost), then it negates expression of substitutability among packages. One way to obviate this problem is to express package substitutability through the use of dummy items (Nisan 2000), although this approach has limitations.

A general alternative to the restricted-subset methods is given by the *restricted-preference* methods, mentioned briefly by de Vries and Vohra (2003). These formats place limitations on what kind of bidding functions may be used, based on assumption or inference on the behavior and preferences of bidders. Rather than reducing the number of bundles that can be bid on in any arbitrary fashion, these methods place limitations on the relationship among the values assigned to various bundles. For example, under most reasonable circumstances it is safe to assume a *nondecreasing* preference restriction: $v_j(S \cup \{i\}) \geq v_j(S)$. In other words, if a bidder is given one more item the value of her bundle does not decrease, and may be forced to bid accordingly. This particular objective function restriction is not strong enough to alleviate much of the computational burden, and the few restrictions explored in the literature that are strong enough to positively affect computations typically do not allow for the expression of complementary bundles (see de Vries and Vohra 2003). Examples of these preference restrictions (without complementarities) include the gross substitutes property, buyer submodularity, and the agents-are-substitutes condition (see, for example, Kelso and Crawford 1982, Ausubel and Milgrom 2002, Bikhchandani and Ostroy 2002, respectively). A preference restriction allowing for compact representation as well as expression of both substitutabilities and complementarities is described by Day and Raghavan (2004b).

4.3.2. Payment Determination. Unlike winner determination, methods for payment determination vary little over the current frontier of multi-item auction design. This is because of the widespread acceptance of the class of VCG mechanisms for honesty-inducing payment determination (for primary sources, see Clarke 1989, Groves 1973, Vickrey 1961). In this issue, however, Chen et al. (2005) question the subtleties of applying this paradigm, and show that some IP formulations of winner determination can lead to inflated payments when using a VCG mechanism naively. Indeed, despite its theoretical beauty, several authors expose drastic problems with VCG payment mechanisms, explaining their scarcity of implementation (see Ausubel and Milgrom 2002, Rothkopf and Harstad 1995, Rothkopf et al. 1990, Sakurai et al. 2000). Among these problems for VCG mechanisms are the vulnerability to false-name bidding, collusive interference, bid-taker cheating, and failure of the payments to support a core outcome, further questioning the widespread acceptance of the VCG prices.

Given the limitations of the VCG payment mechanism, a few alternative payment mechanisms for combinatorial auctions have been proposed that offer a compromise between the pay-as-bid and VCG extremes. iBundle (Parkes 2001) and the ascending proxy auction (Ausubel and Milgrom 2002), for example, both give an alternative to VCG pricing with methods that achieve "bidder-Pareto-optimal payments" within the "core" under certain convexity conditions. Additionally, Parkes et al. (2001) explore how prices may be achieved that approximate VCG payments as closely as possible, preserving some portion of the incentive compatibility while maintaining budget-balance constraints in a combinatorial double auction. An analogous procedure for one-sided auctions is provided by Day and Raghavan (2004a).

Despite their apparent drawbacks, VCG mechanisms maintain several interesting properties as a means of payment determination and greatly simplify the analysis of computational methods via the well-known *revelation principle* (Myerson 1981). Noteworthy is the connection between linear programming dual prices and the calculation of VCG payments (see Bikhchandani and Ostroy 2002). In particular, with the "proper formulation," the VCG payments correspond exactly to the dual variables of the LP relaxation of the winner-determination problem.

4.3.3. Information Flow. In addition to their potential role in payment determination, supporting dual prices may also play a part in the information flow structure of some iterative (i.e., multiround) multiunit auctions, serving as feedback to inform bidders how to proceed in the next round (see Günlük et al. 2005). Similar feedback prices are determined in each round by solving optimization problems in the new formats of both Kwasnica et al. (2005) and Kwon et al. (2005), both presented in this issue. Both of these works represent innovative new designs in the information flow of auctions, although they have conflicting ideas of how updated price information should be captured and presented (i.e., the nature of the information flow).

As a general note regarding information flow, there is a central dichotomy among the different approaches to information flow, dividing the set of all auctions into two types. In one type bidders are asked to submit demand functions (or demand correspondences, more generally), either in response to some current price vector or for a larger set of prices. Single-item English auctions are the simplest auctions of this type, where each demand function reported at the current price is simply willingness to buy or not. The second general information-flow type contains auctions in which the submitted information consists of price offers for various bundles. Although the sealed-bid auctions do fall into this category, use of this type of information transmission is also common in dynamic multiunit auction design.

