



Models for Iterative Multiattribute Procurement Auctions

The Harvard community has made this article openly available. [Please share](#) how this access benefits you. Your story matters

Citation	Parkes, David C., and Jayant Kalagnanam. 2005. Models for iterative multiattribute procurement auctions. <i>Management Science</i> 51(3): 435-451.
Published Version	doi:10.1287/mnsc.1040.0340
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:4039783
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Models for Iterative Multiattribute Procurement Auctions

David C. Parkes

Division of Engineering and Applied Sciences, Harvard University, 229 Maxwell-Dworkin, 33 Oxford Street, Cambridge, Massachusetts 02138, parkes@eecs.harvard.edu

Jayant Kalagnanam

IBM Research Division, T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598, jayant@us.ibm.com

Multiattribute auctions extend traditional auction settings to allow negotiation over nonprice attributes such as weight, color, and terms of delivery, in addition to price and promise to improve market efficiency in markets with configurable goods.

This paper provides an iterative auction design for an important special case of the multiattribute allocation problem with special (preferential independent) additive structure on the buyer value and seller costs. Auction ADDITIVE&DISCRETE provides a refined design for a price-based auction in which the price feedback decomposes to an additive part with a price for each attribute and an aggregate part that appears as a price discount for each supplier. In addition, this design also has excellent information revelation properties that are validated through computational experiments. The auction terminates with an outcome of a modified Vickrey-Clarke-Groves mechanism. This paper also develops Auction NONLINEAR&DISCRETE for the more general nonlinear case—a particularly simple design that solves the general multiattribute allocation problem, but requires that the auctioneer maintains prices on *bundles* of attribute levels.

Key words: multiattribute negotiation; iterative auctions; price-based feedback; Vickrey-Clarke-Groves mechanism; ex post Nash equilibrium; straightforward bidding; procurement

History: Accepted by G. Anandalingam and S. Raghavan, special issue editors; received June 6, 2002. This paper was with the authors 8 months for 3 revisions.

1. Introduction

Multiattribute auctions extend the traditional auction setting to allow negotiation over price and attributes. For example, in a procurement problem, a multiattribute auction can allow different suppliers to compete over both attributes values and price. An *iterative* multiattribute auction determines the outcome dynamically, with agents revising bids in response to feedback, and can usefully reduce the amount of information revealed by agents in the auction. The design of iterative multiattribute auctions forms the focus of this paper.

The use of iterative multiattribute auctions is becoming prevalent in procurement for both services and goods. An interesting example is IT outsourcing. Such contracts need to handle multiple attributes such as: (i) the fraction of the service that is offshored; (ii) the one-time transition cost for the transfer of the operations, and (iii) labor rates for offshore and onshore services. In such auctions, the transition costs are very dependent on the fraction that is offshored, which together with the labor costs contributes to the total cost of a contract. Another interesting example

is the procurement of coal for electric utilities, where multiple attributes such as the heat content (Btu/lb) and the sulphur content (lb/MBtu) directly impact the value of the coal. Suppliers often have the option of blending different types of coals or treating coal to vary these attributes. This allows suppliers to trade off these attributes, depending on their cost structure for these treatments and the value structure of buyers.

Iterative auctions, which allow agents to revise their bids and provide incremental information about preferences, have several advantages over one-shot auctions for procurement problems. First, it is important for bidders to reveal as little information as possible about costs and preferences in a strategic situation such as procurement, because participants are in a long-term competitive relationship. Iterative auctions can elicit cost information from the suppliers on a pure “need to know” basis, and the buyer cannot precisely infer the cost function of sellers. Iterative auctions are also known to outperform sealed-bid auctions in settings in which costs are correlated across sellers, because dynamic feedback provides information to help participants to revise their beliefs

about their own preferences for different outcomes (Milgrom and Weber 1982, Cramton 1998).

The primary contribution of this paper pertains to an auction design for an important special case of the multiattribute allocation problem (MAP) with special preferential independent (PI) additive structure on the buyer value and seller costs. Auction ADDITIVE&DISCRETE (AD) quotes prices in terms of an *additive* component, with a price for each attribute level, and an *aggregate discount*, that applies to an entire bid. The aggregate discount is required whenever the best seller does not dominate (in terms of cost) the other sellers on every attribute. The compact price space in Auction AD provides good information-revelation properties, as validated through computational experiments. As a prelude to developing Auction AD, we also develop Auction NONLINEAR&DISCRETE (NLD). Auction NLD is very general and does not require PI, but requires the auctioneer to maintain prices on bundles of attribute levels.

The PI assumption is quite standard in multiattribute utility theory (Keeney and Raiffa 1993). We find PI to be a compelling model for a class of multiattribute problems in procurement. One common setting is in the context of the procurement of commodity items such as sugar, for which an expressive market already exists (Hohner et al. 2003). In such contexts suppliers usually place bids with the market price as a base price, and with an add-on price for each additional attribute, such as the degree of refinement provided, the location to be delivered to, delivery dates, etc. The price is specified independently for each attribute. Another common setting for PI is for configurable goods such as PCs. A typical desktop has many attributes such as the memory, the processor speed, and the hard drive, each with multiple choices. A very common price structure for these goods is a markup-based price, where configurations are priced with a base price and an add-on price for each attribute (Bichler and Kalagnanam 2003). Auction AD applies when PI holds, and provides particularly useful preference-elicitation properties.

We use linear programming duality theory to design and analyze Auction AD. The auction terminates with the outcome of a modified Vickrey-Clarke-Groves (VCG) (Vickrey 1961, Clarke 1971, Groves 1973) auction. From this, it follows that straightforward strategies form an ex post Nash equilibrium for sellers, against a class of nonadaptive (but not necessarily truthful) buyer strategies. This ex post Nash solution concept is quite robust, because sellers do not need to be informed about the costs of other sellers to follow their equilibrium strategy. We also provide a simple, but useful, bound on the maximal increase in payoff that a buyer can achieve by adopting an adaptive and nontruthful strategy over a truthful strategy. We present computational experiments

that verify the information revelation properties of the auction models. For simulated valuation and cost functions, we compare the amount of information revealed by participants in Auctions NLD and AD. The results demonstrate that the compact additive price space in Auction AD can provide a significant reduction in information revelation over NLD, which in turn can provide a significant saving in information revelation over a one-shot auction.

1.1. Related Work

Che (1993) first studied multiattribute auctions as a model for procurement within the supply chain. Multiattribute auctions have also been studied in the context of bargaining over shared resources between distributed computational agents (Kraus 1997, Jennings et al. 2001).

