

Coordination of Purchasing and Bidding Activities Across Posted Offer and Auction Markets*

Daniel D. Zeng¹, James C. Cox², Moshe Dror³

¹ Department of MIS, The University of Arizona, e-mail: zeng@eller.arizona.edu

² Department of Economics, The University of Arizona, e-mail: jcox@eller.arizona.edu

³ Department of MIS, The University of Arizona, e-mail: mdror@eller.arizona.edu

The date of receipt and acceptance will be inserted by the editor

Abstract In both consumer purchasing and industrial procurement, combinatorial interdependencies among the items to be purchased are commonplace. E-commerce compounds the problem by providing more opportunities for switching suppliers at low costs, but also potentially eases the problem by enabling automated market decision-making systems, commonly referred to as trading agents, to make purchasing decisions in an integrated manner across markets. We are investigating a new approach to deal with the combinatorial interdependency challenges for online markets. This approach relies on existing commercial online market institutions such as posted-offer markets and various online auctions that sell single items. It uses trading agents to coordinate a buyer's purchasing and bidding activities across multiple online markets simultaneously to achieve the best overall procurement effectiveness. This paper presents two sets of models related to this approach. The first set of models formalizes optimal purchasing decisions across posted-offer markets with fixed transaction costs. The second set of models is concerned with the coordination of bidding activities across multiple online auctions.

Key words Multiple Markets–Fixed Transaction Costs–Online Auctions–Equilibrium Bidding

* Research reported in this paper is partially supported by a Proposition 301 Electronic Commerce Grant, University of Arizona Eller College of Business and Public Administration. Earlier and incomplete versions of this paper have appeared in the Proceedings of the Hawai'i International Conference on System Sciences (HICSS-37) and the Proceedings of the Thirteenth Annual Workshop on Information Technology and Systems (WITS '03).

1 Introduction

Industrial procurement constitutes a major component of today's e-commerce and the national economy in general (Woods & Marien 2000). In most industrial procurement settings, a buyer needs to purchase a *bundle* of complementary goods as opposed to individual, unrelated goods (McAfee, McMillian & Whinston 1989, Gaeth, Leven, Chakraborty & Levin 1990). For instance, to assemble a car, an automobile manufacturer needs to first purchase all parts required by the corresponding engineering design. Similarly, a software system integrator needs to acquire licenses of all component software packages before the system integration efforts can be initiated.

Furthermore, many industrial procurement tasks involve complex *combinatorial interdependencies* among the items to be purchased. Such interdependencies are often a result of the presence of multiple design or operation alternatives. Take the example of a software system integrator. Suppose this integrator is charged with the task of developing a customized e-commerce storefront which uses a database system as the backend data repository. In the pre-development procurement phase, the integrator needs to acquire licenses for a database software and a matching Web programming environment. If somehow the decision on which database and Web programming environment should be used is already made, the procurement or licensing task is then concerned with finding the best deal for the single, already decided *bundle*. However, if no such decision is made, the procurement task becomes significantly harder since various choices on database and Web programming environment as well as the compatibility issues between these choices have to be considered. As a result, the decision on whether an item (e.g., a particular database system) should be purchased depends on not only its own price, terms of service contracts, etc., but also other potentially relevant items. Two other widely-cited examples of *combinatorial interdependencies* are: the value of owning a take-off time slot at an airport which depends on whether compatible landing slots can be acquired at other airports (Rassenti, Smith & Bulfin 1982), and licenses for bands of the broadcast spectrum in different geographical areas (Milgrom 2000).

Similar bundling and combinatorial interdependency issues are pervasive in consumer purchasing as well. Often the consumer shops for bundles of complementary goods. Consider, as examples, a digital camera and compatible memory cards, or air tickets, hotel reservations, rental car reservations, and (sometimes) tickets for plays, concerts, and ball games. If choices for these goods exist, then the consumer has to consider the resulting combinatorial interdependency issues.

Traditionally, industrial procurement has been a labor-intensive process (Harding 1990). For each item to be purchased, a member of the procurement staff first identifies potential suppliers from various forms of advertisements, referrals, or prior interactions. Then the staff member initiates contacts with these suppliers to learn more about their products or services and solicits price quotes along with other information such as delivery terms. Procurement or sourcing decisions are typically made after quotes are received and (sometimes) an ensuing negotiation process ends.

The above traditional procurement process often leads to ineffective and costly procurement decisions, especially when bundling and interdependency considerations play a significant role. First, only a small subset of all potential suppliers is identified and included in the procurement process because of the high cost associated with largely manual information search efforts. Second, out of a potentially large number of procurement or sourcing alternatives, only a small portion are explored due to the inherent cognitive limitation of the human procurement decision-maker.

The advent of the Web and various e-commerce technologies including e-procurement promises to revolutionize the way in which business-to-business transactions including procurement are conducted (Hannon 2003). In the emerging electronic marketplace, access to product and supplier information is efficient and cost effective due to online catalogs posted by suppliers. A wide range of catalogs and other related value-added services provided by third-party intermediaries further reduce the information search costs associated with procurement. As a result, information collected on relevant products and suppliers for a given procurement task is expected to be much more comprehensive than that collected manually.

In addition to information access, the Web provides a common platform to carry out many other procurement-related business functions including electronic payment and document and contract management, among others (Laudon & Traver 2002). More significantly, many types of market institutions are directly implemented on the Web and buyers have efficient, simultaneous access to these markets. Two prominent examples of these online institutions are: *posted-offer* markets offered by e-tailers where a seller posts a fixed price for an item and a buyer either takes it (buy the item) or leaves it (not buy the item), and *English auctions* where an item is sold through an ascending-price, real-time auction in which the bidder with the last (i.e., highest) bid buys the auctioned item at a price equal to his or her bid.

Despite all these opportunities enabled by e-commerce, there is evidence that shows that enterprise procurement operations have not yet improved significantly (Pur 1998, Gonsalves 2003). A key challenge is, in an operational sense, how to take advantage of (a) the voluminous amount of product and supplier information available from the Web, and (b) the multiplicity of online markets selling the goods to be purchased. Processing such product and supplier information and making procurement decisions across markets in real time pose serious information and cognitive overload problems to the procurement personnel. For instance, it is almost impossible for a human procurement staff member to actively keep track of dozens of online auction markets and a typically larger number of online posted-offer markets to make the best procurement decisions. When bundling and complex combinatorial interdependencies have to be taken into consideration, procurement information processing and decision-making become even more challenging.