Although the relationship between these two types of submission is readily seen to be inverse (bundles assigned to prices as opposed to prices assigned to bundles), the contexts can be quite different and may not exhibit the same dynamic behavior. A similar interpretation in the language of optimization is that the two approaches are dual. If a primal model's decision variables tell whether a given item is chosen, then certain dual variables may be interpreted as prices for items.

In general, an auction mechanism must direct the flow of information to and from bidders, revealing market information to guide future bidding, and eliciting further bidding information as needed. In a multiunit auction, the auctioneer should recognize the difficulty of formulating bid information (cost of elicitation), requesting just enough information to achieve the desired auction outcome. Other considerations include privacy preservation (mimicking the English auction, which does not reveal the valuation of the winning bidder) and anticollusive design (restricting the release of information that may allow bidders to collude).

4.3.4. Bid Languages. This duality of auction formats draws the first major distinction in the classification of *bid languages*. In single-item auctions it seems to be the only distinction: Do bidders submit demand at a given price (the language of indirect mechanisms) or a reservation price above which demand is zero (the language of direct mechanisms)? Models of multiunit auctions usually assume a straightforward generalization of one of these two languages, but there are reasons for considering alternative bid languages. Because of the difficulty of expression imposed by the exponential bundles problem, various indirect and direct mechanisms for preference revelation are studied in the multi-item auction literature. Many of these approaches build bid expressions through the use of exclusive package bids or *flat bids*.

Definition. A *flat bid* (S, p) is a nonnegative bid p on a set of items S, with no bid on any other bundles.

The direct language considered in the most general combinatorial auction settings involves bidders submitting a reservation price for every single bundle. In this context there is no ambiguity to allow for an exposure problem, but the exponential bundles will be experienced in full force for bidders using this "exclusive or" (*XOR*) of flat bids language, even with only a modest number of items. In general, the problem of bid language design for combinatorial auctions is to mitigate the exponential bundles problem with an expressive system for bid submission. In this context a single bid "sentence" may simultaneously place bids on multiple bundles. Many approaches use flat bids joined by logical connectives, usually *OR* and *XOR*. For example, one could bid ($300 on $\{A, B\}$) *XOR* ($400 on $\{C, D\}$), using the *XOR* to keep from getting both bundles. Nisan (2000) describes the strengths and weaknesses of such languages in detail, while Günlük et al. (2005) elaborate on winner determination for a nested *XOR*-of-*OR* language, as proposed by the FCC.

Determining the appropriate communication format (bid language) for a particular auction may be very application specific. In practice, a language that allows market participants to more easily and effectively express their preferences will facilitate better market outcomes, making the study of bid languages a potentially rewarding avenue of future research.

5. Equilibria and Behavior in Electronic Markets

Much of the multiple-item auction literature focuses on the computational issues of winner/payment determination, mechanism design, and information flow, as outlined above. This is most likely due to the influx of CS and MS scholars since the interesting computational problems became widely known in the early 1990s. Still, the roots of auction theory are well entrenched in the economic field of game theory. Auctions are games of incomplete information, and theoretical studies explain what competitive behavior may be expected in equilibrium.

As noted above, there is no general consensus on how to confront the computational difficulties presented by multiple-item auctions, and it seems that several auction formats will prove useful for implementation in the varying landscape of applications. It should come as no surprise that game theory has had only a limited ability to accurately describe the behavior of the multitude of formats, because each auction presents a unique game for analysis, each presenting a complex equilibrium problem.

One auction that has received special attention is the SAA, due in part to its use in FCC spectrum auctions. In this issue, Engelbrecht-Wiggans and Kahn (2005) investigate the equilibrium properties of this auction, providing a model that explains the observed existence of low-revenue outcomes. Their results concern the need for eligibility rules and careful reserve pricing in order to reduce bidders' ability and incentive to collude. They describe equilibrium behavior in which the small number of bidders are able to anticipate the auction outcome and decide not to compete. In this way they are able to "divide up the pie" without driving up the prices on one another.

Another set of detailed data illustrating collusive bidding in the SAA is given by Cramton and Schwartz (2000). Their analysis focuses on signaling among competitors as a form of tactical collusion to achieve lower payments at the termination of the auction. While competing in simultaneous auctions, bidders are seen to play punishments strategies: Bidder 1 submits a high bid on item *A* (which she is not interested in) in order to drive up the price for Bidder 2, the likely winner of *A*, as punishment for submitting a competitive bid on item *B*, which Bidder 1 will likely win. To be sure that Bidder 2 gets the message, Bidder 1 abuses the precision of expression afforded to the bidders, encoding her initials or an indication of the market for item *B* in the smaller digits of the punishment bid. A bidder employing this strategy may find it advantageous to "name-tag" all her bids with a code in the smallest digits, so that other bidders can identify her and "stay off her turf." Newer rules that protect against this from of signaling have since been implemented in the FCC context, but a bidder's ability to signal her preferences to others in order to achieve a more desirable outcome remains a general concern in multiround auction design.