Early designs emphasized the design of an (buyer) *optimal* auction, to maximize the expected total payoff to the buyer by leveraging beliefs about the costs of sellers. Che (1993) proposed optimal one-shot (sealed-bid) auctions, for a model in which the cost functions of sellers are defined in terms of a single unknown parameter. A buyer provides a *scoring function*, and sellers respond in equilibrium by choosing to supply at a quality level that is efficient given the scoring rule (which itself is not truthful in equilibrium). Che derives an equilibrium in which the buyer states an optimal scoring function. Recently, Branco (1997) extends Che's auction to the case where the seller cost functions are correlated.

The rules in Che's "second-score" auction are those of the one-sided VCG auction (see §2) that forms the basis of our iterative price-based auctions. Rather than focus on the problem of optimal auction design, we consider the problem of *efficient* auction design. We consider the goal of market efficiency to be well suited for the design of stable long-term markets that will form the basis for repeated trade. We expect that efficient markets will come to dominate the electronic market landscape based on our experience with procurement auctions (Hohner et al. 2003). In long-term contract negotiations, the number of suppliers that a company interacts with is very small (typically of the order of 5 to 10), and inefficient allocations across this pool leads suppliers to question the credibility of the buyer to be fair. Even in business-to-business settings this emerges as one of the most important requirements, as reported in the deployments with a large chocolate manufacturer (Hohner et al. 2003). Buyer-optimal auctions are perhaps more appropriate for a one-shot procurement problem, and in a setting in which the buyer has considerable market power. Turning to efficient design also allows a more general model, in which sellers can have an arbitrary cost function. Indeed, optimal auctions are not known even for the special case of PI.

Iterative multiattribute auctions have been considered in prior work. Beil and Wein (2003) propose an iterative variant on Che’s auction for a richer class of parameterized utility functions (this time with K parameters) with known functional forms. The buyer uses K rounds to estimate the sellers’ cost functions, restarting the auction with a different scoring function each time. Sellers are modeled as naive and truthful agents, which allows the buyer to determine the exact seller cost function. For the final $(K + 1)$ st round, Beil and Wein design a scoring function so as to maximize buyer payoff by essentially reporting the same score (within ϵ) for the top two suppliers. Vulkan and Jennings (2000) propose a multi-round efficient auction. Their design differs from our design in that it is not price based, and also because there is no special optimization for the preferential-independence special case.

Recent literature adopts a linear programming approach for the design of iterative *combinatorial auctions*, in which bidders demand different combinations of items (Parkes and Ungar 2000, Bikhchandani et al. 2001, Bikhchandani and Ostroy 2002, de Vries et al. 2003, Mishra and Parkes 2004). However, the Combinatorial Allocation Problem (CAP) and the MAP differ in important ways. First, the preferential-independence special case is well motivated for the MAP, but makes less sense for the CAP. Second, there is private information on both sides of the auction in the MAP. This complicates the auction design problem, because the winner-determination problem now depends on the preferences of the buyer in addition to the revealed bids from sellers. An aggregate price term, in combination with linear prices, has been used previously in Kwon et al. (2004) in the context of the CAP, with the aggregate price used to provide more nuanced price feedback to losing bidders.

1.2. Outline

Section 2 formulates the MAP, and defines a modified VCG auction that is one-shot, but provides a normative basis to guide the design of our iterative auctions. Section 3 introduces Auction NLD and presents theoretical analysis to characterize the performance of the auction. Section 4 defines Auction AD which has a smaller price space—composed of prices on individual attribute levels together with an aggregate discount term—and is applicable to the special case of PI. We also examine a special case in which the aggregate discount term is not required. Section 5 concludes with a computational study of the information revelation properties of the iterative auctions on stylized problems. All proofs are available in the online appendix.

2. The MAP

In the MAP there are n sellers, a single buyer, and the outcome is defined in terms of levels for each of m attributes and a winning seller. Each attribute $j \in \{1, \dots, m\}$ has a finite domain of *discrete* attribute values, Θ_j . For example, $\Theta_j = \{\text{red, yellow, green}\}$, where attribute j represents the color of an item. Let $\Theta = \Theta_1 \times \dots \times \Theta_m$ denote the joint space of attributes. Discrete attribute values are reasonable for the procurement of goods with discrete characteristics, such as *processor speed*, *delivery date*, and *color*. Naturally continuous characteristics such as *weight* and *heat content* must be discretized.

Each seller $i \in I = \{1, \dots, n\}$ has a *cost function*, $c_i(\theta) \geq 0$, for an attribute bundle $\theta \in \Theta$, and the buyer has a *valuation function*, $v(\theta) \geq 0$. We write $c = (c_1, \dots, c_n)$, and $c_{-i} = (c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n)$ to denote the costs without seller i . We assume a *private-values* model, with independently distributed seller costs and buyer value. A private-values model provides a reasonable first approximation for the procurement of goods, because seller costs can be expected to depend on her own local manufacturing base and sellers can be expected to be well informed about the cost of (upstream) raw materials. Later, in §4 we introduce the special case of PI, in which the costs are stated independently for each attribute.

As is standard in the auction literature, all participants are assumed to have quasi-linear utility functions, with utility $u_i(\theta, p) = p - c_i(\theta)$ to seller i for bundle θ at price p and utility $u^B(\theta, p) = v(\theta) - p$ to the buyer. We assume that the buyer will source from a single seller. We focus on the *efficient* multiattribute allocation, in which the objective is to clear the auction to maximize the total value to the buyer, net the seller’s cost:

$$\begin{aligned} \max_{x_i(\theta)} \quad & \sum_{i \in I} \sum_{\theta \in \Theta} x_i(\theta)(v(\theta) - c_i(\theta)) & [\text{MAP(I)}] \\ \text{s.t.} \quad & \sum_{i \in I} \sum_{\theta \in \Theta} x_i(\theta) \leq 1 \\ & x_i(\theta) \in \{0, 1\}, \quad \forall i, \forall \theta. \end{aligned}$$

Setting $x_i(\theta) = 1$ denotes that attribute bundle θ and seller i is selected in the outcome. Let $\Pi(v; c) = v(\theta^*) - c_{i^*}(\theta^*)$ denote the surplus generated in the *efficient outcome*, (θ^*, i^*) , that solves MAP(I). Define the *marginal product* $MP_i = \Pi(v; c) - \Pi(v; c_{-i})$, as the marginal value contributed by seller i to the economy. We will also refer to the efficient outcome (θ^*, i^*) as the *first-best* outcome, and denote the solution to MAP(I $\setminus i^*$) as $(\tilde{\theta}, \tilde{i})$, and refer to this as the *second-best* outcome.

