

A combinatorial procurement auction featuring bundle price revelation without free-riding

Robert W. Day^a, S. Raghavan^{b,*}

^a *Operations and Information Management, School of Business, University of Connecticut, Storrs, CT 06269-1041, USA*

^b *The Robert H. Smith School of Business, University of Maryland, College Park, MD 20742-1815, USA*

Received 26 April 2006; received in revised form 14 August 2007; accepted 8 September 2007

Available online 26 September 2007

Abstract

Combinatorial auctions are currently becoming a common practice in industrial procurement, allowing bidders (sellers of goods and services in the procurement setting) to avoid the risk of selling good or service bundles that are incomplete, inefficient, or excessively expensive to deliver. Two major concerns in combinatorial auction design are the revelation or discovery of market price information over the course of the auction, and the inherent computational difficulty (*NP-hardness*) of the underlying “winner-determination” problem. In this paper we describe a new general auction format maintaining the benefits of the adaptive user-selection approach without the problems of free-riding, inefficiency, or distortionary linear prices. This auction format is particularly well-suited to the largest combinatorial auctions for which winner-determination is computationally tractable, because it provides bundle synergy information that is computable in polynomial time for all interactive phases.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Auction design; Combinatorial auctions; Free-riding; Iterative auctions; Procurement

1. Introduction

Early reports on the use of combinatorial procurement auctions (see [18,23,25] for examples) emphasize the robust ability of auction mechanisms to fairly price goods or services provided to a central buyer, allowing benefits on both sides of the market. Suppliers are able to pursue only combinations of goods or services that are cost-efficient to produce, while the buyer or procurer of services benefits from the price competition among several sellers. De-emphasized in the available literature is the complex interplay between the need for market revelation, allowing

competing participants the ability to learn about one another, and the strategic leverage that becomes available when too much information about a bidder is made available through the auction’s revelation environment.

One of the main strengths of traditional single item auctions is that a bidder has the ability to learn about her own true preferences by observing the behavior of her opponents (see the text by Krishna [20]), a feature that is also beneficial in the combinatorial procurement setting. In the industrial procurement of food materials, for example (see the Mars Inc. procurement auction described in [18]), industrial growers of sugar may each hold some private knowledge on the yield of this year’s crop, which in aggregate will determine the relative scarcity of the commodity and thus influence market behavior. Similarly, several transportation providers participating in

* Corresponding author.

E-mail addresses: Bob.Day@business.uconn.edu (R.W. Day), raghavan@umd.edu (S. Raghavan).

a shipping lane auction (see the case of Sears Inc. described in [23]) may each have some (noisy) forecast of future oil prices, which in aggregate are much more accurate than if treated separately. Since the price of oil may heavily influence the value of future shipping contracts, these transportation firms would very much want to ascertain or estimate the beliefs of their opponents by observing their behavior in a multi-stage auction game. We may conclude from these simple examples that revelation of market information is quite beneficial in the procurement setting, and that a (one-shot) sealed-bid mechanism is undesirable for this reason.

While the importance of revealed information and the potential for over-revelation have been only scarcely explored in the combinatorial auction literature, the issue that has been most discussed is the computational difficulty (*NP-hardness*) of the underlying “winner-determination” problem. Several approaches to this computational problem are discussed in the literature (see [1,14] for surveys), but here we focus instead on the adaptive user-selection techniques which address both the computational difficulties and the revelation properties simultaneously.

1.1. Adaptive user selection techniques

The adaptive user-selection approach was first described in an early combinatorial auction paper by Banks et al. [9] who introduce a decentralized mechanism (AUSM) that shifts computational burden away from the central decision making entity and to the bidding agents themselves. Indeed, this is the essential feature of the adaptive user-selection approach; bidders adapt to a changing auction environment by carefully selecting which packages of auction items¹ is the best to bid on at the current time.

¹ We note that as long as the buyer can set reasonable publicly known upper bounds on prices, a procurement auction can be described as a forward auction (with one seller and many buyers) with the auction items being the rights to provide various products or services, and the bids interpreted as discounts from the posted bounds that the real-life providers are offering to “buy” such rights. This allows us to discuss the revenue-maximizing forward auction methods in the context in which they were first proposed, without losing relevance to the procurement auction discussion at hand. In general, we can switch back and forth between forward and reverse (procurement) auctions with little trouble, but mostly stay within the procurement auction setting for this paper, except, at times, when referring to the forward auction research of others. All of the results presented here may be easily reformulated as results for forward auctions. Thus, the auction designs we propose here can be easily applied to the more commonly discussed (and often more strategic) forward auction environments, such as auctions for spectrum licenses.

Later, this idea takes practical form under Kelly and Steinberg [19] who demonstrate a new adaptive user-selection auction format (PAUSE), in which bidders reveal demand in successive rounds by combining their own bids with standing bids of opponents. At each step the new combination of bids forms a complete assignment of auction items to bidders that offers lower total prices to the buyer than the previously proposed auction outcome. This new combination of bids thus becomes the new incumbent solution until (perhaps) displaced by the next winning combination of bids. Importantly, though each bidder has a potentially difficult problem finding a set of bids that complements her own to form a winning combination, the auctioneer has only the relatively simple job of verifying that the new outcome does not award any item to more than one bidder, and that the new combination indeed decreases buyer payments.

The PAUSE auction pursues *efficient outcomes*² by shifting computational burden to the bidders, who have a greater intuition as to which bundles may be of greatest value to them, and can thus focus their computational energy on relevant bundles. This contrasts a typical sealed-bid combinatorial auction, where each bidder can bid on a great number of bundles, letting the auctioneer shoulder the computational burden of sorting them all out.

Despite the nice properties associated with the decentralized computational approach of PAUSE, there is a strategic problem limiting its applicability. In particular, in order to achieve efficiency, the PAUSE auction requires unsuccessful or displaced bids to be available to all bidders for use in their own computations, so that they may be combined with new bids in order to eventually achieve the efficient solution. In Section 2.3, however, we provide an example demonstrating that this level of revelation is too much; revealing this much information causes “free-riding”, which in turn may cause defensive bidders to curtail activity, thus *sacrificing efficiency*. Incentive-compatible (i.e., truth-inducing) mechanisms have also been proposed as a solution to free-riding in some practical contexts (see [7], for example, where investment in knowledge is seen as a public good subject to free-riding). In the most general settings, however, the incentive-compatible Vickrey–Clarke–Groves (VCG) mechanism and its variants result in prices that are exorbitant for the buyer in a procurement

² An outcome is efficient if it gives the right to supply items to those whose costs are lowest (or in the forward setting allocates goods to those who value them the most). When auctions are designed to approach the efficient outcome, we say they are designed for efficiency.

setting, making it unattractive and typically unused in practice. (See [6] for a discussion of this and other weaknesses of VCG.)

1.2. *The clock proxy auction*

A prominent alternative auction design is the clock–proxy auction, presented recently by Ausubel et al. [4]. This technique combines the transparency of linear prices and the efficiency of a sealed-bid (proxy) auction. Although some may argue against the use of an efficient mechanism, others make a convincing case for the revenue optimality of an efficient mechanism, basing their conclusions on an assumption of unfettered resale on the part of auction winners (see [3]). Since the providers of goods or services in the procurement setting may find it profitable to subcontract out part of their obligations following the auction (making it unwise to assume that the auctioneer could prevent “resale” after the auction), we advocate an efficient mechanism for the present procurement setting, and our mechanism achieves it in much the same way.

As in the adaptive user-selection methods (including the currently proposed format), the clock–proxy’s hybrid-auction paradigm is to reveal price information on individual items, eliminating the need for a vast number of inconsequential bundle bids that might otherwise be submitted in a (one-shot) sealed-bid auction. During the revelation (clock) phase, however, bidders see only aggregate demand and price information for each individual item, thus the free-riding described in the PAUSE auction context will not be observed. We argue, however, that where the revelation of the PAUSE auction is too great, the revelation of the clock–proxy may be too little, not providing the bidders with much accurate information about the prices of bundles.

1.3. *Dual-based feedback mechanisms*

Also relevant from the recent literature are certain auction mechanisms that use individual item prices based on dual prices from the linear programming dual of the (relaxed) winner-determination problem, or some modification thereof. Prominent examples include the mechanisms of Kwon et al. [22] and Kwasnica et al. [21], the latter also borrowing features from the simultaneous ascending auction and the AUSM approach, which will continue to be discussed throughout this paper. These designs are elegant in their approach, but due to the duality gap inherent when complementarities are present, the individual item prices produced by these mechanisms absorb synergies among items into the prices for single items, and

are typically non-unique. These prices may therefore distort information about individual items. In contrast, the single-item prices proposed here are unique and do not incorporate synergistic information, thus they cannot suffer from this type of distortion. Further, these iterative dual-based approaches typically solve a hard computational problem in each round, another weakness relative to the current proposal, which postpones computational difficulty to the final sealed-bid round.

1.4. *A new hybrid auction design*

We are motivated, therefore, to capture some of the benefits of both the adaptive user-selection and the clock–proxy approaches, while rejecting some of their difficulties, and to provide an auction format that offers improvements to other formats proposed in the literature. Specifically, we offer a revelation paradigm that provides price information on any bundle (incorporating synergy information that is lost or distorted under linear pricing), while simultaneously diminishing a bidder’s ability to free-ride relative to the existing adaptive user-selection methods. As such, we introduce a new general auction framework, a hybrid auction with the following properties:

- The opportunity for efficient outcomes when bidding is straightforward.
- Revelation of price information for any complementary bundle without explicit public knowledge of specific bundle bids, diminishing a bidder’s ability to free-ride on the revealed information of others.
- User-selection of relevant bundles during the bundle-revelation stage.
- Computational tractability during all interactive (revelation) phases.
- Reliance on linear (single-item) prices only where appropriate, alleviating the distortion caused by the over-use of linear prices.

To our knowledge, there is no published auction design which incorporates all of these features into a single format. Further, our design has several degrees of freedom, allowing pieces of other auction designs to be incorporated into our own format, allowing the auction to take advantage of alternative bid languages, pricing rules, or other special features when desired.