Recent years have seen the rapid development of automated procurement systems that aim to meet the above challenges (Zeng & Nissen 2004). As common in the literature, we call such systems *trading software agents*, *procurement software agents* or simply *agents* (Joshi & Singh 1999, Greenwald & Stone 2001). To

avoid potential confusion concerning terminology, we briefly state our definition of agent in the context of procurement. Firstly, unlike in mainstream economics or business literature, we exclusively reserve the use of “agent” to refer to a computational entity. Secondly, with respect to the intelligent agent and multi-agent systems literature, we adopt an agent definition in a relatively weak sense (Jennings, Sycara & Wooldridge 1998). In our context, an agent is simply any automated strategic decision-making and execution system which satisfies the following set of conditions (with no reference to the level of decision-making sophistication or “intelligence”). (1) Agents operate in a networked environment. (2) Agents receive delegated procurement tasks from their human users. (3) Agents interact with other agents or human participants directly or indirectly through well-defined online economic institutions. (4) Agents automate some or all aspects of procurement-related transactions.

Agents have been demonstrated to have great potential of further reducing information search efforts and costs associated with various types of process-oriented transactions (Sycara & Zeng 1996, Nissen 2001). However, many significant technical issues have yet to be addressed to develop an effective agent-based e-procurement approach that can fully take advantage of the potentials offered by the electronic marketplace. One such issue is the lack of adequate *modeling* and *computational* support for dealing with realistic procurement tasks. In other words, existing agents provide a sound enabling technological infrastructure but do not yet offer adequate decision-making mechanisms for important procurement scenarios common in practice.

Research reported in this paper is aimed to fill in this important gap for several classes of procurement problems involving bundles and combinatorial interdependencies. We intend to develop analytical models and suggest corresponding computational mechanisms which can in turn be implemented as the core reasoning module of a sophisticated procurement agent.

There is a large body of literature in economics, operations research, marketing, and information systems (e.g., (Vickrey 1961, Milgrom & Weber 1982, Cox, Smith & Walker 1988, Beard, R. B. Ekelund, Ford & Saba 2001, Berry 1994)), which contains models and computational methods applicable to procurement problems involving *single items* sold through different types of market institutions including posted-offer markets and auctions. Our work, on the other hand, focuses on challenges arising from bundling and combinatorial interdependency considerations. Compared with the existing literature on combinatorial interdependencies and related market designs (e.g., (McAfee et al. 1989, Milgrom 2000, Sandholm 2002, Jones & Koehler 2002, Rothkopf 1994)), our work makes different assumptions regarding the underlying marketplace where procurement operations are conducted. The existing literature typically assumes that the seller offers multiple bundles of goods and can design and enforce customized market exchange rules such as various types of combinatorial auctions through which potential buyers interact and transact. We argue, however, that such customized markets exist only for highly specialized niche items (e.g., airport slots, electric power grids, and communication bandwidths). In the foreseeable future, it is unclear how these market institutions will be accepted in general procurement practice. Thus,

instead of focusing on mechanism design, we aim to develop models to guide and coordinate purchasing and bidding activities across multiple *existing* online markets that sell single items to satisfy bundled and combinatorial procurement needs.

The rest of the paper is structured as follows. Section 2 presents a set of models motivated to coordinate purchasing activities across multiple posted-offer markets. We observe that making optimal purchasing decisions with fixed transaction costs are \mathcal{NP} -hard in the strong sense and suggest using the efficient computational methods developed in discrete location theory to make these procurement decisions. Section 3 presents models that can be used to coordinate bidding activities across two online auctions. We focus on two scenarios: (a) two simultaneous first-price sealed-bid auctions and (b) two simultaneous second-price sealed-bid auctions. For both scenarios, we derive the optimal procurement policies. Section 4 briefly discusses related work and procurement agent implementation issues. We conclude the paper in Section 5 by summarizing our research results and pointing out future research directions.

2 Multiple Posted-Offer Markets

In this section, we study procurement in scenarios where the items of interest are sold through multiple posted-offer markets. Such markets are prevalent in e-commerce: many manufacturers or service providers sell their products or services through their own Web sites under a published price schedule; e-tailers also sell an assortment of products online in a similar manner.

Throughout the paper, we assume that bundling requirements or combinatorial interdependencies exist among the items to be purchased. Furthermore, we assume that sellers differ only in one dimension, i.e., *price*, and ignore their differences in other areas such as delivery time and terms, and overall reputation. This allows us to focus on optimization-based formulations with the objective of minimizing the total cost for the given procurement task.

We start this section by formulating the bundle procurement problem under multiple posted-offer markets where each seller charges a fixed transaction fee whenever one or more orders are placed. This model is an abstraction of a common e-commerce practice: many e-tailers offer flat shipping and handling fees either regularly (e.g., `officedepot.com`) or during promotional periods (e.g., `amazon.com`). We then discuss various extensions to the model including how it can be applied to address combinatorial interdependencies and issues with bundled offers. We conclude this section by reporting a computational study conducted to evaluate the economic significance of using the proposed model to make bundle purchasing decisions based on real-world pricing data.

2.1 Bundle Procurement with Fixed Transaction Costs

We study the following procurement problem. The procurement request is a bundle consisting of n items to be purchased. We denote this bundle by set $O = \{1, 2, \dots, n\}$. A set of m sellers (e.g., e-tailers), denoted by $V = \{1, 2, \dots, m\}$,

has been identified as candidate suppliers. Each seller sells at least one of the items in bundle O . For each item i , seller j publishes selling price P_{ij} . (Without loss of generality, if a seller does not sell a particular item, we set the corresponding price to a sufficiently large positive number.) Furthermore, for seller j , if one or more items are ordered, it charges the buyer a fixed transaction fee denoted by S_j , irrespective of the number of items ordered.

The above procurement problem with the objective of minimizing the total procurement cost (item costs plus applicable fixed transaction fees) can be formulated as an integer program. We first introduce the decision variables x_{ij} for $i \in O$ and $j \in V$, all of which are binary. Let $x_{ij} = 1$, if seller j is chosen for item i , and $x_{ij} = 0$ otherwise. We also introduce a set of auxiliary variables y_j for $j \in V$. Let $y_j = 1$, if seller j is chosen for at least one item, and $y_j = 0$ otherwise. Denote by M a sufficiently large constant. We now present the integer program.

$$z_1 = \min \sum_{i \in O} \sum_{j \in V} P_{ij} x_{ij} + \sum_{j \in V} S_j y_j \quad (1)$$

subject to:

$$\sum_{j \in V} x_{ij} = 1 \quad \forall i \in O \quad (2)$$

$$M y_j \geq \sum_{i \in O} x_{ij} \quad \forall j \in V \quad (3)$$

$$x_{ij} = 0, 1 \quad \forall i \in O, \forall j \in V \quad (4)$$

$$y_j = 0, 1 \quad \forall j \in V \quad (5)$$

The objective (1) formalizes the goal of minimizing the total procurement cost, the sum of item purchase costs and the corresponding transaction costs. The constraints (2) guarantee that all individual items in bundle O are ordered from some sellers. The constraints (3) assure that $y_j = 1$ if at least one item is ordered from seller j using the standard modeling technique for If-Then constraints.