The price-based multiattribute auctions defined in this paper implement the outcome of a modified VCG

mechanism. In the VCG mechanism, all agents submit bids to the auctioneer that: (a) determine the efficient trade; (b) compute payments so that each bidder's utility (with respect to her reported valuation) is equal to her marginal product MP_i , i.e., the total value that she contributes to the economy. The VCG mechanism is *strategy-proof*—with truthful bidding a dominant-strategy equilibrium—but it is well understood that it is not *budget balanced* in settings with two-sided private information such as the MAP. Instead, the VCG mechanism will require a payment to the winning seller that is greater than the payment made by the buyer. Indeed, no efficient mechanism can be budget balanced for the MAP (Myerson and Satterthwaite 1983, Krishna and Perry 2000).

We define a modified auction, the one-sided VCG auction, that retains incentives to support truthful bidding on the sell side while achieving budget balance. Truthful bidding is no longer an equilibrium strategy for the buyer in the one-sided auction.

One-Sided VCG Auction

(1) Each seller i bids a cost function \hat{c}_i and the buyer bids a valuation function \hat{v} .

(2) Outcome $(\hat{i}, \hat{\theta})$ is selected in winner determination, to solve MAP(I), based on bids \hat{c}_i and \hat{v} .

(3) The buyer pays $\hat{v}(\hat{\theta}) - \Pi(\hat{v}; \hat{c}_{-\hat{i}})$.

The payment made by the buyer is the VCG payment that the winning seller would receive in the VCG mechanism for this problem, i.e., $\hat{c}_i(\hat{\theta}) + (\Pi(\hat{v}; \hat{c}) - \Pi(\hat{v}; \hat{c}_{-\hat{i}})) = \hat{v}(\hat{\theta}) - \Pi(\hat{v}; \hat{c}_{-\hat{i}})$. The buyer's payment in the (full) VCG mechanism would be $\hat{v}(\hat{\theta}) - \Pi(\hat{v}; \hat{c})$. Here, the buyer makes an additional payment of $\Pi(\hat{v}; \hat{c}) - \Pi(\hat{v}; \hat{c}_{-\hat{i}})$. This provides budget balance, while continuing to support truthful bidding on the sell side. Voluntary participation is also provided: The winning seller's payment is at least her reported cost, and the payment made by the buyer is no greater than her bid price. The price of budget balance is that the auction is not strategy-proof for the buyer. However, we can bound the maximal gain that can be achieved by the buyer in comparison with her utility from reporting her true valuation.

PROPOSITION 1. *The ex post regret to the buyer from truthful bidding in the one-sided VCG auction, given straightforward seller strategies, is at most the marginal product MP_{i^*} , of the efficient seller (defined with respect to reported seller costs).*

To see that this upper bound is tight, consider an instance in which the cost of the second-best bundle, $\tilde{\theta}$, is less to the second-best seller, \tilde{i} , than to the first-best seller, i^* . In this case, the buyer can bid $\hat{v}(\theta^*) = v(\theta^*) + \{\Pi(v, c) - \Pi(v, c_{-i^*}) - \max(0, c_{\tilde{i}}(\tilde{\theta}) - c_{i^*}(\tilde{\theta}))\}$ for the efficient bundle θ^* , and bid truthfully for all other bundles. Given this, $\Pi(\hat{v}; c) = \Pi(\hat{v}; c_{-i^*})$,

and the efficient allocation is implemented, with the buyer taking all the surplus.

In practice, the opportunity to the buyer is more limited because she must bid without perfect knowledge of sellers' bids. Instead, if the buyer has information about the *distribution* from which seller's costs are drawn, then the buyer can play a Bayesian-Nash equilibrium and aim to maximize her expected utility. Trivially, the expected gain in utility over a truthful strategy in this Bayesian-Nash equilibrium is bounded above by $E_{c \sim F(c)}\{MP_{i^*}\}$, where the expectation is taken with respect to the distribution $F(c)$ on seller costs. To see this, notice that for any particular valuation v and costs c the best-case gain for the buyer is MP_{i^*} because the sellers will continue to bid truthfully in equilibrium. This term will be considerably less than the expected utility $E_{c \sim F(c)}\{\Pi(v; c_{-i^*})\}$ from truthful bidding in the usual case of strong competition.

3. Auction NLD: General Preferences

Our first auction NLD generalizes a single-item open-outcry auction, providing a kind of reverse English auction for the MAP. Prices are nonlinear, with combinations of attribute values priced explicitly, and rich enough to provide prices on features that are contingent on the selection of other features. Auction NLD provides a multiattribute auction for general preferences, and determines the efficient allocation without bidders revealing their full cost functions. However, Auction AD (§4), is able to take advantage of the special structure offered by PI and provides a compact price space that leads to more immediate feedback and better information-revelation properties than Auction NLD.

A high-level description of Auction NLD is provided in Figure 1. Auction NLD proceeds in rounds, and maintains an ask price, $p^t(\theta)$, on each bundle $\theta \in \Theta$ of attribute levels. At the start of the auction the buyer makes a claim about her valuation function, $\hat{v} \in \mathcal{V}$. The auctioneer uses this information to solve the winner-determination problem in each round, selecting as a provisional winner the bid that maximizes the buyer's utility given this reported valuation. The auction starts with high initial prices on

Figure 1 Auction NonLINEAR&DISCRETE (NLD)

```

collect a reported valuation,  $\hat{v}$ , from the buyer;
set high initial prices,  $p^1(\theta)$ , on each attribute
bundle  $\theta$ ;
while (active bidding) {
  collect bids on attribute bundles from sellers;
  determine the provisional allocation;
  decrease prices based on losing bids;
}
implement the final provisional allocation;

```

each attribute bundle. We assume that the buyer, or the auctioneer, has conservative prior bounds on the minimal cost for each possible configuration. Initial prices can also just be set so as to be greater than the reported values of the buyer.¹ Prices decrease across rounds, with price changes driven by bids from losing sellers.