The essential features of our three-stage auction are described in the next three paragraphs. From this brief description of typical participation, we see that the general instructions for this auction format are not complicated, with only the Bid Table format in Stage I requiring an appreciable amount of learning to participate.

In Stage I bidders reveal preferences by submitting *bid tables*, through which a bidder can make several flexible offers to provide items while being assured not to provide a group of substitute items which may be costly to provide simultaneously.

Once all bidders are satisfied with their bid tables the auction enters Stage II, in which bidders are asked in turn to submit a bundle bid. To establish bundle prices, the bidder may ask the auctioneer the price on any bundle of interest, and after “probing” several bundles for a price may offer to provide a particular bundle at that price, or may decline to offer a new bundle bid.

Once all bidders decline to offer a bundle bid in Stage II, the auction enters Stage III, in which bidders make any additional bundle bids that (in the procurement setting) are *above* the probed prices at the end of Stage II. Like the final proxy phase in the clock–proxy auction, this stage ensures efficiency and allows for several non-winning bids to be combined forming a winning set of bids if possible. In order to ensure participation in Stage II, Stage III bundle bids that would have been winning at the end of Stage II are rejected. Thus, a bidder cannot wait until Stage III to make a bid that would have been acceptable in Stage II. However, this rule is not enforced on any subset that is in the tentative possession of a bidder at the end of Stage II. Intuitively, at the end of Stage II a tentative winning bid reflects only that a bidder is willing to bid just enough not to be displaced in Stage II, even though the bidder’s true valuation may represent a stronger bid. In Stage III we must allow the bidder to make this stronger bid to ensure efficiency. This overall design allows us to postpone difficult computational problems until the end of the auction, and encourages revelation in earlier stages in order to reduce the computationally difficulty facing the bidders in later stages.

The remainder of the paper is organized as follows. In Section 2 we provide some background information, in particular describing some of the problems with linear price revelation in the clock–proxy auction and a free-riding problem with the PAUSE auction. Section 3 gives an overview of our proposed three-stage auction design, with added detail for each stage provided in Section 4, Section 5, and Section 6. In Section 7 we provide a brief example demonstrating the flow of the three-stage process for a small procurement auction, and provide concluding remarks in Section 8.

2. Background

2.1. Preliminaries

Informally, combinatorial auctions are often motivated by the presence of “synergy” in the preference

structure of a bidder; some items are worth more in a bundle than the sum of what they are worth separately. We now introduce a formal model of a procurement auction, including a precise definition of synergy.

We take as given a set $I = \{1, 2, \dots, N\}$ of N unique indivisible items (goods or services) to be bought by a single buyer, and a set of suppliers/bidders, $J = \{1, 2, \dots, M\}$. In a general combinatorial procurement auction each bidder j can be modeled by a cost function $c_j: 2^I \rightarrow \mathbf{R}$ and a bidding function $b_j: 2^I \rightarrow \mathbf{R}$. To make precise the notion of synergy, the primary motivation for the use of combinatorial auctions, we adopt the following definition of the synergy σ perceived by bidder j on set $S \subseteq I$:

$$\sigma_j(S) = \sum_{i \in S} c_j(\{i\}) - c_j(S)$$

If synergy is positive, the bundle contains complementary services or goods; it costs the bidder less to provide the package than the sum of providing each individual item. If synergy is negative, the bundle contains mostly substitutes; bidder j may prefer, for example, to provide any one of the items more than providing them all together.

This definition of synergy emphasizes the importance of both positive and negative synergy which should both be considered in the design of a combinatorial auction. If negative synergy is ignored (because, perhaps, the immediate benefits to the auctioneer are not as obvious), bidders will feel the need to curtail their bidding activity, because of the possibility of supplying substitute items at too low a price, potentially leading to lost auction efficiency (loosely speaking auction efficiency is a measure of how close the auction allocation is to the efficient outcome or allocation).

The presence of nonzero synergy motivates an auction format in which each bidder discloses some bidding function b_j over bundles, and the auctioneer solves a combinatorial optimization problem to select the bidders that will provide all the items at the lowest cost. This general *winner-determination* problem can be formulated as an Integer Program (IP), with binary variables $x_j(S)$ that equals 1 if and only if bidder j provides bundle $S \subseteq I$:

$$\min \sum_{j \in J} \sum_{S \subseteq I} b_j(S) \cdot x_j(S) \quad (\text{GWD})$$

$$\text{subject to } \sum_{S \supseteq \{i\}} \sum_{j \in J} x_j(S) = 1, \quad \forall i \in I \quad (1)$$

$$\sum_{S \subseteq I} x_j(S) \leq 1, \quad \forall j \in J \quad (2)$$

$$x_j(S) \in \{0, 1\}, \quad \forall S \subseteq I, \forall j \in J \quad (3)$$

Constraints (1) ensure that each item is provided by exactly one bidder, while constraint set (2) prevents the auctioneer from accepting multiple bids from the same bidder, so that the bidding function announced by the bidder remains accurate. An allocation that solves problem (GWD) when bids are replaced by true costs is defined to be an *efficient allocation*. Given two feasible allocations, we may say that the allocation with a lower (better) objective function value has greater efficiency.³

As noted in the introduction, bidders in real-world auctions may have uncertainty about the values of $c_j(S)$ (and hence $b_j(S)$), but may be able to decrease this uncertainty with increased information about their competitors. Additionally, for even a modest number of auction items the sheer number of bundles to consider makes it difficult to find the bundles S worth submitting bids on. We therefore eschew a sealed-bid (direct revelation) version of the auction and focus instead on revealing information that will be useful to the bidders in their challenging role of determining significant bundles and $b_j(S)$ values.

2.2. A problem with linear price signals

Given our motivation to provide bidders with useful market information over the course of the auction, allowing them to learn about their competition through revealed information, one of the most natural feedback mechanisms to consider is a set of linear price signals. Very simply, this would entail a vector of prices, one for each item in the auction, that can be combined linearly to provide the bidder with some idea of the value of a bundle of items at the current state of the auction (such signals are used in the clock–proxy auction of Ausubel et al. [4]). Though linear prices are admirable for their simplicity, synergy is the primary motivation for a combinatorial auction, and as we now show, linear prices may behave badly in the presence of synergy. Roughly speaking, when linear prices are used, the value of synergy must be absorbed into the prices of some individual items, causing the prices to be distorted and misrepresenting the value of a single good or service.

To see that this may be a problem in a real procurement auction, suppose that we are attempting to run a clock auction (the revelation phase of a clock–proxy auction, as

described in [4]) modified from its original form to accommodate the procurement auction environment. Described in this reverse auction setting, the clock auction proceeds as follows: the auctioneer announces a set of linear prices for the items desired by the buyer, and bidding suppliers each announce a vector indicating the number of each item they are willing to supply at this price vector; for any items with excess supply prices are reduced (in proportion to the amount of excess supply), while other prices remain unchanged. Each submitted vector is treated as a binding indivisible offer from the bidder, and the procedure stops when there is no excess supply on any auction item.

Now consider applying this linear price auction format to a procurement auction with three unique items (A , B , and C) demanded by a single buyer, and four potential suppliers with varying economies-of-scale for providing these items. In particular, suppose that any of the four competing firms may supply any single item at a cost of 9 units. Next suppose that the first three bidders are regional suppliers, each of which can supply a particular pair of items at a volume discount, while the fourth, a national supplier, experiences a significant reduction in costs when providing all three items to the buyer. To be precise, let us assume the following privately known cost structure, where blank entries indicate a combination of items that would be too expensive for the corresponding bidder to provide given their current infrastructure:

Bidder	Bundle						
	A	B	C	AB	AC	BC	ABC
1	9	9	9	12			
2	9	9	9		12		
3	9	9	9			12	
4	9	9	9				19

The clock auction begins with the auctioneer announcing a reasonably high price on each item, for example a price of 10 for each of A , B , and C . At this price vector, bidder 1 would make a profit of $8 = 20 - 12$ for providing the bundle $\{A, B\}$ (which we abbreviate AB), or a profit of $1 = 10 - 9$ for providing any single item. Bidder 1 would thus submit the profit maximizing supply vector $(1, 1, 0) = AB$ to the auctioneer, following the announcement of these prices. Similarly, we would expect bidders 2, 3, and 4 to demand⁴ bundles AC , BC , and ABC , respectively.

³ Although we deal with $b_j(S)$ values in our formulation, instead of the true cost $c_j(S)$, the last stage of our auction—when using Vickrey payments or bidder-Pareto-optimal core payments—as described in Section 6 induces behavior that results in an efficient allocation.

⁴ We assume that bidders bid straightforwardly, i.e. at each round a bidder submits a bundle that maximizes her profit. The use of an activity rule, such as the revealed-preference activity rule suggested in [4], discourages a bidder from distorting preferences for bundles, and encourages straightforward bidding, at least in a heuristic sense, during the revelation phase of a hybrid auction.

With these submissions, excess supply will continue to be equal across all items for the entirety of the clock auction, regardless of the size of the price decrement, and thus prices will fall for all items at exactly the same rate. Excess supply will not cease until the prices offered for the items A , B , and C reaches $6 - \delta$ for each, where δ is the price decrement. At this point, the first three bidders will reduce their bid vectors to zero, since no bundle will be profitable. (Bidder 4 would have dropped out when the prices fell just below $19/3$ for each item.) Unfortunately, because of synergy, neither the final nor the penultimate price vector announced by the auctioneer supports a partition of the items in the auction; at the next to last announcement there is excess supply on every item, at the last there is zero supply for all.

The clock-proxy auction controls for this by following the revelation clock phase with a proxy phase, an efficient sealed-bid combinatorial auction that awards winners higher payments in order to stimulate bidding. In this hybrid auction (with both open and sealed-bid phases), all reported bid vectors are considered binding contracts at the corresponding prices. For example, if bidder 4 reported $(1,1,1)$ at the prices $(6.5,6.5,6.5)$, the auction can hold bidder 4 responsible at the end of the auction to supply all three items for a total payment of 19.5, regardless of what bids are submitted in the final proxy round. In this way the auction converts every bid in the revelation (clock) phase into a bid in the final sealed-bid (proxy) phase, assuring efficiency based on the bids that are submitted, and allowing for more submissions that may also increase auction efficiency (i.e., finds an allocation based on submitted bids that decreases the total cost of procuring the items). For this data the proxy auction would find the winning bid from bidder 4, the lowest bid made on ABC , given straightforward bidding strategies on the part of the bidders.