Although a convincing argument is made for the use of anticollusive measures in multiple-item auctions in general, these principles have scarcely been applied to the combinatorial auctions setting. It seems reasonable to expect that much of the intuition from the simultaneous ascending auction does carry over into the realm of package bidding: Using information revealed over the course of the auction, bidders may anticipate the final allocation and curtail competition to achieve the outcome at low prices.

Can collusion disrupt the outcome in other (nonauction) electronic market venues? The work presented here by Campbell et al. (2005) demonstrates that the proposed consumer benefits of electronic markets (specifically, reduced or negligible search costs) may be nullified by an increased opportunity for collusion among eMarket vendors. Their analysis emphasizes the disparity between static (one-shot) and dynamic (infinite time horizon) models of equilibrium behavior in eMarkets, showing in the latter case that competitors can enforce a collusive equilibrium with punishment strategies for those who deviate from the agreed upon high prices, typical of collusion models. Interestingly, this model shows that the firms can maintain these high prices through coordination, even in an environment of highly imperfect monitoring. Because typical B2C eMarkets allow companies to monitor a competitor's prices as costlessly as they are monitored by consumers, we should expect the collusive opportunities to increase from the imperfect monitoring case, as search costs decline in the presence of "e-hubs" and automated Internet search agents. Freely available price information allows firms to enforce the coordinated (collusive) equilibrium prices more easily. The result is that the reduction in rents charged by firms in the face of high search costs are offset by a shift towards a more perfect form of price collusion. This research suggests the need for a new line of antitrust policies (both detection and enforcement) specifically geared to the eMarkets environment.

Similar effects of search costs in eMarkets are provided here by Terwiesch et al. (2005), who explore the behavior of buyers and sellers in Name-Your-Own-Price (NYOP) Internet venues. Although exact equilibria in the NYOP setting are not derived, their model introduces fundamental new understanding of the "bargaining" approach to electronic commerce. Specifically, a model is provided under which NYOP retailers may expect higher profits than in a fixed-price setting, and the optimal-price threshold problem of the retailer is solved exactly. Similarly, the "consumer haggling problem" is modeled, providing the optimal strategy for a consumer based on her valuation and cost of search.

Lest the reader believe that search costs provide the only differentiating feature of eMarkets relative to traditional (nonelectronic) channels of commerce, Viswanathan (2005) investigates (in this issue) the effects of channel flexibility, network externalities, and switching costs in a multichannel model of commerce (i.e., one including traditional, electronic, and hybrid firms). For example, the ability of firms to offer a variety of products with a variety of features over the Internet (channel flexibility) may reduce firms' ability to differentiate themselves over this channel, thus increasing competition and reducing profits for such firms. Viswanathan's model shows an interesting effect that hybrid firms may dampen this effect at equilibrium, reducing competition in the flexible channel (the electronic market) while increasing competition in the less flexible channel (the traditional market).

6. Conclusions

Economists tend to provide models with attractive equilibrium properties and recognizable strategic behaviors, often maintaining limiting assumptions to take the edge off of the computational difficulties. Computationally minded researchers, on the other hand, assume very simple player strategies to assure nice equilibrium behavior, allowing them to focus on the complex bidding and decision-making environments. As such, we assert that the continued progress of eMarket design depends heavily on a unprecedented multidisciplinary collaboration, the inauguration of a new field of research with roots in economics, CS, and MS. We hope that within this special issue you find the seeds for such a collaborative effort, and a common language emerging where once there were several.

Within this issue you will find research on the cutting edge of eMarkets design and analysis. The majority of the material is focused on the study of auctions, reflecting its central standing in the papers submitted for this volume. We hope to have provided the reader with a concise guide to the literature in this ever-growing field of interdisciplinary research. In addition to the auction focused articles contained here, several of the articles provide fresh insight into the unique economics of electronic commerce, its peculiar dynamics, and the opportunities (and dangers) presented by state-of-the-art market mechanisms. We expect that each of these offerings will motivate future research and, perhaps more importantly, educate and influence the public on how business is done electronically.

Acknowledgments

The authors gratefully acknowledge the support of the National Science Foundation, award number DMI-0205489.