In each round, a seller can bid at or below the ask price on one or more bundles or leave the auction. The winning bid from the previous round is automatically retained across rounds. For instance, if the first attribute defines *color* and includes $\Theta_1 = \{red, blue, green\}$, and the second attribute defines *speed* and includes $\Theta_2 = \{slow, medium, fast\}$, then $\{(red, fast, \$100), (blue, fast, \$120)\}$ is a valid bid as long as these bid prices are less than or equal to the ask prices. The new ask price, $p^{t+1}(\theta)$, for a bundle θ that receives a bid from a losing seller, is set to ϵ below the lowest losing bid, with $\epsilon > 0$ defined as the minimal bid increment in the auction. This auction parameter determines the rate of price changes in the auction.

The auction terminates when ask prices have not changed in two consecutive rounds. In most cases, the final price will be less than the buyer’s reported valuation, and the auction implements the final provisional allocation with the buyer making the final payment to the winning seller. However, when there is no efficient trade without the best seller—i.e., when $\Pi(\hat{v}, c_i) \leq 0$ —then the final price will remain above the buyer’s reported value. We handle this case by offering the provisionally winning seller a final price of $p(\theta^*) = \hat{v}(\theta^*)$ (which would be the outcome of the one-sided VCG auction), where θ^* is the bundle in the final provisional allocation. If the seller accepts this offer, then the auction closes with this outcome. Otherwise, the auction terminates with no trade.

3.1. Theoretical Analysis

Auction NLD terminates with the outcome of the modified VCG auction for sellers, which brings *straightforward bidding* into an ex post Nash equilibrium. Straightforward bidding, or myopic best response, is defined as follows:

DEFINITION 1. A seller’s myopic best-response (MBR) bidding strategy, $mbr(\hat{c}_i)$, is to bid at the ask price for all bundles that have nonnegative utility and are within ϵ of maximizing her utility, given current ask prices, and for some cost function \hat{c}_i .

¹ Technically, correct convergence is guaranteed whenever initial prices are at least the lowest cost across all sellers on the bundle plus the marginal product of the winning seller MP_{i^*} (to support the VCG outcome).

In straightforward bidding a seller is myopic, bidding as though the current auction prices are final, and ignoring the effect of her bid on future prices. We establish the following equivalence between Auction NLD and the one-sided VCG auction:

PROPOSITION 2. *Auction NLD with straightforward bidding terminates with the efficient outcome and the payment in the one-sided VCG auction for a truthful buyer, and as the minimal bid increment ϵ goes to zero.*

This equivalence brings straightforward bidding into an ex post Nash equilibrium for sellers, in the sense that best response to prices in every round is the optimal strategy for a seller whatever the costs of other sellers and whatever the reported valuation of the buyer, as long as the other sellers also follow a straightforward strategy. Ex post Nash is a robust solution concept because sellers need not be informed about the costs of other sellers to compute their best response.²

Formally, let $g(s^*(c))$ denote the outcome of an auction when bidders with costs $c = (c_1, \dots, c_n)$ follow straightforward bidding strategy $s^*(c)$ and the outcome of the auction (an allocation and payments) is defined with outcome rule $g(s(c))$, given joint strategy s and costs c . Then, strategy s^* is an ex post Nash equilibrium if

$$u_i(g(s^*(c))) \geq u_i(g(s', s^*(c_{-i}))) \quad (1)$$

for all bidders i and all costs c_{-i} and all costs c_i , where s' is an arbitrary bidding strategy and $u_i(g(s(c))) \in \mathbb{R}$ is the utility to bidder i for an outcome.

THEOREM 1. *Truthful MBR is an ex post Nash equilibrium for sellers in Auction NLD as the minimal bid increment goes to zero, whatever the reported valuation \hat{v} of the buyer.*

We also have the following immediate corollary:

COROLLARY 1. *Auction NLD is efficient when the buyer bids truthfully, and the maximal benefit to a buyer from some nontruthful strategy is no greater than the marginal product of the efficient seller.*

3.2. Competitive Equilibrium Prices

To complete the analysis of Auction NLD, we define *competitive equilibrium* (CE) prices and show the

² Ex post Nash makes a weaker knowledge assumption than that required for a Bayesian-Nash equilibrium, in which the distribution on seller costs must be common knowledge. However, ex post Nash makes a stronger knowledge assumption than that required for the dominant-strategy equilibrium in the one-sided VCG auction, in which sellers need not even believe that other sellers will be rational.

auction terminates with *maximal* CE prices that support the outcome of the one-sided VCG auction.

As is standard, we say that prices $p(\theta)$ and feasible MAP solution (θ', i') are in CE if bundle θ' maximizes the utility for the buyer and the winning seller i' at the prices, and the prices are less than cost on all bundles for all other sellers. It is easy to see that CE prices always exist; for instance, consider the instance of prices $p(\theta) = \min_i \{c_i(\theta)\}$, together with the efficient allocation. Moreover, by linear-programming duality we see that the CE prices support an efficient outcome.

Let π_i denote the *maximal payoff* to seller i at prices p , defined as $\pi_i = \max_{\theta \in \Theta} [p(\theta) - c_i(\theta), 0]$. Let π^B denote the *maximal payoff* to the buyer at prices p , defined as $\pi^B = \max_{\theta \in \Theta} [v(\theta) - p(\theta), 0]$. Formally, we require $\pi^B = v(\theta') - p(\theta')$ and $\pi_{i'} = p(\theta') - c_{i'}(\theta')$, with $\pi_i = 0$ for all $i \neq i'$ in CE.

DEFINITION 2. The maximal CE prices, \bar{p}_{ce} , maximize the price on the efficient bundle, θ^* , across all CE prices.

Maximal CE prices provide the winning seller with the best possible revenue across all prices that support the efficient trade in equilibrium. Intuitively, the seller cannot receive a greater payment without some other seller being able to step in and make the buyer an offer that she will prefer. Maximal CE prices can be constructed by adjusting CE prices to maximize $p(\theta^*)$ while maintaining the conditions required for CE.

LEMMA 1. Prices $\bar{p}_{ce}(\theta^*) = c_{i^*}(\theta^*) + MP_{i^*}$, with $\bar{p}_{ce}(\theta') = \min_{i \neq i^*} \{c_i(\theta')\}$ on all $\theta' \neq \theta^*$ are maximal CE prices.

Given this characterization, the following equivalence is immediate.

PROPOSITION 3. The payments in the one-sided VCG auction are implemented in the maximal CE.

Notice that price $\bar{p}_{ce}(\theta^*) = c_{i^*}(\theta^*) + MP_{i^*}$ is exactly the payment to the winning seller in the one-sided VCG auction.