We now state and prove the main result regarding the complexity of solving the above integer program.

Theorem 1 *The problem of finding the minimum cost procurement plans as defined by (1)-(5) is \mathcal{NP} -hard in the strong sense.*

Proof. The integer program defined by (1)-(5) is equivalent to the mixed integer program formulation of the uncapacitated facility location (UFL) problem (Mirchandani & Francis 1990). The \mathcal{NP} -hardness proof for UFL is based on polynomial reduction from the vertex cover problem, a well-known \mathcal{NP} -hard problem ((Mirchandani & Francis 1990), Theorem 3.1). \square

Theorem 1 suggests that it is unlikely that a polynomial or even pseudo-polynomial algorithm exists that can optimally solve the above procurement problem. Practically, this means that finding an optimal procurement plan is impossible when the

number of items and the number of potential sellers are modestly large, especially when procurement decisions have to be made in a relatively short time frame. Fortunately, there exist several classes of polynomial time heuristic methods developed in the discrete location theory literature that can be directly brought to bear upon this computational challenge (Mirchandani & Francis 1990, Vazirani 2001, Jain, Mahdian, Markakis, Saberi & Vazirani 2003). Although these methods do not guarantee optimal solutions, they are capable of producing high-quality solutions and some of them have guaranteed error bounds.

Note that if none of the sellers charge fixed transaction costs, i.e., $S_j = 0$ for all $j \in V$, the procurement problem defined by (1)-(5) has an obvious polynomial solution: for each item, buy it from the seller who offers the lowest selling price for that item. Intuitively, the absence of fixed transaction costs decomposes the bundle procurement problem into a collection of independent single-item buying problems which are easy to solve.

2.2 Model Extensions

We now extend the model developed in the previous section to capture procurement decisions in more complex procurement scenarios. In the first scenario, a buyer needs to buy a bundle of goods. For each constituent item, however, the buyer now can choose between different brands. These brands are all functionally equivalent but are sold under different prices from different sellers.

In the second scenario, we consider full-fledged combinatorial interdependencies among the items to be purchased. In this case, the procurement requirements are not specified as a fixed bundle. Rather, they are given as multiple alternative bundles which lead to different utilities. The buyer has to find the bundles that maximize the difference between their utility and purchasing cost.

In the third scenario, we model another common e-commerce pricing practice: many online posted-offer markets waive the shipping and handling charge if the size of the order placed exceeds a predetermined threshold. When this threshold is not met, the buyer needs to pay the fixed per-shipment transaction fee.

In the fourth and last scenario, we are concerned with an emerging trend of e-commerce practice in which sellers offer bundled goods. The best example of such practice in the retail setting is `amazon.com`, which offers book bundles on the fly to potential customers when they search and browse its online book catalog.

2.2.1 Bundle Procurement with Multiple Brands The notation used to formalize the bundle procurement problem with multiple brands is based on that used in Section 2.1. Below we describe the notational differences.

We assume that for item $i \in O$, there is a nonempty set A_i of brands from which the buyer can choose. We use a new subscript k to indicate these brands. The decision variables are now x_{ijk} indicating whether seller j is chosen for brand k of item i , and the problem input parameters regarding item prices are P_{ijk} indicating the price seller j offers on brand k of item i . The problem of minimizing the total procurement cost with multiple brands can then be formulated as follows.

$$z_2 = \min \sum_{i \in O} \sum_{j \in V} \sum_{k \in A_i} P_{ijk} x_{ijk} + \sum_{j \in V} S_j y_j \quad (6)$$

subject to:

$$\sum_{j \in V} \sum_{k \in A_i} x_{ijk} = 1 \quad \forall i \in O \quad (7)$$

$$M y_j \geq \sum_{i \in O} \sum_{k \in A_i} x_{ijk} \quad \forall j \in V \quad (8)$$

$$x_{ijk} = 0, 1 \quad \forall i \in O, \forall j \in V, \forall k \in A_i \quad (9)$$

$$y_j = 0, 1 \quad \forall j \in V \quad (10)$$

The objective (6) indicates the goal of minimizing the total procurement cost, which is the sum of item purchase costs and applicable fixed transaction costs. The constraints (7) guarantee that exactly one brand for each individual item in bundle O is ordered from some seller. The constraints (8) assure that $y_j = 1$ if at least one item is ordered from seller j .

The integer program defined by (6)-(10) is a generalization of the model from Section 2.1. Observe that if a seller is chosen to provide for a brand of item i , this brand has to be the cheapest brand offered by this seller for item i . Therefore, the above program can be reduced to the one studied in the previous section, thus making the heuristic methods from discrete location theory applicable.

2.2.2 Procurement with Combinatorial Interdependencies We first formalize the notion of combinatorial interdependencies. Consider the set of n distinct items that a buyer is interested in purchasing $O = \{1, \dots, n\}$. Denote the power set of O by 2^O , consisting of all subsets of O . In general, any combinatorial interdependencies can be fully captured by a utility function in the form of $u : 2^O \rightarrow R^+$, where R^+ represents the set of nonnegative real numbers.

In practice, however, it is difficult to fully specify utility function u due to its size. The following alternative three-step approach can be taken instead.

- Step 1. The (risk-neutral) buyer provides a set F of bundles of potential interest. For each bundle, the buyer specifies its utility. (Such information constitutes a partial definition of u ; we assume that $u(\cdot) = 0$ for all other elements of 2^O (that is, the elements in set $2^O \setminus F$)).
- Step 2. For each bundle in F , the buyer applies the models developed in the previous two sections to identify the lowest procurement cost.
- Step 3. The buyer selects the bundle(s) with maximum difference between its utility and its lowest procurement cost.

Note that the efficiency of this enumeration approach largely depends on how quickly a high-quality procurement plan can be found for a given bundle (as studied in the previous sections). Using fast heuristics to solve these underlying pro-

curement problems, we envision that dealing with combinatorial interdependencies does not pose new significant computational challenges. The key issue is how to conveniently elicit utility function u in practical settings.