References

Anandalingam, G., R. Kwon, L. Ungar. 2002. An efficient approximation algorithm for combinatorial auctions. Working paper, Systems Engineering, University of Pennsylvania, Philadelphia, PA.

Ausubel, L., P. Cramton. 2002. Demand reduction and inefficiency in multi-unit auctions. Working paper, University of Maryland, College Park, MD.

Ausubel, L., P. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers Theoret. Econom.* **1**(1), http://www.bepress.com/bejtefrontiers/vol1/iss1/art1.

Ausubel, L., P. Cramton, P. Milgrom. 2005. The clock-proxy auction: A practical combinatorial auction design. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chapter 5. MIT Press, Cambridge, MA, Forthcoming.

Ball, M., G. Donohue, K. Hoffman. 2005. Auctions for the safe, efficient and equitable allocation of national airspace systems resources. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chapter 22. MIT Press, Cambridge, MA, Forthcoming.

Banks, J., J. Ledyard, D. Porter. 1989. Allocating uncertain and unresponsive resources: An experimental approach. *RAND J. Econom.* **20**(1) 1–25.

Bikhchandani, S., J. Ostroy. 2002. Ascending price Vickrey auctions. *Games Econom. Behavior*, http://dx.doi.org/10.1016/50899-8256(03)00141-6.

Campbell, C., G. Ray, W. Muhanna. 2005. Search and collusion in electronic markets. *Management Sci.* **51**(3) 497–507.

Carare, O., M. Rothkopf. 2005. Slow Dutch auctions. *Management Sci.* **51**(3) 365–373.

Chen, R., G. Janakiraman, R. Roundy, R. Q. Zhang. 2005. Efficient auction mechanisms for supply chain procurement. *Management Sci.* **51**(3) 467–482.

Clarke, E. 1989. Multipart pricing of public goods. *Public Choice* **8** 19–33.

Cramton, P. 2003. Electricity market design: The good, the bad, and the ugly. *Proc. 36th Hawaii Internat. Conf. System Sci.*, IEEE Press, Piscataway, NJ.

Cramton, P., J. Schwartz. 2000. Collusive bidding: Lessons from the fcc spectrum auctions. *J. Regulatory Econom.* **17** 229–252.

Day, R., S. Raghavan. 2004a. Generation and selection of core outcomes in sealed-bid combinatorial auctions. Working paper, University of Connecticut, Storrs, CT.

Day, R., S. Raghavan. 2004b. Matrix bidding in combinatorial auctions. Working paper, University of Connecticut, Storrs, CT.

Deltas, G., R. Engelbrecht-Wiggans. 2005. Naive bidding. *Management Sci.* **51**(3) 328–338.

de Vries, S., R. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS J. Comput.* **15**(3) 284–309.

de Vries, S., J. Schummer, R. Vohra. 2003. On ascending Vickrey auctions for heterogeneous objects. Working paper, Northwestern University, Evanston, IL.

Ding, M., J. Eliashberg, J. Huber, R. Saini. 2005. Emotional bidders—An analytical and experimental examination of consumers' behavior in priceline-like reverse auctions. *Management Sci.* **51**(3) 352–364.

DotEcon-Ltd. 2001. Auctioning airport slots: A report for HM treasury, and the department of the environment, transport and the regions. Technical report, www.dotecon.com.

Elmaghraby, W., P. Keskinocak. 2003. Combinatorial auctions in procurement. C. Billington, T. Harrison, H. Lee, J. Neale, eds. *The Practice of Supply Chain Management*. Kluwer Academic Publishers, Norwell, MA.

Engelbrecht-Wiggans, R., C. Kahn. 2005. Low-revenue equilibria in simultaneous ascending-bid auctions. *Management Sci.* **51**(3) 508–518.

Friedman, L. 1956. A competitive bidding strategy. *Oper. Res.* **4** 104–112.

Fujishima, Y., K. Leyton-Brown, Y. Shoham. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. *16th Internat. Joint Conf. Artificial Intelligence* (*IJCAI*), Stockholm, Sweden.

Groves, T. 1973. Incentives in teams. *Econometrica* **41** 617–631.

Günlük, O., L. Ladányi, S. de Vries. 2005. A branch-and-price algorithm and new test problems for spectrum auctions. *Management Sci.* **51**(3) 391–406.

Hohner, G., J. Rich, E. Ng, G. Reid, A. Davenport, J. Kalagnanam, H. S. Lee, C. An. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces* **33**(1) 23–35.

Kelly, F., R. Steinberg. 2000. A combinatorial auction with multiple winners for universal service. *Management Sci.* **46**(4) 586–596.