In what follows, prices p are said to ϵ -dominate prices p' if $p(\theta) + \epsilon \geq p'(\theta)$ for all θ . Also, we define CE prices with respect to the reported \hat{v} of the buyer.

LEMMA 2. Auction NLD maintains ask prices that ϵ -dominate the maximal CE prices when all sellers follow a truthful MBR bidding strategy.

Auction NLD must terminate when agents are rational, because while the auction remains open the price $p^t(\theta)$ falls in each round on at least one attribute bundle, θ , for which a seller has submitted a bid. To keep bidding, this seller must have $\hat{c}_i(\theta) \leq p^t(\theta)$, which places a lower bound on the minimal price that can support bids. Termination follows because we assume that the number of agents, the number of

attributes, and the number of attribute levels are all finite. Finally, Auction NLD terminates in the maximal CE prices.

LEMMA 3. Auction NLD terminates with maximal CE prices when sellers follow a truthful MBR strategy, and as the minimal bid increment goes to zero.

3.3. Number of Rounds to Terminate in Auction NLD

It is interesting to characterize the number of rounds that Auction NLD can take to reach termination. Let m denote the number of attributes, and let $l = \max_j |\Theta_j|$ denote the maximal number of discrete levels of any single attribute. The size of the attribute bundle space is $O(l^m)$. Let $V_{\max} = \max_{i \leq n} [\max_{\theta} p^1(\theta) - c_i(\theta)]$, where $p^1(\theta)$ is the initial ask price on bundle θ . The number of rounds in NLD is polynomial in $V_{\max}/1/\epsilon$, but exponential in l and m .

PROPOSITION 4. Auction NLD converges in $O(l^m V_{\max}/\epsilon)$ rounds, with minimal bid increment ϵ and MBR strategies.

As a result of the exponential price space, the number of rounds is worst-case exponential in m , the number of attributes. In contrast, Auction AD, which is applicable for the special case of PI, has a linear price space and will terminate after a number of rounds that scales as $O(lm)$.

3.4. Adaptive Buyer Strategies

The basic version of Auction NLD, as described, requires the buyer to commit to a valuation function $\hat{v} \in \mathbb{V}$ at the start of the auction. In a simple variation, we can allow the buyer to bid *incrementally*, providing the minimal amount of additional information that is needed in each round to determine the winner and to adjust prices. For instance, the buyer could be presented with the set of bids received from sellers, and asked which bid it prefers. Sunderam and Parkes (2003) considered this approach, providing a proxy agent to intermediate between the bidders and the auctioneer. A response to a new query places additional constraints on the buyer's valuation. The proxy can use this partial information to respond when it has enough information, and fall back to the buyer and ask for additional information when necessary.

In practice, making the auction dynamic on the buy side would require striking a balance between the need to reduce information revelation from the buyer with a need to protect the interests of sellers. Incremental bidding provides new opportunities for strategic behavior by buyers. In particular, a buyer can use sellers' bids to adapt her strategy and improve her final payoff. We already saw in §2 that a well-informed buyer can do better than truth revelation. A new concern for sellers is that bids during the auc-

tion can reveal information that can help the buyer to adapt her strategy. An auction that permits dynamic buyer strategies as well as seller strategies should at least ensure that the buyer's revealed preference information is *consistent* across rounds.

4. Auction AD: Preferential Independence

We now consider the special case of PI, in which the cost of each seller can be stated separately for each attribute, with the total cost defined as the additive sum across attributes. Preferential independence is quite standard in multiattribute utility theory (Keeney and Raiffa 1993), and is relevant in our setting of procurement auctions whenever the value and cost is separable across attributes. We provide some motivating scenarios below.

We introduce Auction AD, which has a linear price space and maintains prices on the individual level of each attribute as well as an aggregate discount price that applies to an entire bid. The total ask price for a bundle of attribute levels is calculated as the sum of the level prices minus this discount term. The aggregate discount term is necessary when the efficient supplier does not dominate the other suppliers on every attribute. We interpret Auction AD as implementing a primal-dual algorithm for the MAP, with straightforward bidding leading to an efficient solution. In §5, we show that the auction solves problems with less information elicitation from participants than Auction NLD.

The outline of this section is as follows. We first define PI. We then describe the auction, state its main theoretical properties, and provide an extended example for a simple problem. Finally, we define CE prices and complete the formal theoretical analysis, again demonstrating an equivalence between maximal CE prices and the one-sided VCG auction.

4.1. Preferential Independence

PI imposes the additional assumption that a seller's cost for a level of an attribute does not depend on the levels set on the other attributes, and similarly for the value of the buyer. We find PI compelling for a class of multiattribute problems in procurement, including: for *commodity items*, where the attributes (refinement level, location, delivery dates, etc.) can be readily assigned independent costs and values; and for *configurable goods* such as PCs, where the attributes (memory, processor speed, etc.) can be readily assigned independent costs and values.

Formally, we now define an attribute bundle θ in terms of $x = (x_1, \dots, x_m)$. For each attribute j , we have $x_j \in \{0, 1\}^{|\Theta_j|}$, and $\sum_k x_{jk} \leq 1$ so that at most one level is selected. Setting $x_{jk} = 1$ indicates that level k of

attribute j is selected. The cost function for seller i can now be expressed as:

$$c_i(\theta) = \sum_{j \leq m} \sum_{k \leq |\Theta_j|} c_{ijk} x_{jk}, \quad (2)$$

where c_{ijk} is the marginal cost to seller i if level k of attribute j is selected. The valuation function for the buyer can now be expressed as:

$$v(\theta) = \sum_{j \leq m} \sum_{k \leq |\Theta_j|} v_{jk} x_{jk}, \quad (3)$$

where v_{jk} is the marginal value to seller i if level k of attribute j is selected. With this, the MAP problem can be formulated as:

$$\max_{x_{ijk}, y_i} \sum_{i \leq n} \sum_{j \leq m} \sum_{k \leq |\Theta_j|} (v_{jk} - c_{ijk}) x_{ijk} \quad [\text{MAP-PI}]$$

$$\text{s.t.} \quad \sum_{k \leq |\Theta_j|} x_{ijk} \leq y_i, \quad \forall i, \forall j \quad (4)$$

$$\sum_{i \leq n} y_i \leq 1 \quad (5)$$

$$x_{ijk}, y_i \in \{0, 1\},$$

where $y_i = 1$ if seller i is the winner, and $x_{ijk} = 1$ if level k of attribute j is supplied by this seller.