But the purpose of the open descending-price clock phase is to reveal useful market information, aiding bidders in placing bids, possibly allowing them to place different bids in the sealed-bid final round than they would have without this market information (assuming that they would update their cost structure based on information revealed in the auction). How useful is the information reported to bidders at the end of the clock phase? Assuming a small enough decrement the auctioneer is essentially reporting the price vector $(6,6,6)$ for individual items, in this example. But this price vector is only informative about bundles of size two for this data, and is otherwise misleading. A bidder might be led to believe that a single item bundle could be produced at a cost of around 6 when the actual cost is

9, or that a new bid on the entire set of three items must be below 18 in order to be competitive, when it need only be just below 19.

In fact it is not difficult to see (and is well known) that in the presence of synergy, any set of linear prices will fail to support the efficient solution as a competitive equilibrium for this data. Bidders with incompatible bundles will drive the prices too low for the winners to accept (as bidder 4 cannot accept 18, here). We conclude that linear price signals are misleading and not appropriate in the context of bundle synergy, and that market information should be conveyed to the bidder in some other way.

In the PAUSE auction, bids on bundles become available information for the bidders after the first stage (which accepts bids on single items only), avoiding the problem of linear prices present in the clock-proxy auction. Since bids are binding for the remainder of the auction, a rational bidder will not bid below the cost of a single item in the first stage, and so the first stage of the PAUSE auction should end with single item prices of 9 at the lowest for any item. Assuming bidders are fully aggressive, the next stage will begin with bidders allowed to combine bids on bundles of size 2 (two item bundles) with an available bid of 9 on any single item by itself.

To begin the bundle bidding stage, bidder 1 may combine a standing bid of 9 on item C by itself with a bid to provide AB at a price of 17, thus offering a lower combined total ($=26$), than the buyer would have paid at the end of the first stage ($=27$). Bidder 2 may respond by combining a standing bid of 9 on B by itself with a bid to provide AC at price 16, again lowering the total cost to the buyer. This process will continue through the substage of bids on bundles of size 2, with bidders 1, 2, and 3 each submitting bids on bundles of size two only to be undercut by the next bid.

At no point in this procedure are the price signals unclear: the bidder knows exactly how low to bid for the bundle of size two in order for it to become provisionally winning (i.e. winning until someone else displaces this bid with a lower combined total). Assuming fully aggressive bidders, this may take place until one of these three bidders bids 12 on a bundle of size 2 (combined with the complementary bid of 9 on a single item), at which point none of the other regional bidders will want to undercut (assuming they will not change their preferences based on revealed information).

At this point (when no one wants to bid on a bundle of size two) the PAUSE auction accepts bids on bundles of size three. Since the buyer is now being offered all items at a price of $21 = 9 + 12$, bidder 4 can now achieve the efficient solution by bidding $21 - \delta$, keeping some of her true cost hidden, depending on the size of the

minimum price decrement δ . Again there is no misinformation about the current state of the auction; bidder 4 knows exactly what price to offer for all three items in order to become winning, at least provisionally.

So we see that the adaptive user-selection methods such as PAUSE (and the format proposed here) avoid the problem experienced by linear price signals in the presence of synergy. The amount of information revealed in these methods is enough for bidders to make accurate decisions, and the format in which auction information is conveyed is not distortionary or misleading. We will see in Section 2.3, however, that the amount of information conveyed by the adaptive user-selection methods as described in the previous literature may reveal *too much* information, that is, enough to cause strategic interference.

2.3. A free-rider problem

In the process of providing synergy expression, some combinatorial auction formats introduce what is commonly known as the *threshold problem*, the potential inability of an individual bidder to displace an inefficient package bid without the coordinated efforts of other bidders.

Take for example a three bidder auction with four items A, B, C , and D . Suppose that $c_1(AB) = c_2(CD) = 12$, $c_3(ABCD) = 27$. (Let us assume prohibitively large costs for all other bidders and bundles.) The threshold problem occurs if a bid of 27 on all four items is made by bidder 3 and tentatively accepted by the auctioneer. Neither of the other two bidders is able to submit a package bid individually that will knock out this bid (provide lower cost to the buyer), but if they were simultaneously to bid 13 each on AB and CD , respectively, they can together under-price the bid of 27 and achieve a more efficient solution (and maintain a profit of 1 each). A combinatorial auction should therefore address this threshold issue, by providing the bidders a mechanism for coordinating their bids in order to “gang up” on bidders they cannot defeat individually.

There is, however, a disincentive to honest revelation in these circumstances known as the *free-rider problem*. If bidder 1 suspects or knows that $c_2(CD) = 12$, she may try to only bid as low as 14 on AB hoping that bidder 2 will shoulder more of the burden necessary to overcome bidder 3. If bidder 2 takes the same approach, bidding only 14 hoping that bidder 1 will bid 12, the two may fail to overcome the bid of 27 by combining for a larger aggregate bid of 28.

To address the threshold problem Banks et al. [9] and Kelly and Steinberg [19] suggest the use of a publicly

available stand-by queue containing voluntarily submitted package/price bids that bidders may select to combine with their own bids to make winning package bids. Though this technique does address the threshold problem, it is in fact open to the difficulties of the free-rider problem because too much information is available to aid free-riding in the public queue.

The PAUSE auction is able to avoid the free-rider problem for this particular example by accepting bids on bundles in increasing size order (larger bundles after smaller ones). However, free-riding problems also occur with bundles of the same size, causing trouble for the PAUSE auction. To demonstrate this, let us alter our previous free-rider example slightly, by changing bidder 3 to have costs $c_3(BC) = 10$, and adding a fourth bidder with $c_4(S) = 8 |S|$ for any bundle S .

The PAUSE auction will commence with bids on individual items, and conclude with a lowest price of 8 on every individual item. Next the PAUSE auction is open to bids on bundles of size 2, where each bidder must combine existing bids with her own to achieve a lower cost solution for the buyer. In this bundle bidding stage, each submitted bundle bid becomes “registered” for others to possibly use later, leading to the possibility of free-riding.

To begin the bundle bidding, bidders 1, 2, and 3 each bid 15 on their respective bundles of interest, combined with the existing bids of 8 each from bidder 4, lowering the cost to the buyer from 32 to 31 in every case. The buyer selects one of these “composite” bids (i.e. a collection of bids covering every item) and registers the new bundle bids. Seeing the registered bid of bidder 1, bidder 2 may combine her own existing bid of 15 on CD with bidder 1’s bid of 15 on AB , offering still lower total costs of 30. (Bidder 1 could make the same offer.) Bidder 3, however, cannot combine with another bundle bid and must bid down to 14 on BC in order to lower total cost (with the bids of 8 each on A and D). This continues in much the same way until bidder 3 is bidding 12 on BC with bidder 4’s help to achieve a cost of 28, while bidders 1 and 2 each propose to use the other’s registered bid of 14 together with her own bid of 14 to achieve 28.

At this point we encounter the free-riding problem. Bidders 1 and 2 can see that bidder 3 is continuing to bid on BC while they must work together to displace this bid. Confronted with the decision to continue lowering their bundle bids or to stop, bidders 1 and 2 find themselves in a prisoner’s dilemma: if both go on bidding they can achieve a profit of 1 each, but if one goes on while the other does not, the one who stops bidding will achieve a profit of 2 while the other earns

Stage	Format	Problem Treated
I	Bid Table Auction	negative-synergy exposure
II	Bidders "probe" particular packages for a price, Set a reservation value for unsuccessful probes	positive-synergy exposure
III	Auctioneer conducts proxy auction with "honesty constraints" from Stage II	threshold, free-rider

Fig. 1. Outline of the Procurement Auction.

zero. As is typical with a prisoner's dilemma, the only Nash equilibrium is for both bidders to stop bidding, achieving zero profit each. Because continued bidding allows the other to free-ride, we arrive at a solution in which bidder 3 wins *BC* and the buyer pays too much (27 rather than 26). Or, in other words we arrive at a less efficient solution to the auction.

More desirable is a mechanism through which these bidders can participate and arrive at the efficient solution, sharing the burden of overcoming the combined bids of bidders 3 and 4 equitably, without being given enough information about the bids of others to induce a free-riding situation. In this paper, we introduce a new way to reveal bundle synergy information through price signals only, addressing the free-rider problem by a more controlled revelation of bidder information, and addressing the threshold problem in a final sealed-bid (proxy) auction, assuring a more efficient auction outcome.

3. An overview of the three-stage auction design

The mechanism we propose proceeds in three stages, as shown in Fig. 1. The first two stages are designed to reveal cost information while mitigating the *exposure problem*, a phenomenon widely recognized as a primary reason for the use of combinatorial auctions. If items are provided singly (not as complete packages), bidding on individual items *exposes* the bidder to the possibility of bidding too low for some sets of items. The *negative-synergy exposure problem* occurs when a bidder reveals honest cost information about substitute items and is exposed to the risk of providing several low cost individual items that cost more when taken together (negative synergy). The *positive-synergy exposure problem*⁵ occurs when a bidder reveals honest cost information about complementary items and is exposed

to the risk of getting paid too little for an incomplete set of complements (unrealized positive synergy or economies-of-scale). In either case cautious bidders will have an incentive not to tell the truth and may be reluctant about revealing cost information.

3.1. Stage I

The first stage of our auction addresses the negative-synergy exposure problem, and fills the same role as the first stage in the PAUSE auction. In order to reveal cost information systematically, it is natural to reveal the bids on single items first. Where the PAUSE auction employs a simultaneous ascending auction (SAA) for this purpose, we introduce an improvement, the Bid Table Auction, which allows the bidder to avoid the negative-synergy exposure that is present when the (SAA) is used. As we will see, the strategy space of the simultaneous ascending auction is a strict subset of the Bid Table Auction. This new format offers more possibilities to those who are interested, while anyone wishing to participate as if the format was a SAA is free to do so.