2.2.3 Conditional Free Shipping We assume that Seller j waives S_j if the order amount exceeds T_j . We introduce a set of auxiliary variables w_j for $j \in V$. Let $w_j = 1$ if the order amount with seller j exceeds T_j and $w_j = 0$ otherwise. The following integer program formalizes the bundle procurement problem under this type of conditional free shipping.

$$z_3 = \min \sum_{i \in O} \sum_{j \in V} P_{ij} x_{ij} + \sum_{j \in V} S_j (y_j - w_j) \quad (11)$$

subject to

$$\sum_{j \in V} x_{ij} = 1 \quad \forall i \in O \quad (12)$$

$$M y_j \geq \sum_{i \in O} x_{ij} \quad \forall j \in V \quad (13)$$

$$M(1 - w_j) > T_j - \sum_{i \in O} x_{ij} P_{ij} \quad \forall j \in V \quad (14)$$

$$x_{ij} = 0, 1 \quad \forall i \in O, \forall j \in V \quad (15)$$

$$y_j = 0, 1 \quad \forall j \in V \quad (16)$$

$$w_j = 0, 1 \quad \forall j \in V \quad (17)$$

The objective (11) formalizes the goal of minimizing the total procurement cost. When the order amount exceeds threshold T_j , $y_j - w_j = 0$, indicating qualified free shipping. The constraints (14) ensure that w_j is set to zero if the threshold T_j is not met. Computationally, the integer program defined by (11)-(17) is significantly harder to solve than the base model presented in Section 2.1. We are currently developing and experimenting with various types of computational methods to solve this problem. The related computational study is beyond the scope of this paper.

2.2.4 Bundle Offers We assume that Seller j offers a collection of item bundles. Denote this collection (set) by B_j and its cardinality n_j . Individual bundles are indexed by $k = 1, 2, \dots, n_j$. Selling prices are now associated with bundles as opposed to single items. We use P_{kj} to denote the published selling price for the k^{th} bundle offered by Seller j . We introduce an additional set of parameters a_{ijk} to simplify the notation. Let $a_{ijk} = 1$ if the k^{th} bundle in B_j contains item i , and $a_{ijk} = 0$ otherwise. The decision variables x_{kj} now have the following interpretation: Let $x_{kj} = 1$ if the k^{th} bundle in B_j is bought and $x_{kj} = 0$ otherwise. The following integer program formalizes the bundle procurement problem with sellers offering bundled goods.

$$z_4 = \min \sum_{j \in V} \sum_{1 \leq k \leq n_j} P_{kj} x_{kj} + \sum_{j \in V} S_j y_j \quad (18)$$

subject to

$$\sum_{j \in V} \sum_{1 \leq k \leq n_j} a_{ijk} x_{kj} \geq 1 \quad \forall i \in O \quad (19)$$

$$M y_j \geq \sum_{1 \leq k \leq n_j} x_{kj} \quad \forall j \in V \quad (20)$$

$$x_{kj} = 0, 1 \quad \forall j \in V, 1 \leq k \leq n_j \quad (21)$$

$$y_j = 0, 1 \quad \forall j \in V \quad (22)$$

In this formulation, constraints (19) ensure that each of the items to be purchased is contained in at least one of the purchased bundles. Note that the base model presented in Section 2.1 is a special case of this integer program defined by (18)-(22) through limiting each element of B_j (an offered bundle) to a singleton set containing only one item. We are currently applying several computational methods developed in combinatorial auction research to solve this model (de Vries & Vohra 2003) and will report our computational experience in a future paper.

2.3 A Computational Evaluation

To investigate the extent to which the total procurement costs can be reduced when following the optimal bundle procurement decisions, we have conducted a computational study using real-world pricing data retrieved from various e-tailers and online infomediaries. This section summarizes the setup of this computational study and its main findings.

Two sets of e-commerce data were collected and used in our study. The first set, collected mainly from `bestwebbuys.com` on July 31, 2003, contains per-shipment shipping and handling charges as well as per-item listing prices (including per-item shipping and handling charges when applicable) for used books. These books were selected to cover several major book categories such as business, biographies, fictions, etc. The second set, retrieved from `pricescan.com` on August 9, 2003, contains similar pricing information for a range of computer hardware including keyboards, computer mice, memory modules, among others. The following table summarizes some key statistics characterizing these two datasets. In column labels, “AVE” stands for sample mean (average) and “STD” for sample standard deviation.

Dataset	Total # of items	Total # of sellers	(AVE, STD) for # of sellers per item	(AVE, STD) for price per item	(AVE, STD) for per-shipment cost per seller
Book	2317	8	(2.57, 0.77)	(9.07, 3.49)	(3.33, 0.26)
Hardware	100	60	(10.60, 3.66)	(39.02, 6.67)	(7.99, 4.43)

To demonstrate the impact of adopting the optimal bundle procurement decisions, we use the following common procurement policy as the comparison base: For each item in the bundle to be purchased, buy it from the seller that offers the *lowest total cost* including both the per-item charge and the fixed transaction cost. Obviously, this simple policy can result in suboptimal decisions because it ignores the fact that S_j is charged only once even if more than one item are ordered from seller j . Nevertheless, this policy is intuitive and captures the essence of the strategy used by all existing comparative shopping agents developed mainly for single item purchasing. We refer to this policy of the myopic, heuristic nature as the MYO policy and the optimal bundle procurement policy governed by the model presented in Section 2 as the OPT policy.

We conducted two separate computational experiments using the used book and computer hardware datasets, respectively, to compare the performance of OPT against that of MYO for bundle size ranging from 2 to 20. In each experiment, we first randomly generated 200 bundle requests for each bundle size under study. For instance, for bundle size 2, each request was a pair of randomly selected items (without replacement) from all available items. Then for each bundle request, we computed both MYO and OPT policies and calculated the resulting total procurement costs. In our experiments, the OPT policies were computed using the CPLEX integer program solver. (It took CPLEX about 4 seconds to solve a problem with bundle size 20 in our experiments on a Sun Ultra 5 workstation; the time needed to compute MYO is negligible.) Figures 1(a) and 1(b) summarize the key results of these experiments.

For books, as shown in Figure 1(a), the buyer can save 6.64% for bundle size 2 when using OPT as opposed to MYO. This number is the average of percentage savings over the 200 randomly generated 2-book bundle problem instances. As shown in Figure 1(b), for bundle size 2, the maximum percentage saving achieved by OPT out of these 200 cases is 40.10%. The general trend for the book dataset seems to indicate that as the bundle size increases, the percentage saving enabled by using OPT decreases in both the average and best cases. For instance, with bundle size 20, the average percentage saving drops to 0.50% with the maximum being 0.66%, suggesting that the buyer should adopt the MYO policy since OPT brings almost no value but is much more difficult to compute. The results based on the hardware dataset, however, exhibit completely different behavior. As illustrated in Fig. 1, for bundle size 2, the average percentage saving is merely 1.09%. As the bundle size increases, however, the average percentage saving enabled by using OPT now increases in general. For bundle size 20, the percentage saving reaches 7.67% (a dollar value of \$52.34). It is also interesting to observe from Fig. 2 that the maximum percentage savings for the hardware dataset do not fluctuate as much as those for the book dataset as the bundle size increases.