We allow the buyer to opt out of one or more attributes whenever the winning seller cannot supply any level of that attribute at a cost below value. This is common practice in pricing schemes used for services or configurable goods. For example, desktops have multiple attributes (memory, processor, hard drive, etc.), and the most common pricing scheme is based on using a base price and markup. In such situations, the opt-out option for one or more attributes (say memory) is equivalent to the buyer choosing the default level in the base model (e.g., 128 MB).

4.2. Auction AD

A high-level description of Auction AD is provided in Figure 2. Prices are basically linear additive on attribute levels, but with an additional (anonymous) discount that applies to an entire bid. Auction AD maintains prices $p_{jk}^t \geq 0$ on individual levels k of each attribute j in round $t \geq 1$, along with price discount $\Delta^t \geq 0$ that applies to the total price on any bundle of attribute levels. The overall ask price on bundle $x = (x_1, \dots, x_m)$ is defined as

$$p^t(x) = \left(\sum_{j \leq m} \sum_{k \leq |\Theta_j|} p_{jk}^t x_{jk} \right) - \Delta^t.$$

The role of the discount is to support the efficient allocation in a price equilibrium when the best seller does not dominate (in terms of buyer value and seller cost) the other sellers on every attribute. For instance,

Figure 2 Auction ADDITIVE&DISCRETE (AD)

```

collect a reported valuation,  $\hat{v}$ , from the buyer;
set high initial prices,  $p_{jk}^1$ ;
initialize price discount  $\Delta^1 = 0$ ;
while (active bids) {
  collect bids from sellers;
  determine the provisional allocation;
  determine the most-preferred levels on each
  attribute;
  decrease level prices,  $p_{jk}^t$ ;
  increase the discount  $\Delta^t$  if stalled;
}
implement the final provisional allocation;

```

one supplier of IT outsourcing might be preferred by the buyer on *offshore labor rates* and *professional services*, while another supplier might be preferred in terms of *minimal response times*.

The auctioneer sets initial level prices to be some value greater than the buyer's reported level values, and initializes the price discount to zero. The auction is parameterized with a minimal bid increment, $\epsilon > 0$, which determines the rate at which prices are decreased across rounds. At the start of the auction the buyer provides a reported valuation function, \hat{v} , to the auctioneer.³ In each round, bids are collected from sellers, with the winning bid from the previous round resubmitted automatically (at the previous bid price). This mimics the dynamics in the English auction, where the current winner does not need to match the new ask price, but is allowed to let her bid stand until competing bidders match the new ask price.

Each seller can bid on one or more levels on each attribute, with a bid price that matches or improves on the current ask price on that level. A seller can also skip an attribute altogether, and might choose to when the price is above cost across all levels. A seller must always match, or improve on, the current price discount asked in the auction. The full discount applies, even when the supplier is only bidding on a subset of the attributes. A seller can also exercise a *last-and-final* bid on an attribute in any round. This allows the seller to bid ϵ above the ask price on any level of that attribute. However, once exercised, a seller cannot improve her bid on any level on that attribute in future rounds, although she can still improve her bids on other attributes, and she can still offer a larger discount.

When selecting a provisional winner, the auctioneer (on behalf of the buyer), must choose a level from within each attribute included in the bid. For instance, in a setting with a *color* attribute and a *speed* attribute,

a typical bid might state $\{(red, \$50), (yellow, \$80)\}$, $\{(fast, \$100), (slow, \$20)\}$, with *discount* \$40. In response, the buyer might consider combining $(red, slow)$ for a total price of $\$50 + \$20 - \$40 = \30 . A *provisional allocation* is determined in each round to maximize the buyer's utility given her stated valuation and given current bids. Ties are broken in favor of the current provisional winner.

Price-Update Rules. In order to describe the price-update rules, we first need a language to talk about characteristics of attributes, levels, and sellers. First, we need to distinguish between attributes that are *in play* and *not in play*. To be in play, at least one bid must be received on the attribute at a level price below the buyer's reported value for that level (the discount in a bid is ignored in this characterization). We also characterize a level within an attribute as a *most-preferred* level, which is a level that is within ϵ (the price increment in the auction) of maximizing the buyer's utility at the current ask prices. Finally, a seller is said to be *active* on an attribute if she bids at or below the ask price on one or more levels. A seller is said to be *active overall* if she is either the current winner in the auction, or she is active on one or more attributes.

Each attribute is then considered in turn, with the price-update rule selected to depend on whether or not the attribute is in play:

Not In Play. Consider two subcases.

(*not-a*) If a losing seller does not bid on *any* most-preferred level for this attribute, then set the ask price on that level to ϵ below the bid price from this seller or the current ask price, whichever is smaller.

(*not-b*) If all bids from sellers include a most-preferred level and at least one losing seller is active on the attribute, then set the new ask price on each level to ϵ below the lowest bid price from any seller on this attribute.

In Play. (*in*) Set the new ask price on any level that receives a bid to ϵ below the lowest bid price received, while all active sellers are also active on this attribute.

The price discount is also adjusted according to the following rule:

Discount. If ask prices on attribute levels have remained unchanged for two successive rounds, and there are still two active sellers, then increase the *discount* to ϵ above the largest discount that was bid in the last round by a losing seller.

Rule *not-a* is used to decrease prices on levels on attributes that are not in play and that receive bids from sellers that are not competitive on the attribute. Rules *not-b* and *in* are used to drive competition between sellers that might be competing on different levels on an attribute. Note that both of these rules consider bids from the provisional winner when adjusting prices. Rule *not-b* requires that all bids are

³ As in Auction NLD, the auction can also be operated with incremental bidding from the buy side. The buyer must provide enough information in each round to guide winner determination and price adjustment.

on most-preferred levels, but does not require that all active sellers are active on the attribute. Rule *in*, on the other hand, does require that all active sellers are active on this attribute. Rule *discount* provides competitive price pressure between sellers that are competing on different attributes (which precludes the use of the standard in-play rule because they would not both be active on the same attribute).

In the usual case, the price offered by the winning seller at the end of the auction is less than the buyer’s reported value on all attributes and the auction implements the final allocation at this price. On the other hand, if an attribute is not in play, then the auction considers the level on which the winner is most competitive and offers the winner a final price equal to the buyer’s reported value on that level. The winner can either accept, or choose to opt out of this attribute.

We define a straightforward MBR strategy for a seller in the auction as:

DEFINITION 3. A seller’s MBR strategy, $mbr(\hat{c}_i)$, is to bid on all levels of each attribute that are within ϵ of maximizing her utility while the total price is greater than cost, for some fixed cost \hat{c}_i .