In the *Bid Table Auction*, a bidder is guaranteed the ability to simultaneously lower bids on mutually exclusive alternative items without being forced to provide more than one of these items (thus these items are pure substitutes for this bidder). The bidder expresses this mutual exclusivity by placing the bids for mutually exclusive items in a single column of her *bid table*. Affording the bidder several such columns to work with allows her to express a variety of substitutable preferences. An example of a bid table for a shipping lane auction is illustrated in Fig. 2. In practice, bid tables may be adjusted over several rounds to achieve a final set of bid table entries (such as those in Fig. 2) if bidders are initially unsure of their bid table bids.

The essential feature of a bid table is that only one bid may be accepted from each row (associated with a single item) and only one from each column (associated with a group of mutually exclusive offers). Empty entries may be assumed to be prohibitively large (expensive) offers or simply no offer at all. For this example, a transportation company is bidding on three trips from origin *A* to

⁵ These positive-synergy and negative-synergy definitions of the exposure problem represent a refinement of the definition of the exposure problem given or implied in the literature. As with the refined definition of the threshold problem presented in Section 6, these new distinctions offer a better classification, allowing the auction format to function more smoothly by treating these problems separately.

	W	X xor Y	Z1	Z2
A → B (1)	15			
A → B (2)	10			
A → B (3)	15			
C → D (1)		15		
C → D (2)		17		
E → F (1)			12	18
E → F (2)			12	18
E → F (3)			12	18
G → H		20		

Fig. 2. A Bid Table for a Shipping Lane Auction.

destination *B* (varying perhaps on pick up times: morning, afternoon and evening), two trips from *C* to *D*, three trips from *E* to *F*, and one from *G* to *H*. Column labels might represent bidder resource availability. For example column *W* might represent a truck or driver available to make one trip from *A* to *B*. Similarly, *Z1* and *Z2* might represent two trucks or drivers both of which are available to make one trip each from *E* to *F*.

By bidding as in Fig. 2, the company is expressing that it will only make a single trip from *A* to *B* (since there are offers for these trips only in a single column and only one entry per column may be accepted by the buyer). These bids in column *W* also express that the company would prefer the afternoon option (2), which it would provide at a lower cost (by \$5) than other two options (1 and 3), each of which it would provide at \$15. The company's reasons for this distinction remain private, but it is easy to imagine that a single trip fits neatly into an existing schedule, which already has a driver returning an empty truck from *A* to *B* leaving in the afternoon, for example. This portion of the bid table neatly demonstrates the protection against the negative-synergy exposure problem; the three trips from *A* to *B* are of lesser value when taken together for this bidder (i.e., experience negative synergy) so only one bid table column is used for all three. In a SAA, on the other hand, a bidder must either pick which one of these three to bid on (perhaps picking poorly relative the bids of others) or expose herself to the possibility of providing more than one of these substitute items by bidding on more than one of them.

Similarly, a bidder may express mutual exclusivity among *different* routes by placing bids in the same column. Here the bids in the column labeled *X xor Y* show that the bidder is willing to provide only one of the routes from *C* to *D*, or the single route from *G* to *H*, but not both, at the indicated prices.

If the bidder instead wants to treat two or more items as *partial* substitutes, bids may be spread over several

columns, as in *Z1* and *Z2*. In this case the bidder is expressing the desire to have an increasing price schedule on three options for routes from *E* to *F*. Because the auction is searching for the bid minimizing allocation, it will first accept the bid from the lower price column, *Z1*, paying the bidder \$12 for a single route from *E* to *F*. Since only one bid may be accepted per column, if this bidder is awarded two routes from *E* to *F*, the second bid must be accepted from the column *Z2*, pricing the second route higher at \$18. Clearly, many analogous increasing price schedules over subsets of *I* are possible using this technique.

If the bidder wants instead to express a volume discount (a natural possibility in a combinatorial procurement auction) the bid table format cannot accommodate this expression. Indeed, only substitutable preferences may be described with bid tables, while volume discounts necessarily represent a complementarity among auction items to the provider. Stage I of our auction is designed specifically to reveal substitutable preferences, because the flexibility of substitutable preferences eases the computational needs of the following rounds, which reveal complementary preferences, a more computationally difficult endeavor. In fact, with preference expression that satisfies the “gross substitutes⁶ property”, the computation of Walrasian equilibrium prices can be performed rapidly (i.e., poly-time) for bid table submissions, being not much more computationally difficult than solving the underlying assignment problem. Lehmann et al. [24] provide the first proof of the gross substitutes property for the OXS bidding language (i.e., OR-of-XOR-of-Singletons, equivalent to the language of bid table submission, but more difficult to read in aggregate) while Day and Raghavan [12] provide the relevant algorithm for the winner-determination problem and computation of equilibrium prices.

Stage I of the auction proceeds in an iterative fashion with bidders submitting and adjusting bid tables like the one in Fig. 2. Following the submission of bid tables by all bidders the auctioneer announces a current winning allocation and a price for each item. These prices are dual prices to a slightly modified primal allocation problem and are equal to the highest Walrasian equilibrium prices (as shown in [12]). These prices

⁶ In the context of procurement auctions with quasi-linear preferences, the gross substitutes property states that if prices fall on some items, then a bidder would not wish to provide less of any alternative items for which the prices did not fall. Roughly speaking, this ensures that complementarity is not expressed in the bidders' valuation functions.

convey aggregate information about the value of the items to all bidders as individual commodities; a new entrant would have to accept a price less than or equal to the Walrasian price to provide that item. After observing the current state of the market in terms of Walrasian prices, bidders may then adjust and resubmit their bid table entries according to the following eligibility rules:

- an entry may not increase
- an entry in a non-winning bid table row must be lowered to at most the current price for that row minus one bid increment, or else a “last and best” entry may be made that cannot be adjusted again for the remainder of the auction.
- entries in a winning row of a bid table need not be adjusted.

Such rules are necessary for this stage to ensure that the auction proceeds at a healthy pace and so that bidders may not lay in waiting without revealing their intentions. In [12] we show that the dynamic (descending) version of the Bid Table Auction converges to the same equilibrium prices as the direct revelation Bid Table Auction under the assumption of straightforward bidding.

3.2. Stage II

When a round of Stage I occurs in which no bidder chooses to decrease any bid table entry, the auction enters Stage II, designed to mitigate the positive-synergy exposure problem by accepting all-or-nothing package bids. In this stage each bidder is asked in turn to submit a package of items and a price which guarantees to increase auction efficiency (reduce the objective of the winner-determination problem) if accepted. If this bid survives the remainder of the auction, the bidder provides the items in the package, receiving no less than the specified price, as well as any items she may still be providing *individually* from Stage I at the prices determined there. Because a particular package may be intended to take the place of a set of bid table columns (those that bid on the items in that package), a bidder is allowed to specify those bid-table columns so that the auctioneer may remove them from her bid table. This will ensure that multiple substitute items are not provided at a low cost unintentionally, one from the bid table and one from the package bid.

Using the techniques described in Section 5, each bidder in Stage II can “probe” the market to determine what price must be offered for a particular bundle of items. Each “probe” is simply a query to the auctioneer,

“In order to become the (tentative) winner of bundle S , what price do I have to offer?” Once she has found the subset that she is most willing to provide at the current suggested price, she may submit this package/price pair which is guaranteed to be winning, until (possibly) it is “knocked out” by another bidder’s package bid. Note, each probe is almost computationally costless, requiring only the solution of an easy assignment problem. Thus, in practice a bidder may devise a (polynomial-sized) list of “interesting” bundles to probe automatically at each round of Stage II. As in PAUSE and AUSM, the computational burden is shifted to the bidder to compile this list of interesting bundles, and the development of decision support tools for this purpose in specific environments remains the subject of future study.

Though the order in which package bids should be accepted in Stage II is not set in stone, we suggest that probing may proceed in a round-robin fashion, where the order of bidding is generated randomly for each round, making sure that the last participant in one round is not the first participant in the next. If the results of Stage I (i.e. the final bid tables) and each of the provisionally winning Stage II bids are released to the bidders, they may continue to computationally probe the data while waiting their turn. This would allow bidders to explore more scenarios and potential business plans, so that many package bids may be evaluated in preparation for a turn to bid.

On the other hand, privacy of individual information may be of paramount importance in many applications. Thus we allow for the possibility that Stage II may be conducted with the provisionally winning Stage II bids and final Stage I bids known only to the auctioneer. In this case more time should be given to each bidder for probing, since this would be the bidders’ only way to learn about the competition. When such a scenario occurs in practice, the bidder may have a list of “relevant” bundles which are automatically probed at the beginning of the turn, giving the human users as much time as possible to think about these prices.

Whether the bidders are allowed access to provisional winner information or this information is withheld for privacy at the cost of auction speed, Stage II continues until a round occurs in which every bidder declines to submit a new bid at the probe prices. A bidder who declines to bid in one round may become active again later, leaving room for implicit cooperation when a competitor’s non-overlapping bid helps knock out one that does overlap a bundle of interest. Remaining inactive, however, causes a risk of closing Stage II, resulting in restrictions on the level of bidding that can occur in Stage III.

3.3. Stage III

As pointed out in Section 2, previous adaptive user-selection auctions advocate a publicly available stand-by queue in order to mitigate the threshold problem and promote efficiency. Eschewing this approach because of its openness to the free-rider problem, we instead suggest the use of a *third auction stage* to mitigate both of these problems. Rather than consider collections of non-winning package bids in our own Stage II, our auction format does not address the threshold problem explicitly until Stage III. The reason for doing this is similar to the reason for the separation of Stage I and Stage II: We do not begin to admit package bids in Stage II until after we know that such bids will beat the individual item prices derived in Stage I; similarly, we do not want to consider a collection of individual package bids displacing some other collection of package bids until no more improving individual package bids are possible.