We are currently analyzing the datasets and related experimental results in depth to account for the above observations. In particular, we are investigating whether approximation algorithms and competitive analysis techniques (Vazirani 2001, Borodin & El-Yaniv 1998) may provide a useful framework for such an analysis and help develop a theory that can predict under which situations and to what extent OPT will (significantly) outperform MYO. Besides these potentially




Fig. 1 Average Cost Savings OPT/MYO




Fig. 2 Maximum Cost Savings OPT/MYO

interesting further developments, our computational study reveals that significant savings are possible in realistic bundle procurement settings. Both paired t -test and Wilcoxon test confirm that the differences between OPT and MYO are statistically significant for all datasets and bundle sizes under study.

3 Multiple Auction Markets

3.1 Motivation and Basic Assumptions

In recent years, online auctions have gained wide acceptance in the electronic marketplace. Both consumer products and industrial goods are routinely traded

through online auction houses in large volumes (Hannon 2003). In particular, auction markets have been rapidly developed in a number of vertical industrial segments such as auto parts (`covisint.com`) and chemicals (`chemical.net`).

Auctions have been extensively studied in the economics and game theory literature (Klemperer 1999). A subfield of auction theory studies combinatorial interdependencies and product complementarities (Sandholm 2002, Jones & Koehler 2002, Rothkopf 1994). A basic assumption made by most combinatorial auction work is that there exists a combinatorial auction market through which a seller and multiple buyers interact. In such an auction, the seller offers a range of product bundles and buyers or bidders bid for them based on their utility functions and their knowledge about other bidders.

Combinatorial auctions have many inherent theoretical appeals. In addition, significant practical lessons concerning auction setup and effectiveness have been learned through their recent applications in areas such as selling radio spectrum rights and trading electricity power (Ausubel, Cramton, McAfee & McMillan 1997, Klemperer 1999).

We project, however, that combinatorial auctions will not have immediate impact on consumer purchasing or industrial procurement based on the following arguments. First, from a technological standpoint, developing and managing general-purpose combinatorial auction markets is significantly harder than it is for auction markets that sell single items due to the complex trading rules associated with combinatorial auctions. It is also unclear how well combinatorial auctions will scale as a market mechanism when the number of goods and good bundles as well as the number of potential buyers grow, despite recent developments in the computational aspects of combinatorial auctions (Sandholm 2002, Jones & Koehler 2002). Second, individual sellers typically do not offer the wide range of goods and bundles for sale to satisfy potential buyers' procurement needs. Third, from a buyer's perspective, participating in a combinatorial auction requires significant computational expertise to make hard bidding decisions (Sandholm 2002).

At the same time, auction markets selling single items have been well developed and offer potentially substantial savings for buyers. In effect, buying single, independent items from various auction markets is becoming a standard practice in industrial procurement. What is lacking is theoretically sound guidelines and rules that can be used to coordinate bidding activities across these multiple auction markets to satisfy procurement requirements with combinatorial interdependencies.

This section presents models developed to study these across-market coordination issues in two specific auction settings. For simplicity, in both settings, we assume that the buyer is interested in two items. Furthermore, we assume that there exists a simple form of interdependency among these two items: some extra positive utility is generated when both items are acquired (in addition to the sum of their individual utilities). In the first setting, each of the items is sold through an independently run first-price sealed-bid auction in which the bidder with the highest bid buys the auctioned item at a price equal to his or her bid. In the second setting, each item is sold through a second-price sealed-bid auction in which the bidder with the highest bid buys the auctioned item at a price equal to the second-highest

bid. We develop below models that prescribe how a risk-neutral procurement agent should bid on these two auctions to maximize its expected utility.

3.2 Bidding in Two First-Price Sealed-Bid Auctions

Consider two first-price sealed-bid auction markets, each selling a distinct item that may interest buyers or bidders. Denote by OBJ_i the item sold through auction market $i = 1, 2$, respectively. Assume that n bidders compete for these two items. We consider private-value auctions in which each bidder's valuation of all item bundles, in this case, $\{\text{OBJ}_1\}$, $\{\text{OBJ}_2\}$, and $\{\text{OBJ}_1, \text{OBJ}_2\}$, is private and known only by this bidder. Denote by U_j bidder j 's valuation function.

In order to gain some initial insights into the structure of this dual auction market without performing a complex, full-fledged strategic analysis, we study a game against nature formulation of the above problem with the following simplifying assumptions.

- All bidders $j = 1, 2, \dots, n$ are risk-neutral.
- Complementarity between OBJ_1 and OBJ_2 exists *only* for bidder 1. In other words, U_1 is super-additive, i.e., $U_1(\{\text{OBJ}_1, \text{OBJ}_2\}) > U_1(\{\text{OBJ}_1\}) + U_1(\{\text{OBJ}_2\})$, whereas U_j for other bidders $j = 2, 3, \dots, n$ are additive, i.e., $U_j(\{\text{OBJ}_1, \text{OBJ}_2\}) = U_j(\{\text{OBJ}_1\}) + U_j(\{\text{OBJ}_2\})$. We denote by δ the extra utility generated for bidder 1 by acquiring both items, i.e., $\delta = U_1(\{\text{OBJ}_1, \text{OBJ}_2\}) - U_1(\{\text{OBJ}_1\}) - U_1(\{\text{OBJ}_2\})$.
- The highest bid among those submitted by bidder $j = 2, 3, \dots, n$ for OBJ_1 is a random variable that follows a known cumulative distribution function (CDF) $G_1(\cdot)$. Denote by $g_1(\cdot)$ the corresponding probability density function (PDF).
- The highest bid among those submitted by bidder $j = 2, 3, \dots, n$ for OBJ_2 is a random variable that follows a known CDF $G_2(\cdot)$. Denote by $g_2(\cdot)$ the corresponding PDF. Furthermore, these two highest bids are independently distributed.

In a full-fledged strategic analysis based on equilibrium concepts such as Bayesian-Nash equilibrium (Fudenberg & Tirole 1991), the distributions of the bids from bidder 1's rivals are derived from their equilibrium bidding strategies. In a game against nature formulation, the bids of rival bidders are treated as part of the uncertain environment (i.e., nature) which is characterized by the two distribution functions ($G_1(\cdot)$ and $G_2(\cdot)$).