Auction AD shares the same appealing properties that we have demonstrated for Auction NLD.

THEOREM 2. Truthful MBR is an ex post Nash equilibrium for sellers in Auction AD as the minimal bid increment goes to zero, whatever the reported valuation \hat{v} of the buyer, for the case of PI.

We can establish that straightforward bidding is an equilibrium by demonstrating the same equivalence between the outcome of Auction AD and the outcome of the one-sided VCG auction when sellers follow MBR strategies and when PI holds. A detailed proof of this equivalence is presented in the online appendix (available at <http://mansci.pubs.informs.org/ecompanion.html>).

Moreover, Auction AD is efficient when the buyer is truthful, and we again inherit the simple bound on the maximal gain to the buyer from some adaptive (nontruthful) strategy.

COROLLARY 2. Auction AD is efficient when the buyer bids truthfully, and the maximal benefit to a buyer for a nontruthful strategy in Auction AD is no greater than the marginal product of the efficient seller, for the case of PI.

4.3. Auction AD: Simple Example

It is helpful to illustrate the auction dynamics on the simple two-attribute example in Table 1. The buyer wants a *fast, red* car, but is also willing to settle for a *slow, red* car. There are three sellers, and Seller 2 is most competitive for this buyer: There is a potential gain from trade of $\$100 + \$100 - \$80 - \$40 = \$80$ when the buyer buys the *fast, red* car from Seller 2.

Table 1 A Simple Two-Attribute MAP

	Color Red	Speed	
		Fast	Slow
Value (\$)	100	100	60
Cost 1 (\$)	120	80	55
Cost 2 (\$)	80	40	40
Cost 3 (\$)	60	70	45

Table 2 illustrates the state of the auction in each round, as well as the MBR bidding strategies. We assume a bid increment of \$20, and simulate the auction for truthful MBR strategies. Level *red* is most preferred for Attribute 1 in each round, because it is the *only* level for this attribute. Level *fast* is most preferred for Attribute 2 in Round 1, and both *fast* and *slow* are most preferred in future rounds. The price on Attribute 1 is held up at \$100 until Seller 1 drops out of the auction, while Sellers 2 and 3 compete down the price on Attribute 2. Seller 1 drops out in Round 7, which starts new price competition on Attribute 1. The price discount is used to prevent deadlock, which would otherwise occur in this example because there are no linear prices that support the efficient allocation.

The auction terminates with the the efficient outcome (*red, fast*), to Seller 2, and with a payment by the buyer of \$120. The payment in the one-sided VCG auction would be $\$120 + \$10 = \$130$, where \$120 is the cost of Seller 2 for (*red, fast*) and \$10 is her marginal product (i.e., the buyer would take \$70 of the \$80 gain from trade). Auction AD implements this outcome for a small enough minimal-bid increment, ϵ . What follows is a detailed round-by-round description of the behavior of the auction on this example.

Round 1. Prices are initialized to \$120 for each level that is greater than the buyer’s reported value. At this price, every seller bids on *red*, Seller 2 bids for *fast* and *slow* (because she has the same cost for each), and Sellers 1 and 3 bid for *slow*, which costs less than *fast* but has the same ask price. Seller 2 becomes the provisional winner, because the buyer’s value is higher for *fast* than *slow*.

Round 2. Before price adjustment, neither attribute is yet in play, so consider rules *not-a* and *not-b*. Rule *not-b* is used for Attribute 1 because all bids are most preferred, and the ask price on *red* drops to \$100. Rule *not-a* is used for Attribute 2 because both losing bidders submitted a bid on *slow*, which is not most preferred, and the price drops to \$100. Seller 2’s winning bid from the previous round is repeated (at ϵ above the new ask price on *red* and on *slow*). Seller 1 submits a last-and-final bid on Attribute 1, which is now priced below her cost. Both Sellers 1 and 3 bid on both levels of Attribute 2, because the price difference

Table 2 Auction AD on the MAP in Table 1

Round	Rules	Attribute 1			Attribute 2			Discount				
		Red	Fast	Slow								
1		price bid 1	120	2*	3	120	2*	1	2	3	0	
2	<i>not-b not-a</i>	price bid 1 + ε	100	2 + ε	3*	120	2	3*	1	2 + ε	3	0
3	— <i>not-b</i>	price bid 1 + ε	100	2*	3	100	2*	3 + ε	1	2	3 + ε	0
4	— —	price bid 1 + ε	100	2*	3	100	2*	3	1	2	3	0
5	— <i>in</i>	price bid 1 + ε	100	2	3*	80	2 + ε	3*	1	2 + ε	3	0
6	— —	price bid 1 + ε	100	2	3*	80	2	3*	1	2	3	0
7	— <i>in</i>	price bid	100	2*	3	60	2*	3 + ε		2	3 + ε	0
8	<i>in</i> —	price bid	80	2 + ε*	3	60	2*	3 + ε		2	3 + ε	0
9	— —	price bid	80	2 + ε*	3	60	2*	3 + ε		2	3 + ε	0
10	<i>discount</i>	price bid	80	2 + ε	3*	60	2	3 + ε*		2	3 + ε	20
11	— —	price bid	80	2*	3	60	2*	3 + ε		2	3 + ε	20
12	<i>in</i> —	price bid	60	2 + ε*	3	60	2*	3 + ε		2	3 + ε	20
13	— —	price bid	60	2 + ε*	3	60	2*	3 + ε		2	3 + ε	20
14	<i>discount</i>	price bid	60	2 + ε*		60	2*			2		40

Notes. Bid increment \$20. MBR seller strategies, truthful buyer. Prices in each round are those at the *start* of the round before bids are received. Key *not-a*, *not-b*, *in*, and *discount* indicates which rule was used to adjust prices from the previous round, for Attribute 1 and 2, respectively. “+ε” indicates a bid that is ε above the current ask price. “*” indicates the provisional allocation.

between *fast* and *slow* is now within ε of their cost difference. Seller 3 becomes the new provisional winner, outbidding Seller 2.

Round 3. Attribute 1 is now in play, but the price is not adjusted because Sellers 1 and 2, although active in Round 2, were not active on Attribute 1 (bidding at ε above the ask price). This violates the condition for rule *in*. Attribute 2 is still priced above value, and not in play. Rule *not-b* is used because all levels are most preferred, and the prices are decreased to \$100 and \$80 on levels *fast* and *slow*, respectively. Seller 3 repeats her bid from the previous round, while Seller 2 matches the new ask price and becomes the new provisional winner.