In order to address the threshold problem and the free-rider problem simultaneously, we conduct a sealed-bid auction in Stage III which takes bid information from Stages I and II and attempts to improve upon the allocation found at the end of Stage II. When a bidder probes a particular set in Stage II, she is told a price that she must accept in order to become the provider of this bundle, at least for the time being. If she chooses to accept the reported price at the time of probing, she becomes the provisional winner of this package deal. If she decides instead that her cost for this bundle is higher than the amount necessary to win that package unilaterally, she may be forced to wait until Stage III to submit her bid on this bundle, hoping that it may then become part of the winning *collection* of bids.

In order to assure participation in Stage II, a bid on a particular bundle in Stage III must be more than would become winning in Stage II. If a bidder passes on the price of every bundle in a round of Stage II, she forfeits her right to bid the current probe amount (or lower) for any bundle in Stage III. Further, every Stage III bundle bid is screened to ensure that it would not have been winning in Stage II; thus an unprobed bid is treated just as a probed bid rejected by the bidder. For an honest bidder, this rule will not be restrictive when the current (probe) price for a bundle is lower than her true cost for that bundle. We therefore refer to these restrictions as *honesty constraints*.

If a bidder is winning a particular bundle bid at the termination of Stage II, however, the auction has not necessarily discovered the true cost of that bundle to the bidder. The auction therefore puts no honesty constraint on a Stage III bundle bid for a bidder that is the provider of that bundle at the end of Stage II. This provides an

additional incentive for participation in Stage II; a bundle won in Stage II may be protected by an arbitrarily low bid in Stage III, despite limitations on the bids of competitors.

Within Stage III, the auctioneer will determine an allocation that is more efficient than the one achieved in Stage II (if one exists), and compute competitive prices, such as Vickrey–Clarke–Groves payments or bidder–Pareto-optimal core payments (see [5,13,17,27]). Such pricing mechanisms usually ensure positive profit to each supplier, except perhaps where market competition is too tight (i.e. when a bidder must supply at the cost of her bid because any higher payment could be contested by a competitor who is willing to provide for less).

As in the clock–proxy auction, the final sealed-bid round allows the entire auction to inherit the efficiency and incentive properties of the payment rule used in the final stage. A major concern is “jump-bidding” in intermediate rounds, in which bidders wait until later rounds (or stages) to reveal their preferences. Just as with the revealed preference activity rule used to mitigate jump-bidding in the clock–proxy auction, the honesty constraints proposed here assume that bidders bid honestly and consistently, and constrain the later activity of bidders based on this assumption of honesty.

Strategically, this design makes free-riding as difficult as possible. When a bidder decides not to place any further bids on any bundles in Stage II (i.e., when a bidder passes on the opportunity to bid when it is her turn), she relinquishes the ability to bid any amount below the probe price in the final sealed-bid round for any bundle she is not winning at the end of Stage II. This makes it unwise not to accept the probe price except when it is truly too low.

Next, faced with the decision of what to report for her valuation of this bundle in Stage III, she cannot test the waters to see how much of a free-ride is possible. Though she may decide to report a valuation that is more than her true cost for this bundle, the higher the price she reports the greater the probability that she will regrettably miss an allocation for which she provides this bundle at positive profit. Honest bidders, on the other hand, have the opportunity to achieve an efficient displacement if one exists, and will be paid fair prices determined by whichever sealed-bid pricing algorithm is pre-announced by the auctioneer.

3.4. Summary

The three-stage auction design we have described has several attractive properties. For the remainder of the paper we will elaborate upon the theoretical landscape in which this design functions, as well as providing more

details into how each stage will proceed, and an illustrative example. Before delving deeper into this exposition, we list a few of the desirable properties which have been or will be discussed.

Stage I

- The winner-determination problem can be solved quickly (polynomial time in the number of items and bidders).
- Bidders may bid competitively on several substitute items at once, without the risk of providing more than desired.
- Individual item price signals are revealed, diminishing the number of package bids that are necessary in Stages II and III.
- The gross substitutes property is automatically upheld throughout Stage I.

Stage II

- Bidders may bid on packages without exposure to the risk of being paid too little for an incomplete, high-cost package (e.g. one with unrealized economies-of-scale).
- Each bundle “probe” is computable in polynomial time, as is the auctioneer’s validation of a winning bundle/price pair.
- Non-participation in Stage II constrains activity in Stage III, making free-riding and demand reduction risky decisions.
- A bidder may guarantee the final right to provide the package she is awarded at the end of Stage II, but may be charged to beat out her Stage III competition.

Stage III

- If non-winning bids from Stage II can be combined to overcome a provisionally winning package bid, this will be corrected in the final sealed-bid round.
- An efficient allocation is computed given all revealed bid information.
- Competitive prices are determined which share the cost of displacement fairly among those who benefit, alleviating the free-rider problem. (We advocate bidder-Pareto-optimal core payments, though VCG or any other payment mechanism may be used.)

4. Stage I: bid tables vs. SAA

An important feature of our model is the order in which different types of preference information is revealed over the course of the auction, postponing the

solution of difficult computational problems until they are assured to be necessary. For example, the all-or-nothing package bids of Stage II introduce a high level of computational difficulty (we discuss the *NP*-completeness of the feasibility problem with such bids in Section 5). By first executing a Bid Table Auction in Stage I, we eliminate from consideration any package bid which doesn’t at least beat the “individual prices” for the items in the package. We place Stage I before Stage II in order to drive down these prices on individual items as much as possible before considering any package bids, eliminating as many as possible from consideration since they are more difficult to handle computationally. We then shift the computational burden to the bidder who must decide which bundle to bid on in Stage II.

This idea of first driving up individual prices and then shifting computational burden to the bidders is present in the PAUSE auction of Kelly and Steinberg [19]. However, the use of the SAA in their own Stage I does not address the negative-synergy exposure problem, unlike the Bid Table Auction. Indeed, if a Bid Table Auction was simply used to replace the SAA in Stage I of Kelly and Steinberg’s model, bidders would experience less negative-synergy exposure and be able to bid more aggressively on individual items. The resulting lower prices on individual items will in turn ease more of the computational burden of package bidding. This can be shown easily by noting that if bidders bid with only a single active entry in each bid table column they are participating exactly in the SAA.

The following simple example illustrates this point: suppose that $c_j(\{A\}) = c_j(\{B\}) = 3$, but $c_j(\{A, B\}) = 8$. Restricting bidders to the SAA is equivalent to allowing only one active entry per column in a bid table, thus the lowest this bidder would bid in SAA is equivalent to the bid table:

A	4	
B		4

since any lower would expose the bidder to providing A and B at less than cost on the bundle $\{A, B\}$. The Bid Table Auction makes no such restriction, allowing for the bid:

A	3	5
B	3	5

If, for example, this bidder is pitted against another identical bidder, the SAA can guarantee individual item prices no less than \$4 each, while the protection from negative synergy exposure introduced by the Bid Table

Auction allows the prices to drop to \$3 each. Since the purpose of the first stage of either auction is to reveal lowest price information on single items (in the procurement environment), one that can provably achieve lower individual item prices clearly dominates. This effect strengthens the Stage I results relative to the SAA, because many bids in Stage II need not be considered if lower prices are established in Stage I.

Replacing the SAA with a Bid Table Auction has other benefits besides just achieving lower prices. In Stage I of the three-stage auction proposed here, new market entrants will have a greater opportunity to make simultaneous aggressive bids relative to the SAA, forcing market incumbents to either let them into the market or share the burden of keeping them out. Fair entry into the market place has been an expressed consideration in the adoption of auctions in many environments (for example in the airport landing-slot context described in [8]), and we feel that the use of the Bid Table Auction will also enhance the ability of new entrants to compete in the market by decreasing the opportunity for the “bullying” punishment strategies described by Cramton and Schwartz [10] for the SAA. In this behavior (observed in actual SAA implementations), an aggressive bidder drives up the price (in a forward auction) on a competitor as punishment for placing a competing bid. In the Bid Table Auction this aggressive behavior is not possible; only price signals are revealed to the bidder at each iteration, making it hard to tell who to bully. Further, each bidder simultaneously competes on several substitutable items without exposure to receiving more than needed; a bullying bidder would have to drive up the price on all items for the punishment strategy to be effective, increasing her chances of winning an item that was bid on only to punish a competitor.

As a final note in this comparison, the Bid Table Auction may be easier to participate in (despite a higher fixed-cost of learning a more complicated format), because bidders who know their lowest costs for substitutable items may simply submit their final bid table and wait while others play iteratively and watch the prices go down. Indeed, the iterative nature of the bid table auction is provided to allow for some learning about the preferences of others. For bidders who are well informed of their own cost structure and do not need to learn much about the bids of others, we would expect Stage I to conclude after very few rounds (perhaps one). The SAA, however, will always need several rounds for bidders to convey their preferences. Thus the Bid Table Auction allows bidders to convey their preferences for substitute items more succinctly, allowing them to get to

Stage II where they will bid on complementary packages more quickly, and achieves better (lower) individual items prices in the process.

The winner-determination problem for Stage I is simply an assignment problem. That is, it assign items to bidders in order to minimize the cost of procuring these items, such that every item is assigned to exactly one bidder, and so that at most one item is assigned to any column k in bidder j 's bid table:

$$z = \min \sum_{(i,j,k) \in I \times J \times K_j} b_{ijk} \cdot x_{ijk} \quad (A)$$

$$\text{subject to } \sum_{i \in I} x_{ijk} \leq 1, \forall (j, k) \text{ with } j \in J \text{ and } k \in K_j$$

$$\sum_{j \in J} \sum_{k \in K_j} x_{ijk} = 1, \forall i \in I \quad (4)$$

$$x_{ijk} \geq 0, \forall i, j, k$$

where each b_{ijk} represents bidder j 's bid on item i in her k th bid table column, and K_j indexes all columns of bidder j . This problem is well known to be total-unimodular, (see for example [26]) so that we do not need explicit integer constraints on the decision variables x_{ijk} which equals 1 exactly when the item i is assigned to bidder j 's k th column. Indeed, the linear program (A) automatically solves the corresponding Integer Program to integer optimality (and does so for integer right-hand-sides (RHS), allowing also for an auction for multiple identical items.)