We now derive bidder 1's optimal bidding functions in the above game against nature formulation. To simplify the notation, denote by u and v the bid bidder 1 submits to auction market 1 and 2, respectively; also denote by x and y bidder 1's valuation of OBJ_1 and OBJ_2 , respectively. Under the first-price auction, the expected payoff of bidder 1, $EP1(u, v)$, can be calculated as follows.

$$\begin{aligned} EP1(u, v) &= (x - u) \text{Prob}\{\text{bidder 1 wins auction 1}\} \\ &\quad + (y - v) \text{Prob}\{\text{bidder 1 wins auction 2}\} \\ &\quad + \delta \text{Prob}\{\text{bidder 1 wins auctions 1 \& 2}\}. \end{aligned}$$

The probability of bidder 1 winning auction 1 equals $G_1(u)$ and the probability of bidder 1 winning auction 2 equals $G_2(v)$. Thus the above equation can be rewritten as follows.

$$EP1(u, v) = (x - u)G_1(u) + (y - v)G_2(v) + \delta G_1(u)G_2(v). \quad (23)$$

To maximize bidder 1's expected payoff, the first-order conditions for optimal bids (u^*, v^*) must be satisfied. (For ease of exposition, we ignore boundary conditions and various technical assumptions regarding $G_1(\cdot)$ and $G_2(\cdot)$.) The following lemma states these conditions.

Lemma 1 *Bidder 1's optimal bidding functions (u^*, v^*) on two first-price, sealed-bid auction markets have to satisfy the following conditions:*

$$x - u^* + \delta G_2(v^*) = G_1(u^*)/g_1(u^*) \quad (24)$$

$$y - v^* + \delta G_1(u^*) = G_2(v^*)/g_2(v^*). \quad (25)$$

In some special cases, closed-form bidding functions can be derived based on the above conditions. An example is given below.

Example 1 Assume that there are 2 bidders and that bidder 2's bids for both OBJ_1 and OBJ_2 are drawn from a uniform distribution with support $[0, 1]$. Equations (24) and (25) can then be simplified as follows.

$$\begin{aligned} x - 2u^* + \delta v^* &= 0 \\ y - 2v^* + \delta u^* &= 0. \end{aligned}$$

Solving the above system of linear equations (and also considering applicable boundary conditions), we obtain the following optimal bidding functions:

$$u^* = \begin{cases} \min(1, \frac{2x+\delta y}{4-\delta^2}) & \text{if } 0 \leq \delta < 2, \\ 1 & \text{if } \delta \geq 2. \end{cases} \quad (26)$$

$$v^* = \begin{cases} \min(1, \frac{2y+\delta x}{4-\delta^2}) & \text{if } 0 \leq \delta < 2, \\ 1 & \text{if } \delta \geq 2. \end{cases} \quad (27)$$

Note that when $\delta = 0$, the bidding functions reduce to $u^* = x/2$ and $v^* = y/2$. They are precisely a special case of the classical Vickrey solution for individual first-price auctions (Vickrey 1961). We also observe that when δ increases, both u^* and v^* are nondecreasing.

3.3 Bidding in Two Second-Price Sealed-Bid Auctions

We now study how to bid in two second-price sealed-bid auction markets. Consider two such markets, each selling a distinct item of interest. Following the notation from the previous section, we denote by OBJ_i the item sold through auction market $i = 1, 2$, respectively. Assume that n risk-neutral bidders compete for these two items under the private value assumption. Denote by U_j bidder j 's valuation function. Furthermore, we assume that U_1 is super-additive and other $n - 1$ valuation functions are all additive. We use δ to denote the extra utility generated for bidder 1 by acquiring both items.

In addition, we assume that it is known to any bidder $j^* \neq j$ that $U_j(\{\text{OBJ}_1\})$ and $U_j(\{\text{OBJ}_2\})$ are independently distributed according to known CDFs $F_1(\cdot)$ and $F_2(\cdot)$, respectively. These CDFs are assumed to have PDFs, $f_1(\cdot)$ and $f_2(\cdot)$.

The above model is amenable to strategic analysis using Bayesian-Nash equilibrium. We first derive the optimal bidding strategy for bidder $j = 2, 3, \dots, n$. These bidders treat auctions 1 and 2 independently because of the structure of their valuation functions. In each auction, the dominant strategy for these bidders is to bid their values (Milgrom & Weber 1982). In effect, how bidder 1 will bid has no impact on these bidders' optimal bidding behavior.

We now derive the optimal bidding strategy for bidder 1 under the assumption that all other bidders will bid their values.

Observe that, since bidders $j = 2, 3, \dots, n$ bid their values, which are independently distributed according to $F_1(\cdot)$ and $F_2(\cdot)$, the distributions of the highest bids of these bidders are given by CDFs of the order statistics, $(F_1(\cdot))^{n-1}$ and $(F_2(\cdot))^{n-1}$. We use W and Z to denote these two random variables, respectively. In addition, we use x and y to denote bidder 1's valuation of OBJ_1 and OBJ_2 , respectively.

Using \mathbb{I} to denote an indicator function, we write the expected payoff of bidder 1 under his bids u, v as follows.

$$\begin{aligned}
 EP2(u, v) &= \mathbf{E}_W[(x - W)\mathbb{I}_{W < u}] + \mathbf{E}_Z[(y - Z)\mathbb{I}_{Z < v}] + \\
 &\quad \mathbf{E}_{W, Z}[\delta \mathbb{I}_{W < u} \mathbb{I}_{Z < v}] \\
 &= \int_{-\infty}^u (x - W) d(F_1(W))^{n-1} \\
 &\quad + \int_{-\infty}^v (y - Z) d(F_2(Z))^{n-1} \\
 &\quad + \delta (F_1(u))^{n-1} (F_2(v))^{n-1}.
 \end{aligned} \tag{28}$$

When $\delta = 0$, bidder 1's expected payoff $EP2(u, v)$ can be decomposed into two independent parts, producing the standard dominant truth-revealing strategy $u^* = x$ for auction 1 and $v^* = y$ for auction 2. To maximize $EP2(u, v)$ in general, the following first-order conditions on optimal bids (u^*, v^*) must be satisfied.

$$\begin{aligned}(n-1)(x - u^* + \delta(F_2(v^*))^{n-1})(F_1(u^*))^{n-2}f_1(u^*) &= 0 \\ (n-1)(y - v^* + \delta(F_1(u^*))^{n-1})(F_2(v^*))^{n-2}f_2(v^*) &= 0.\end{aligned}$$

Under standard regularity assumptions, these conditions can be further simplified, given in the following lemma.