Kelso, A., V. Crawford. 1982. Job matching, coalition formation and gross substitutes. *Econometrica* **50** 1483–1504.

Kleindorfer, P., L. Li. 2005. Multi-period, VaR-constrained portfolio optimization in electric power. *Energy J.* **26**(1) 1–26.

Krishna, V. 2002. *Auction Theory.* Academic Press, San Diego, CA.

Kwasnica, A., J. Ledyard, D. Porter, C. DeMartini. 2005. A new and improved design for multi-object iterative auctions. *Management Sci.* **51**(3) 419–434.

Kwon, R., G. Anandalingam, L. H. Ungar. 2005. Iterative combinatorial auctions with bidder-determined combinations. *Management Sci.* **51**(3) 407–418.

Ledyard, J., M. Olson, D. Porter, J. Swanson, D. Torma. 2002. The first use of a combined value auction for transportation services. *Interfaces* **32**(5) 4–12.

Levine, D. 2003. Industrial procurement auctions with expressive competition. CombineNet, Inc. Presentation, INFORMS, Atlanta, GA.

Logistics.com. 2001. Logistics.com named to Forbes Magazine's "best of the Web: B2b" list for second straight year. Press release, Logistics.com.

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *Management Sci.* **89**(5) 1063–1080.

McAfee, R. 1993. Auction design for personal communications services. PacTel Exhibit in PP Docket No. 93-253.

Milgrom, P., R. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50** 1089–1122.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res.* **6** 58–73.

NEXTOR. 2002. Workshop: National airspace system resource allocation: Economics and equity, Queenstown, MD. Organizers: M. Ball, G. Donohue, K. Hoffman. http://www.nextor.org/workshop2.html.

NEXTOR. 2004. Workshop: Government, the airline industry and the flying public: A new way of doing business, Queenstown, MD. Organizers: M. Ball, G. Donohue, K. Hoffman. http://www.nextor.org/Jun04.html.

Nisan, N. 2000. Bidding and allocation in combinatorial auctions. *Proc. 2nd ACM Conf. Electronic Commerce*, ACM Press, New York, 1–12.

O'Neill, R., P. Sotkiewicz, B. Hobbs, M. Rothkopf, J. W. Stewart. 2005. Efficient market-clearing prices in markets with nonconvexities. *Eur. J. Oper. Res.* **164**(1) 269–285.

Parkes, D. 2001. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. Ph.D. thesis, Computer and Information Science, University of Pennsylvania, Philadelphia, PA.

Parkes, D., J. Kalagnanam, M. Eso. 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges. *17th Internat. Joint Conf. Artificial Intelligence* (*IJCAI'01*), AAAI Press, Menlo Park, CA, 1161–1168.

Rassenti, S., V. Smith, R. Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell J. Econom.* **13**(2) 402–417.

Rothkopf, M., R. Harstad. 1995. Two models of bid-taker cheating in Vickrey auctions. *J. Bus.* **68** 257–267.

Rothkopf, M., A. Pekeč, R. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Sci.* **44** 1131–1147.

Rothkopf, M., T. Teisberg, E. Kahn. 1990. Why are Vickrey auctions rare? *J. Political Econom.* **98** 94–109.

Sakurai, Y., M. Yokoo, S. Matsubara. 2000. An efficient approximate algorithm for winner determination in combinatorial auctions. *Proc. 2nd ACM Conf. Electronic Commerce*, (*EC*-00). ACM Press, New York, 30–37.

Sandholm, T., S. Suri, A. Gilpin, D. Levine. 2005. CABOB: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Sci.* **51**(3) 374–390.

Scarf, H. 1960. Some examples of global instability of the competitive equilibrium. *Internat. Econom. Rev.* **1** 157–180.

Schneur, A. 2004. Emptoris Inc. Presentations, NEXTOR Workshop: Government, the Airline Industry and the Flying Public: A New Way of Doing Business. Queenstown, MD.

Terwiesch, C., S. Savin, I. Hann. 2005. Online haggling at a name-your-own-price retailer: Theory and application. *Management Sci.* **51**(3) 339–351.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16** 8–37.

Viswanathan, S. 2005. Competing across technology-differentiated channels: The impact of network externalities and switching costs. *Management Sci.* **51**(3) 483–496.

Wu, D., P. Kleindorfer. 2005. Competitive options, supply contracting, and electronic markets. *Management Sci.* **51**(3) 452–466.

Zurel, E., N. Nisan 2001. An efficient approximate allocation algorithm for combinatorial auctions. *Proc. 3rd ACM Conf. Electronic Commerce*. ACM Press, New York, 125–136.