In [12] we provide a simple algorithm to obtain maximal Walrasian equilibrium prices on each item starting from the values of the dual variables to problem (A), and we propose that these prices be reported by the auctioneer at each round of Stage I. Briefly summarized, the generic constraint of the linear programming dual to (A) is of the form $p_i - s_{jk} \leq b_{ijk}$ corresponding to each x_{ijk} variable, where p_i may be interpreted as a price for item i and s_{jk} may be interpreted as the surplus at column j, k . By duality, this constraint will be satisfied at equality when $x_{ijk} = 1$, properly defining the surplus on a column, while each constraint otherwise says that the price paid by the winner of an item must be low enough so that each competing bid table column cannot achieve higher surplus by providing the same item at a lower price.

We [12] show that if we first solve the primal problem (A), the resulting dual prices correspond to a Walrasian equilibrium, but not necessarily the highest Walrasian equilibrium. To compute the highest Walrasian equilibrium prices, we may re-solve the dual problem after relaxing any “self competition” constraints (in which one

bid table column would set the price for an item won by the same bidder) based on our previous solution to the primal problem. The technique is a natural generalization of the methods of Demange et al. [15] (who use an assumption of unit demand, i.e., a single bid table column per bidder), and an example of the procedure's results can be seen in Section 7.

5. Stage II: package bidding and probing

Preferences for complementary bundles introduce a potential for computational difficulty into the auction process. The problems of determining the winning bids and prices for a given round of Stage I can be achieved in polynomial time, as shown in [12]. In order to develop Stage II, in which we wish to include package bidding (for complementary packages), we must consider the computational difficulties inherent when complements are introduced.

Consider scenarios in which bidders report demand correspondences as in Gul and Stachetti [16] and Ausubel [2]. In these (forward) auctions bidders report which bundles they would be willing to take at current prices (their demand correspondences), after which the auctioneer selects a feasible allocation with one bundle from each correspondence, if possible. The prices are changed in order to lower demand, and the process repeated until such a feasible allocation exists. Both assume the gross substitutes property, but one might wonder if their demand correspondence reporting technique might apply to a more general model of bidders. The following theorem (proved in [11]) recognizes a difficulty in this scenario: with a slightly more general set of bidder preferences finding a feasible allocation for the subproblem at each round is *NP*-complete.

Theorem 5.1. If the demand correspondence $\{\emptyset, S\}$ is admissible for any $S \subseteq I$, then the problem of determining whether there exists a feasible allocation of items among the bidders according to these reports is *NP*-complete.

The above theorem states that if a there exist complementarities (positive synergies) that would cause a bidder to demand all of a particular bundle or nothing at all, then the auctioneer's problem of even determining the existence of a feasible solution is *NP*-complete. Thus, any extension of bidder preferences which allows for all-or-nothing bids may expect computational difficulties. As it turns out (see [11]), all-or-nothing bidding of this type (and the resulting computational difficulty) does not occur when the gross substitutes condition holds; if a bidder currently demands both S and $\emptyset$, she must also demand every subset of S under gross substitutes.

Demand correspondence reporting can present a computational difficulty *at every round* when the gross substitutes condition does not hold. Our overall approach to an auction with complements is therefore as follows: use the bid-table format in Stage I to enforce gross substitutes and reveal price information, and then accept package bids to promote price discovery and bundle price revelation while avoiding the computational burden. The method for avoiding this computational burden (until Stage III) is to only consider “strong” bundle bids in Stage II, those which can immediately become winning without being combined with bundle bids of others that are not currently winning.

Bidders in Stage II will “probe” various sets to determine whether they are willing to place a “strong” bundle bid to compete for this set, or instead to indicate that they have reached a lower bound on the price they are willing to accept for this set. For the first bidder in Stage II, this will proceed as follows.

The bidder picks a set of items S that seems attractive for probing, using the current prices as a first approximation of what bundles are most attractive. When the bidder reports this set, the auctioneer (an appropriately designed software package) probes it to determine a price. Since no other package bids have been accepted, this task is equivalent to finding a price such that auction efficiency increases if the package bid is accepted and the other items are awarded according to the reports from Stage I.

Given the first bidder in Stage II is probing some bundle S , let z^* be the objective value from (A) at the end of Stage I, and let z_{-S} be a solution to the allocation problem (A) after changing the RHS of each demand constraint in the set (4) to zero $\forall i \in S$. This is of course the efficient solution to the auction in the absence of the items in S .

The auctioneer may then report to the bidder a price of $p_S = z^* - z_{-S} - \delta$. If the bidder is willing to accept this price, then auction efficiency is increased one decrement⁷ δ by this bidder providing S at p_S and all other items being provided according to the assignment solution yielding z_{-S} . Finding this assignment and corresponding value of z_{-S} is not computationally burdensome: we merely solve a smaller assignment problem than was solved at each round of Stage I. A bidder may continue to probe several bundles before choosing one to bid on, if any.

In order to continue this procedure past the first bidder in Stage II, we must first provide notation for the set of packages for which package bids have been provisionally accepted thus far, S . Now, when a bidder

⁷ In practice, δ may be adjusted to control the pace of the auction, with larger values early on giving way to smaller values later in Stage II.

probes set S , the auctioneer takes the current auction value z^* , and computes z_{-S} , a solution to the primal allocation problem with the set of restrictions:

$RHS = 0 \forall$ constraints in (4) with $i \in S$ or $i \in T$

where $T \in S$ with $S \cap T = \emptyset$

plus the current bid values on all bundles T with $S \cap T = \emptyset$.

In words, the auctioneer computes an optimal assignment of items to bid table entries, given that S is reserved for the current bidder and that a currently winning bid on any subset T not overlapping with S reserves the items it is already winning. Given this assignment of items not in S , the probing bidder must offer to provide the set S at a price that results in an increase in efficiency by δ . Note that if we tried to find an optimal assignment of the items not in S while additionally considering the use of previously knocked out bundle bids, the problem again becomes *NP-hard*. We therefore intentionally leave out all bundle bids except the ones that are currently winning and do not overlap the bundle being probed.

The auctioneer now reports a probe value $p_S = z^* - z_{-S} - \delta$ as before. If this price is accepted, the new bid will push any overlapping bids out of the current allocation, so that any items in a displaced package bid but not in S must be allocated according to the assignment information from Stage I. The bidder will pay the amount necessary to compensate the auction for the drop in efficiency from displacing any currently accepted package bids and reverting to bid table values for leftover items. Again the computations are easily handled as restricted versions of the assignment problem (A). Since all such restrictions are simply bounds on *RHS* values for (A) (or equivalently removing a set of variables), and by the total unimodularity of problem (A), the polynomial solvability is not sacrificed. After each iteration in which a bid on probed package S is accepted, the set S is updated by adding S and removing any package T with $S \cap T = \emptyset$.

One criticism of the PAUSE auction is that the upper-hand is given to bidders with greater computational resources (and hence those with greater financial resources) who can construct composite bids (an *NP-hard* task) more easily than less endowed competitors. The auction proposed here, however, provides no such advantage. With computational tractability maintained throughout the bundle-revelation stage, computations can be left in the hands of the auctioneer. Bidders will each have an equal opportunity to probe the auction and need not be given access to existing bundle bids for

simulations or other strategic computations favoring those with greater resources.

In the Stage II package auction (as described thus far in this section) there is still a potential exposure problem. A bidder may provide a particularly desirable package from a bid in Stage II, but then also provide several other individual items that were not awarded as part of a package. This will occur, for example, if the bidder uses a particular column to bid on a certain set of substitute items and then wins the right to provide one of these items as part of a package bid. This problem easily disappears if we allow (or force) bidders to specify which of their bid table columns corresponds to (bids on) the items in the package bid. We may then simply remove (temporarily ignore) these columns from consideration in the calculation of z_{-S} . This *Package/Column Designation* suggests a reasonable user-interface for Stage II probing; the bundle to be probed may be specified by simply clicking on the bid table entry corresponding to the row of the desired item and the column to be removed if this package bid is accepted. The auctioneer removes the corresponding columns from the bidder's bid table until (perhaps) the bundle bid is no longer winning. This assures that the protection against negative-synergy exposure is maintained in Stage II as it was intended in Stage I.

This removal of bid table columns suggests a further incentive to participation in the Bid Table Auction of Stage I. As more packages are accepted into the final allocation, more and more bid table columns will be removed. Further, certain items may remain "loose," not being incorporated in any package bid. As fewer bid table columns remain active in the auction, competition on these loose items diminishes, resulting in higher final prices for items that are eventually won from bid table entries. Intuitively, a bid table entry is very flexible and may be applied at any time; by making this flexible offer to provide an item in the early rounds, a bidder may be rewarded with an attractive price in the final allocation.

6. Stage III: the threshold problem, the free-rider problem and efficiency

The auction format described in the present work differs from earlier combinatorial auction literature in its treatment of the threshold problem. As mentioned above, some advocate a publicly available stand-by queue of non-winning bids which an interested bidder may combine with her own to form a winning bid (see [9,19]). We argue that this revelation of information is beyond what is necessary to achieve efficiency and allows for tactical manipulation with little if any benefit. As previously mentioned, this tactical manipulation comes in the form of free-riding. A

bidder with the ability to view the valuations of others will find it in her best interests to reveal as little valuation as is necessary to form a winning bid, placing as much burden as possible on those who have revealed their valuations to the stand-by queue. Intelligent bidders may therefore cautiously shade down their own revelation to the queue, possibly not revealing enough information for an efficient allocation to be reached. Although we do not provide a formal game-theoretic analysis, it should be clear that the presence of the stand-by queue introduces two potential problems related to dishonest revelation: bidders taking advantage of the revealed information of others (free-riding) and the inability to achieve efficiency due to cautious shading (free-riding defense strategies).