Lemma 2 *Bidder 1's optimal bidding functions (u^*, v^*) on two second-price, sealed-bid auction markets have to satisfy the following conditions:*

$$x - u^* + \delta(F_2(v^*))^{n-1} = 0 \quad (29)$$

$$y - v^* + \delta(F_1(u^*))^{n-1} = 0. \quad (30)$$

We can easily verify an intuitive property: As δ increases, both u^* and v^* are nondecreasing.

We now summarize the Bayesian equilibrium bidding functions for the model studied in this section: Bidder 1's equilibrium bidding functions are characterized by Lemma 2; all other bidders bid their values. Below we present a special case where u^* and v^* have simple, closed-form solutions.

Example 2 Assume that there are 2 bidders ($n = 2$), and that bidder 2's values for both OBJ_1 and OBJ_2 are drawn from a uniform distribution with support $[0, 1]$. Equations (29) and (30) can then be simplified as follows.

$$\begin{aligned}x - u^* + \delta v^* &= 0 \\ y - v^* + \delta u^* &= 0.\end{aligned}$$

We then obtain the following equilibrium bidding functions:

$$u^* = \begin{cases} \min(1, \frac{x+\delta y}{1-\delta^2}) & \text{if } 0 \leq \delta < 1, \\ 1 & \text{if } \delta \geq 1. \end{cases} \quad (31)$$

$$v^* = \begin{cases} \min(1, \frac{y+\delta x}{1-\delta^2}) & \text{if } 0 \leq \delta < 1, \\ 1 & \text{if } \delta \geq 1. \end{cases} \quad (32)$$

4 Related Work

In recent years, the literature on comparative online shopping from posted-offer markets has been steadily growing, studying both economic decision-making and technological issues (Tsvetovaty, Gini, Mobasher & Wieckowski 1997, Doorenbos, Etzioni & Weld 1997, Montgomery, Hosanagar, Krishnan & Clay 2003). Our work presented in Section 2 shares the same research goals and methodologies as those that focus on economic decision-making mechanisms (e.g., (Montgomery et al. 2003)). One differentiating feature of our work is that we explicitly study

bundling and combinatorial interdependency issues and consider fixed transaction costs, whereas others focus on single-item purchasing but take into consideration issues such as network speed. In effect, our work and the existing online shopping literature are complementary in that our models and the existing models can be readily integrated to create practical online procurement agents that can take advantage of online posted-offer markets.

There is a vast literature on auctions (Klemperer 1999). Here we only discuss several lines of auction research that are directly related to our work reported in Section 3. Most combinatorial auction work assumes the existence of combinatorial auction markets which sell a full range of product bundles (Sandholm 2002, Rothkopf 1994). Recent years have also seen the increasing acceptance of simultaneous ascending auctions in selected applications (Milgrom 2000). Such auctions can be viewed as staged, simultaneously-run English auctions and can mitigate some of the problems with full-fledged combinatorial auctions following, e.g., the Groves-Clarke pivot mechanism. However, it is still unclear whether simultaneous ascending auctions will be suitable for general industrial procurement tasks.

Researchers have also attempted to use a sequence of single-item auctions to deal with product bundling issues (e.g., (Boutilier, Goldszmidt & Sabata 1999)). A comparison between such sequential auction mechanisms and those proposed in this paper will be of significant interest.

Several recent papers study two-product complementarities in various types of auction markets (Brusco & Lopomo 2002, Sherstyuk 2002, Brusco & Lopomo 2001). Under certain restrictive assumptions regarding the size of the extra utility generated by acquiring both products (i.e., δ in Section 3.2), these authors are able to perform a full strategic analysis and obtain equilibrium bidding functions. However, these results do not seem to be generalizable for cases where these restrictions on δ are removed.

The models presented in Section 2 and 3 provide a starting point towards operationalizable, pragmatic decision-making mechanisms that can serve as the core reasoning module of automated procurement agents. We briefly discuss related agent implementation issues. Developing procurement agents for posted-offer markets is fairly straightforward given the current state-of-art Web technologies. With limited Web content parsing capabilities, such agents can relatively easily retrieve the pricing information along with the applicable shipping costs for items of interest from e-tailers, comparative shopping sites, and relevant intermediaries (Zeng & Nissen 2004). Procurement algorithms and heuristics can be implemented either as an integral part of these agents or as some generic service accessible to these agents. Such agents can also (semi)automate various transactional aspects of procurement (Nissen 2001).

Developing procurement agents that can bid effectively across multiple auction markets, however, seems to pose many challenges. Real-time access to information about relevant auction sessions from one or more online auction houses is not difficult when using tools such as those available from `auctiontammer.com`. Obtaining information regarding other buyers in terms of their value distributions,

however, is challenging. One possible approach is to construct empirical distributions using data from past auctions selling similar items.

One of the key challenges in developing online procurement agents is evaluation. How do we know whether and to what extent the models developed actually improve the overall effectiveness of the industrial procurement process? How will the human procurement personnel interact with such automated procurement systems? A promising experimentation-based approach to (partially) address this important evaluation challenge has been attempted by the multi-agent research community (Greenwald & Stone 2001). This community organizes an annual trading agent tournament involving multiple fully-automated trading agents developed by various research groups. These agents are charged with the task of assembling travel packages for a fixed number of customers. They interact with each other through multiple auctions of different types run by the tournament organizer, and are rated by their performance calculated as the difference between the sum of customer utilities (measured by dollars) and the total cost of travel packages. The strategies used by winning agents in this simulated environment clearly provide useful insights about effective agent design in real-world applications. At the same time, such a simulated competitive environment serves as an infrastructure to evaluate empirically any procurement agents before their adoption or even trial in real-world applications.

5 Conclusion and Future Work

This paper presents two sets of models that can guide procurement agents to make optimal procurement decisions in a number of scenarios in which bundling and combinatorial interdependency considerations play an important role. The first set of models formalizes optimal purchasing decisions across posted-offer markets with fixed transaction costs. We discuss the computational complexity of these models and suggest methods based on discrete location theory to deal with the related computational issues. The second set of models focuses on the coordination of bidding activities across two sealed-bid auctions. For each of these models, we derive the conditions that optimal bidding functions have to satisfy.

The work reported in this paper is the beginning of a long-term research agenda aimed at developing practical online procurement mechanisms that can deal with bundling and combinatorial interdependency issues. There remain many open research questions of practical importance in this area of study. We conclude this paper by summarizing some of them we are currently working on.

- We are currently developing a framework to coordinate bidding activities across multiple English auctions when complementarities between items are present. An important complication with multiple English auctions is that these auctions may end at different times. For some auctions, the end time is fixed and known before the auction starts. For others, the end time is unknown and may even be directly influenced by the chosen bidding strategy (e.g., as in soft-ending English auctions). We are in the process of analyzing all these possible scenarios and developing the corresponding contingent bidding strategies.