Even if for some altruistic reason we were to assume truthful revelation (thereby nullifying the free-rider problem) another problem remains with the stand-by queue approach. Maintaining the availability of this bid information throughout the auction seems to diagnose the threshold problem as a barrier to bidding throughout the auction, when in fact the threshold problem should be recognized as a barrier to efficiency only in the final allocation. A bidder in an auction with a stand-by queue may spend valuable computation time attempting to find a collection of available queue bids to combine with her own, only to have the successfully combined coalitional bid knocked out in subsequent rounds. To clarify the situation, we advocate a new definition of the threshold problem:

Definition:

The *Threshold Problem* occurs when a collection of bids, each of which cannot become winning unilaterally, can be combined to reach an efficient final allocation.

There are two changes here from the definition given loosely in Section 2 (and in the literature at large.) One distinction is the word “final” which makes it unnecessary in particular to consider a set of bids which may be combined to displace a bid that eventually is not winning (as is the case in the intermediate phases of AUSM [9] or PAUSE [19].) A second change is “each of which cannot become winning unilaterally.” If a bidder is willing to submit a bid which can become winning unilaterally, we would want her to do so, rather than hiding this low valuation only to submit this into the final sealed-bid phase when it is too late for her opponents to react. To encourage revelation during Stage II we assume that a bidder has bid truthfully and hold her to it in Stage III via bounds on her Stage III bids. These “honesty constraints” make it risky to hide preferences in Stage II, as doing so will limit one’s ability to bid aggressively in Stage III.

These observations suggest that the threshold problem should not be addressed until a final sealed-bid auction, regardless of what type of sealed-bid auction is used. For

the remainder of this section we describe the specific winner-determination problem for a sealed-bid auction tailored to the multi-stage auction as developed to this point. In particular, we want to maintain the bid table entries as bid information, allow for package bidding, and uphold the Package/Column designations for each package bid as described in Section 5. We now describe an integer programming formulation for determining winners, treating all revealed bid information as binding.

In addition to each (final) bid table entry b_{ijk} , we will have (from Stages II and III) some indexed list of L package bids $\{(B_1, S_1), \dots (B_b, S_l), \dots (B_L, S_L)\}$, where each (B_l, S_l) represents a bid of B_l for set S_l . In practice a sophisticated bid language may be used that bids on several packages simultaneously, but for now we maintain bids on an arbitrary number of exclusive packages for the sake of generality. This set of package bids will contain any bid that is winning at the end of Stage II, and any other bids submitted for Stage III after having been screened to ensure that each B_l could not become winning at the end of Stage II.

We introduce the notation C_l to denote the set of columns that a bidder has associated with a particular package bid (B_l, S_l) , and 0–1 decision variable y_l which equals one if and only if package bid (B_l, S_l) is accepted. Also, we will wish to express that if (B_l, S_l) is a winning package bid then the bid table columns associated with tasks in C_l will become inactive (i.e. not win any items). To find an optimal allocation we solve the following allocation problem for Stage III, which we denote (A_3) :

$$\min \sum_{(i,j,k) \in I \times J \times K_j} b_{ijk} \cdot x_{ijk} + \sum_{l=1 \text{ to } L} B_l \cdot y_l \quad (A_3)$$

$$\text{subject to } \sum_{i \in I} x_{ijk} + \sum_{l | (j,k) \in C_l} y_l \leq 1, \forall (j,k)$$

$$\text{with } j \in J \text{ and } k \in K_j$$

$$\sum_{j \in J} \sum_{k | (j,k) \in C_l} x_{ijk} + \sum_{l | i \in S_l} y_l = 1, \forall i \in I$$

$$\sum_{l | B_l \text{ is made by } j} y_l \leq 1, \forall j$$

$$x_{ijk} \in \{0, 1\}, \forall i, j, k$$

$$y_l \in \{0, 1\}, \forall l$$

This modification of formulation (A) (from Section 5) places the y_l variables so that if a package bid wins its specified collection of items, then the columns associated with this bid will not win other items, and no other

bid which includes any of those items may be accepted. The third set of constraints assures an *XOR* relationship among the bundle bids.

The solution to (A_3) gives the final efficient winners and their awarded bundles, but does not tell us what prices to charge the bidders for the items they receive. Pay-as-bid is one extreme possibility, but clearly this pricing rule provides the strongest incentives for bid shading, risking a less efficient outcome. The other extreme possibility is to adopt a VCG payment structure which is guaranteed to be incentive-compatible (induce truthful bidding).

There are, however, several problems with VCG payments when positive synergy is present. Most notably, the VCG outcome is often not a “core” outcome; a losing group of bidders may be able to show a consistent set of submitted bids that offers to provide the buyer with all items at lower total cost. Intuitively, the VCG mechanism “pays” the bidders in order to induce truth-telling, and in the presence of complementary items, these payments can grow to absurd levels in regard to perceived fairness (and what is rational for the buyer).

The recent work of Ausubel and Milgrom [5], Day and Raghavan [13], Hoffman et al. [17], and Wurman et al. [27] advocate the use of “bidder-Pareto-optimal core” pricing, which does pay rents to bidders in order to induce truth-telling, but not such high rents as to cause this absurdity sometimes associated VCG payments. Loosely speaking, with core pricing no coalition of bidders could offer to provide the items to the buyer at lower total cost, and have every member of the coalition prefer this offer to the outcome of the auction (or be at worst indifferent). It is worth noting though, that for some instances the VCG payments are in the core, in which case the bidder-Pareto-optimal prices correspond exactly to the VCG prices, as is the case for the following example of Section 7. In other cases, bidder-Pareto-optimal core payments are different from the VCG payments, in which case truthful revelation is not a dominant strategy for bidders. However, as we show in [13], at Nash equilibrium in a bidder-Pareto-optimal core auction each bidder will bid “close to truthfully,” resulting in an efficient allocation.

7. Example: a shipping lane auction

We now present an illustrative example to demonstrate how the stages of our auction design might proceed, using a fairly simple model of synergies. Suppose that a firm is holding a procurement auction for shipping lanes, and bidders *X*, *Y*, and *Z* have come to bid on the available transportation contracts *A*, *B*, *C*, *D*, and

E. The network associated with these routes is depicted graphically in Fig. 3.

Imagine that each of these bidders has costs for each of these items (contracts) individually and perceives some natural synergies among routes. In particular, any route that forms a 2-cycle will cost 3 units less than the assignment of the individual contracts for bidder *X*, and similarly every 3-cycle will save bidder *X* 7 units. These decreases in costs may have to do with driver scheduling and the fact that a cycle diminishes the number of empty loads for which a shipping company may still have to pay their drivers. Among the contracts available in this auction there is only one 2-cycle (*A* – *B*) and only one 3-cycle (*A* – *C* – *D*), but let us imagine that bidder *X* has an existing contract to move shipments along the path *F*, allowing this company to experience the benefits of a 3-cycle (*C* – *E* – *F*) from obtaining contracts on just *C* and *E* in the present auction. For the other bidders, we suppose that bidder *Y* has no existing contracts, but perceives a cost decrease of 7 units on a 2-cycle and 4 units on a 3-cycle, and bidder *Z* also has no existing contracts and experiences a decrease in costs of 6 on a 2-cycle and 9 on a 3-cycle.

For the sake of brevity, we do not show iterations of Stage I, in which players alter their bid tables, but instead present the following set of final Stage I bid table entries:

	Bidder X				Bidder Y			Bidder Z			Prices	
<i>A</i>	11	11	11	11	13	13	13	12	12	13	13	12
<i>B</i>	14	12	14		13	13	13	11	11	12	12	13
<i>C</i>	14	12	14		13	13	13	15	15	16	16	13
<i>D</i>			16	13	16	16	16	17	17	18	18	16
<i>E</i>			22	18	24	24	24	23	23	24	24	23

Bidder *Y* has simple preferences: at most three routes could be supported with an additive cost structure for items won individually. Bidder *Z* similarly has additive costs for the first two routes provided, but must expand (perhaps train a new driver) at a cost of 1 for the third or fourth route awarded (thus the third and fourth columns include an extra unit of cost.) Bidder *X* has a slightly more complex cost structure: routes *B* and *C* both




Fig. 3. Shipping Lanes to be Auctioned.

originate at the same city, and are hence partially substitutable; by putting the lowest cost for these two routes in the same column, bidder X has made certain that both of these routes will not be provided at their lowest cost. A similar partial substitutability is expressed for D and E , while A could be provided at cost 11 regardless of what other routes are provided. Given these bid table entries, the winning entries (shown in bold) are the tentative winners of a contract on the respective routes, at the indicated highest Walrasian equilibrium prices. All but D are awarded at a profit, because D could be provided by bidder X if bidder Y were to be paid any more than 16.

Together with the synergies described above, these bid tables describe the following costs for bundles. (Note: we need only list bundles which contain positive synergy; individual routes may be added on at their bid table value.)

	AB	ACD	$ABCE$	CE
Bidder X	20	29	41	23
Bidder Y	19	38		
Bidder Z	17	37		

In Stage II, each bidder asks the auctioneer (probes) for a price quote on any bundle, designating that a lowest cost assignment of their bid table columns be removed if the price offer is accepted. For simplicity, we assume straightforward bidding in which the bidder always chooses a bundle that offers the most profit, as long as that profit is non-negative. The rounds of Stage II proceed in random order as follows, with accepted prices shown in bold:

Rounds	AB	ACD	$ABCE$	CE
Round 1				
Bidder Z	21	37		
Profit	4	0		
Bidder Y	20	37		
Profit	1	−1		
Bidder X	18	32	49	29
Profit	−2	3	8	6
Round 2				
Bidder Y	18	35		
Profit	−1	−3		
Bidder Z	18	34		
Profit	1	−3		
Bidder X	16	30	47	29
Profit	−1	1	6	6
Round 3				
Bidder Z	16	32		
Profit	−1	−5		
Bidder Y	16	33		
Profit	−3	−5		

Thus Stage II ends with bidder X winning the bundle $ABCE$, which includes a known 2-cycle and allows this firm to complete a 3-cycle with its existing contract on F , while bidder Y is awarded a contract on route D from its bid table value.