- Section 3 presents structural results regarding the conditions that the optimal bidding functions have to satisfy. We are currently developing practical computational methods to calculate these optimal bids under a broad spectrum of value distributions.
- Hybrid markets such as a combination of a posted-offer market and an English auction have recently emerged (e.g., eBay’s buy-it-now option). We are working on models that can deal with such hybrid markets as well as those consisting of a heterogeneous collection of auctions.
- Evaluation is an important component of the type of research reported in this paper. We are currently extending the preliminary evaluation for posted-offer markets reported in Section 2.3. We are also collecting data from online auction sites and plan to perform a similar computational study to evaluate the proposed bidding coordination mechanisms across online auction markets.

References

- Ausubel, L. M., Cramton, P., McAfee, R. P. & McMillan, J. (1997). Synergies in wireless telephony: Evidence from the broadband pcs auctions, *Journal of Economics and Management Strategy*.
- Beard, T. R., R. B. Ekelund, J., Ford, G. S. & Saba, R. S. (2001). Price-quality trade-offs and welfare effects in cable television markets, *Journal of Regulatory Economics* **20**(2): 107–123.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation, *Rand Journal of Economics* **25**(2): 242–262.
- Borodin, A. & El-Yaniv, R. (1998). *Online Computation and Competitive Analysis*, Cambridge University Press.
- Boutilier, C., Goldszmidt, M. & Sabata, B. (1999). Sequential auctions for the allocation of resources with complementarities, *Proceedings of IJCAI-1999*, pp. 527–533.
- Brusco, S. & Lopomo, G. (2001). Collusion via signalling in simultaneous ascending bid auctions with heterogeneous objects, with and without complementarities, technical report.
- Brusco, S. & Lopomo, G. (2002). Simultaneous ascending bid auctions with budget constraints, technical report.
- Cox, J. C., Smith, V. L. & Walker, J. W. (1988). Theory and individual behavior of first-price auctions, *Journal of Risk and Uncertainty* **1**: 61–99.
- de Vries, S. & Vohra, R. V. (2003). Combinatorial auctions: A survey, *INFORMS Journal on Computing*.
- Doorenbos, R. B., Etzioni, O. & Weld, D. S. (1997). A scalable comparison-shopping agent for the world-wide web, in W. L. Johnson & B. Hayes-Roth (eds), *Proceedings of the First International Conference on Autonomous Agents (Agents’97)*, ACM Press, Marina del Rey, CA, USA, pp. 39–48.
- Fudenberg, D. & Tirole, J. (1991). *Game Theory*, The MIT Press.
- Gaeth, G. J., Leven, I. P., Chakrabourty, G. & Levin, A. M. (1990). Consumer evaluation of multi-product bundles: An information integration analysis, *Marketing Letters* **2**(1): 47–57.
- Gonsalves, A. (2003). Purchasing executives see increases in e-procurement, TechWeb News, January 15.
- Greenwald, A. R. & Stone, P. (2001). Autonomous bidding agents in the trading agent competition, *IEEE Internet Computing*.

- Hannon, D. (2003). Bringing e-sourcing to a decentralized spend, *Purchasing Magazine*, May 15.
- Harding, M. (1990). *Profitable Purchasing — An implementation handbook for Just-In-Time*, Industrial Press Inc.
- Jain, K., Mahdian, M., Markakis, E., Saberi, A. & Vazirani, V. V. (2003). Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp, *Journal of the ACM* **50**(6): 795–824.
- Jennings, N., Sycara, K. & Wooldridge, M. (1998). A roadmap of agent research and development, *Autonomous Agents and Multi Agent Systems*.
- Jones, J. L. & Koehler, G. J. (2002). Combinatorial auctions using rule-based bids, *Decision Support Systems* **34**(1): 59–74.
- Joshi, A. & Singh, M. P. (1999). Multiagent systems on the net, *Communications of the ACM*.
- Klemperer, P. (1999). Auction theory: A guide to the literature, *Journal of Economic Surveys* **13**(3): 227–286.
- Laudon, K. C. & Traver, C. G. (2002). *E-Commerce: Business, Technology, Society*, updated edition edn, Addison-Wesley.
- McAfee, R. P., McMillian, J. & Whinston, M. D. (1989). Multiproduct monopoly, commodity bundling, and correlation of values, *Economic Journal* **104**(2): 371–383.
- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction, *Journal of Political Economy*.
- Milgrom, P. R. & Weber, R. J. (1982). A theory of auctions and competitive bidding, *Econometrica* **50**(5): 1089–1122.
- Mirchandani, P. B. & Francis, R. L. (1990). *Discrete Location Theory*, Wiley-Interscience Publication.
- Montgomery, A. L., Hosanagar, K., Krishnan, R. & Clay, K. B. (2003). Designing a better shopbot, *Management Science*.
- Nissen, M. E. (2001). Agent-based supply chain integration, *Journal of Information Technology Management* **2**(3): 289–312.
- Pur (1998). Purchasing online, *The Magazine of Total Supply Chain Management*.
- Rassenti, S. J., Smith, V. L. & Bulfinch, R. L. (1982). A combinatorial auction mechanism for airport time slot allocation, *Bell Journal of Economics* **13**: 402–417.
- Rothkopf, M. H. (1994). *Models of Auctions and Competitive Bidding*, Vol. 6 of *Handbooks in Operations Research and Management Science*, Elsevier Science, chapter 19.
- Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions, *Artificial Intelligence* **135**: 1–54.
- Sherstyuk, K. (2002). Simultaneous ascending auctions with common complementarities, U. of Hawaii at Manoa technical report.
- Sycara, K. & Zeng, D. (1996). Coordination of multiple intelligent software agents, *International Journal of Cooperative Information Systems* **5**(2 & 3): 181–211.
- Tsvetovaty, M., Gini, M., Mobasher, B. & Wieckowski, Z. (1997). Magma: An agent-based virtual market for electronic commerce, *Journal of Applied Artificial Intelligence*.
- Vazirani, V. V. (2001). *Approximation Algorithms*, Springer.
- Vickrey, W. (1961). Counterspeculation, auctions and competitive sealed tenders, *Journal of Finance* **16**(1): 8–37.
- Woods, J. A. & Marien, E. J. (eds) (2000). *Supply Chain Yearbook: 2001 Edition*, McGraw-Hill, New York.
- Zeng, D. D. & Nissen, M. E. (2004). *The Internet Encyclopedia, Vol 2*, John Wiley & Sons, Inc., chapter Intelligent Agents, pp. 192–203.