Bidders next enter any bid that is higher than the final probe price they receive at the end of Stage II, except for bidder X , who may submit any bid on any subset of $ABCE$ due to the winning Stage II bid on this bundle. The auctioneer then determines the efficient solution in Stage III and computes VCG prices for each awarded bundle contract. Bidder X wins CE with a bid of 23 and a VCG payment of 34; bidder Y wins D with a bid table entry of 16 and a VCG payment of 16, and bidder Z wins AB with a bid of 17 and a VCG payment of 18. Using the technique described in [13] one can show that the VCG payments are in the core, and thus need not be adjusted.

In this example we have briefly demonstrated the overall flow of the three-stage auction, from bid tables to probing for bundles to a final sealed-bid round. The example shows that some bidders may experience a profit while others may only be able to win a contract by bidding away any visible profit. It also shows the need for the final sealed-bid round, with Stage II ending at a less than efficient solution, and shows that with forthcoming bidders the honesty constraints do not prohibit the bids necessary to achieve efficiency.

8. Concluding remarks

We have presented a new auction design that incorporates features from both the PAUSE auction of Kelly and Steinberg [19] and the clock–proxy auction of Ausubel et al. [4]. In particular, our method emphasizes revelation of price information on individual items in order to reduce the number of bundles for which bids will be necessary in the (NP -hard) winner-determination problem, and to allow for the revelation of aggregate information in common value settings. Unlike the clock–proxy auction and the dual-based pricing approaches, our mechanism for revealing prices on individual items does not incorporate non-linear (or non-additive) synergy information into linear item prices, as this leads to meaningless and distorted item prices. Unlike the PAUSE auction, free-riding is limited by a disciplined approach to bundle price revelation. We do, however, reveal price information over bundles, and as in the PAUSE auction this revelation takes place with only little computational burden to the auctioneer. (In contrast to PAUSE, this takes place through the repeated solution of assignment problems in our case, a polynomial-time procedure.) Like the clock–proxy auction, the use of a final sealed-bid

round ensures efficiency and allows us to encourage truthful bidding throughout the auction with a “second-price” final payment mechanism.

With the increasing use of combinatorial auctions for procurement (and other settings) the design of a computationally viable auction that incorporates price discovery (both for items and bundles), while mitigating many of the problems such as free-riding and the threshold problem, is of significant practical interest. The auction proposed in this paper addresses these issues in a novel and comprehensive way, building upon both PAUSE and the clock–proxy auction, yet avoiding some of the criticisms of each. To our knowledge, no other auction proposed so far is able to address all of the issues addressed by our auction. We look forward to experiments and practical implementations of our auction format in the future, in both the forward and reverse (procurement) auction contexts.

Acknowledgments

The authors gratefully acknowledge the summer research support from their respective institutions, as well as the National Science Foundation, award number DMI-0205489.

Appendix A

Proof of Theorem 5.1

Proof. We work in the forward auction setting for consistency with the literature to which we are referring. The problem of allocating items according to demand correspondences is as follows: given a set of demand correspondences $D_1, D_2, \dots, D_M$, each a set of subsets of the set of all items I , select one subset S_j from each D_j such that $S_j \cap S_{\bar{j}} = \emptyset$ whenever $j \neq \bar{j}$ and $\bigcup_{j \in J} S_j = I$. It is easy to verify that this problem is in *NP*. Given a set of S_j s we need only verify that $S_j \cap S_{\bar{j}} = \emptyset$ for all $\frac{M(M-1)}{2}$ pairings of distinct bidders and that $\bigcup_{j \in J} S_j = I$. Clearly these verifications can occur in polynomial time. To show that the problem is *NP*-complete, we transform a generic set partitioning feasibility problem into an instance of the allocation feasibility problem. Given a set of objects X and a family of subsets of X , $F = \{X_1, X_2, \dots, X_p\}$, the set partitioning feasibility problem is to find a subset F' of F such that $X_j \cap X_{\bar{j}} = \emptyset$ for any distinct X_j and $X_{\bar{j}}$ in F' and $\bigcup_{X_j \in F'} X_j = X$. To transform this into an instance of the demand correspondence allocation feasibility problem simply let $I = X$ and create a bidder j for each X_j where j 's demand correspondence = $\{\emptyset, X_j\}$. $\square$

References

- [1] G. Anandalingam, R. Day, S. Raghavan, The landscape of electronic market design, *Management Science* 51 (3) (2005) 316–327.
- [2] L. Ausubel, An efficient dynamic auction for heterogeneous commodities, *American Economic Review* 26 (3) (2006) 602–629.
- [3] L. Ausubel, P. Cramton, The optimality of being efficient. Working paper, University of Maryland, Jun 1999.
- [4] L. Ausubel, P. Cramton, P. Milgrom, The clock–proxy auction: a practical combinatorial auction design, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), *Combinatorial Auctions*, chapter 5, MIT Press, 2006, pp. 115–138.
- [5] L. Ausubel, P. Milgrom, Ascending auctions with package bidding, *Frontiers of Theoretical Economics* 1 (1) (2002) 1–42.
- [6] L. Ausubel, P. Milgrom, The lovely but lonely vickrey auction, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), *Combinatorial Auctions*, chapter 1, MIT Press, 2006, pp. 17–40.
- [7] S. Ba, J. Stallaert, A. Whinston, Optimal investment in knowledge within a firm using a market mechanism, *Management Science* 47 (9) (2001) 1203–1219.
- [8] M. Ball, G. Donohue, K. Hoffman, Auctions for the safe, efficient and equitable allocation of national airspace systems resources, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), *Combinatorial Auctions*, chapter 22, MIT Press, 2006, pp. 507–538.
- [9] J. Banks, J. Ledyard, D. Porter, Allocating uncertain and unresponsive resources: an experimental approach, *Rand Journal of Economics* 20 (1) (1989) 1–25.
- [10] P. Cramton, J. Schwartz, Collusive bidding: lessons from the FCC spectrum auctions, *Journal of Regulatory Economics* 17 (2000) 229–252.
- [11] R. Day, Expressing Preferences with Price-Vector Agents in Combinatorial Auctions. PhD thesis, Applied Mathematics and Scientific Computation, University of Maryland, College Park, 2004. url: <http://users.business.uconn.edu/bday/thesis.pdf>.
- [12] R. Day, S. Raghavan, Assignment preferences and combinatorial auctions, Working paper, University of Connecticut, Apr 2006, url: <http://users.business.uconn.edu/bday/AsP.pdf>.
- [13] R. Day, S. Raghavan, Fair payments for efficient allocations in public sector combinatorial auctions, *Management Science* (2007) (to appear).
- [14] S. de Vries, R. Vohra, Combinatorial auctions: A survey, *INFORMS Journal on Computing* 15 (3) (2003) 284–309.
- [15] G. Demange, D. Gale, M. Sotomayor, Multi-item auctions, *Journal of Political Economy* 94 (4) (1986) 863–872.
- [16] F. Gul, E. Stachetti, The English auction with differentiated commodities, *Journal of Economic Theory* 92 (1999) 66–95.
- [17] K. Hoffman, D. Menon, S. van den Heever, T. Wilson, Observations and near-direct implementations of the ascending proxy auction, in: P. Cramton, Y. Shoham, R. Steinberg (Eds.), *Combinatorial Auctions*, chapter 17, MIT Press, 2006, pp. 415–450.
- [18] G. Hohner, J. Rich, E. Ng, G. Reid, A. Davenport, J. Kalaganam, H.S. Lee, C. An, Combinatorial and quantity-discount procurement auctions benefit mars, incorporated and its suppliers, *Interfaces* 33 (1) (2003) 23–35.
- [19] F. Kelly, R. Steinberg, A combinatorial auction with multiple winners for universal service, *Management Science* 46 (4) (2000) 586–596.
- [20] V. Krishna, *Auction Theory*, Academic Press, 2002.
- [21] A. Kwasnica, J. Ledyard, D. Porter, C. DeMartini, A new and improved design for multi-object iterative auctions, *Management Science* 51 (3) (2005) 419–434.

- [22] R. Kwon, G. Anandalingam, L. Ungar, Iterative combinatorial auctions with bidder-determined combinations, *Management Science* 51 (3) (2005) 407–418.
- [23] J. Ledyard, M. Olson, D. Porter, J. Swanson, D. Torma, The first use of a combined value auction for transportation services, *Interfaces* 32 (5) (2002) 4–12.
- [24] B. Lehmann, D. Lehmann, N. Nisan, Combinatorial auctions with decreasing marginal utilities, *Games and Economic Behavior* 55 (2) (2006) 270–296.
- [25] Logistics.com. Logistics.com named to forbes magazine's "best of the web: B2b" list for second straight year, 2001.
- [26] G. Nemhauser, L. Wolsey, *Integer and Combinatorial Optimization*, John Wiley and Sons, 1999.
- [27] P. Wurman, J. Zhong, G. Cai, Computing price trajectories in combinatorial auctions with proxy bidding, *Electronic Commerce Research and Applications* 3 (4) (2004) 329–340.



Robert W. Day is an Assistant Professor of Operations and Information Management in the School of Business at the University of Connecticut. He received his Ph.D. in Applied Mathematics with a concentration in Operations Research from the University of Maryland, College Park in 2004. His dissertation focused on combinatorial auctions received INFORMS' Dantzig dissertation award in 2005. Day continues to study combinatorial auctions and other OR applications, including Healthcare OR.



S. "Raghu" Raghavan is Associate Professor of Management Science at the Robert H. Smith School of Business, University of Maryland, with a joint appointment at the Institute for Systems Research. He has published on a wide variety of topics (including telecommunications, auctions & electronic markets, and data mining) and in numerous academic outlets such as *Management Science*, *Operations Research*, *Decision Support Systems*, and the *INFORMS Journal on Computing*. He holds two patents, and has won several awards for his work. These include (i) the Dantzig award for the best doctoral dissertation, (ii) the INFORMS Computing Society Prize for innovative contributions to the field of data mining, and (iii) the Glover-Klingman prize for the best paper published in the journal *Networks*. He enjoys teaching optimization and decision modeling courses, and was awarded the Legg-Mason Teaching Innovation award in 2007 by the Smith School. Dr. Raghavan obtained his Ph.D. in Operations Research from the Massachusetts Institute of Technology.