Round 4. Attribute 2 is now in play, but there are no price changes by rule *in* because Seller 1 was not active on Attribute 1, and Seller 3 was not active on Attribute 2. No longer the provisional winner, Seller 3 improves her bids, but Seller 2 remains the provisional winner (ties are broken in favor of the current winner).

Round 5. All sellers were active on Attribute 2 during Round 4 and rule *in* is used to adjust prices to \$80 and \$60 on levels *fast* and *slow*, respectively. Seller 2 repeats her bids from the previous round, while Seller 3 matches the new ask prices and becomes the new provisional winner.

Round 6. No price changes. Seller 2 improves her bid, to match the ask prices. Seller 3 remains the provisional winner.

Round 7. All sellers were active on Attribute 2 during Round 6, and rule *in* is used to adjust prices to \$60 and \$40 on levels *fast* and *slow*, respectively. Seller 3 repeats her bids from the previous round, while Seller 2 matches the new ask prices and becomes the new provisional winner. Seller 1, unable to compete on either attribute, drops out of the auction.

Round 8. The price on *red* is reduced to \$80, with rule *in*, now that Seller 1 has dropped out of the auction. Seller 2 repeats her bids. Seller 3 matches the price change on Attribute 1, but is unable to compete on Attribute 2 and submits a last-and-final bid.

Round 9. Ask prices are unchanged. Seller 3 is preferred on Attribute 1, but Seller 2 is preferred on Attribute 2 (and remains the provisional winner).

Round 10. Deadlock is broken by increasing the price discount. Seller 2 repeats her bids from the previous round, while Seller 3 improves her bid and matches the requested \$20 discount, becoming the new provisional winner.

Round 11. Ask prices are unchanged. Seller 2 improves her bid, matching the ask prices and becoming the new provisional winner.

Round 12. Both sellers were active on Attribute 1 in Round 11, and so the ask price drops to \$60. Seller 3 improves her bid, matching this price change, but Seller 2 remains the winner.

Round 13. Ask prices are unchanged, bids are unchanged. Seller 2 remains the provisional winner.

Round 14. Deadlock is broken by increasing the price discount. Seller 3 drops out, unable to compete at this new discount. Seller 2 repeats her previous bid, and the auction terminates with a winning bid of \$80 for *red*, \$60 for *fast*, with a price discount of \$20 for an overall price of \$120.

4.4. Theoretical Analysis

In this section, we establish that Auction AD terminates with maximal CE prices and that these prices correspond to the outcome of the one-sided VCG auction. The structure of the proof is the same as for Auction NLD.

The price space in Auction AD consists of prices on levels, together with an additive price discount. We will provide an integral formulation of MAP-PI, from which we define dual prices that correspond to prices in Auction AD.

$$\max_{x_{ijk}, x_{jk}^B, x_i, y_i} \sum_{j \leq m} \sum_{k \leq |\Theta_j|} v_{jk} x_{jk}^B - \sum_{i \leq n} \sum_{j \leq m} \sum_{k \leq |\Theta_j|} c_{ijk} x_{ijk} \quad [\text{MAP}]$$

$$\text{s.t.} \quad \sum_{k \leq |\Theta_j|} x_{ijk} \leq x_i, \quad \forall i, \forall j \quad (6)$$

$$x_i \leq y_i, \quad \forall i \quad (7)$$

$$\sum_{i \leq n} y_i \leq 1 \quad (8)$$

$$\sum_{i \leq n} x_{ijk} \geq x_{jk}^B, \quad \forall j, \forall k \quad (9)$$

$$x_i \leq 1, \quad \forall i \quad (10)$$

$$\sum_{k \leq |\Theta_j|} x_{jk}^B \leq \sum_{i \leq n} y_i, \quad \forall j \quad (11)$$

$$x_{ijk}, x_{jk}^B, x_i, y_i \geq 0.$$

In addition to variables x_{ijk} and y_i from formulation MAP-PI in §4.1, we introduce x_{jk}^B to define the level selected by the buyer and x_i as an additional variable

to indicate which seller wins. The objective function is stated in terms of x_{jk}^B for the levels selected by the buyer, and x_{ijk} to denote the levels selected by the sellers. Then, Constraints (4) and (5) are restated as Constraints (6) and (7), with Constraints (8) and (9) to model that the buyer can purchase only levels offered by a seller. We explicitly allow $\sum_k x_{ijk} = 0$ for the efficient seller, for which $y_i = 1$, because a seller need not select a level for every attribute. Valid inequalities (10) and (11) are introduced to isolate additional dual variables with a useful economic interpretation.

LEMMA 4. Linear program MAP is integral.

To construct the dual, introduce variables π_{ij} , Δ_i , π^B , p_{jk} , π_i , and π_j^B to correspond with Constraints (6), (7), (8), (9), (10), and (11).

$$\min_{\pi_{ij}, \Delta_i, \pi^B, p_{jk}, \pi_i, \pi_j^B} \pi^B + \sum_{i \leq n} \pi_i \quad [\text{DMAP}]$$

$$\text{s.t.} \quad \pi^B \geq \sum_{j \leq m} \pi_j^B + \Delta_i, \quad \forall i \quad (12)$$

$$\pi_j^B \geq v_{jk} - p_{jk}, \quad \forall j, \forall k \quad (13)$$

$$\pi_i \geq \sum_{j \leq m} \pi_{ij} - \Delta_i, \quad \forall i \quad (14)$$

$$\pi_{ij} \geq p_{jk} - c_{ijk} \quad (15)$$

$$\pi_{ij}, \Delta_i, \pi^B, p_{jk}, \pi_i, \pi_j^B \geq 0.$$

Variables p_{jk} can be interpreted as the price on level k of attribute j , and variable Δ_i can be interpreted as the price discount to seller i . We show that an optimal dual solution exists in which Δ_i is the same for all agents, and write $\Delta_i = \Delta$ for all i . (Variable Δ corresponds with the price discount in Auction AD.)

DEFINITION 4. Prices (p_{jk}, Δ) and feasible MAP solution (θ', i') are in CE if bundle θ' simultaneously maximizes the payoff to the buyer and seller i' at the prices, and every bundle is priced less than cost for all other sellers.

To establish that formulations MAP and DMAP characterize a CE, we show that the complementary slackness (CS) conditions between feasible dual and feasible primal solutions correspond to conditions for CE.

Given prices p_{jk} and discount Δ , the solution to DMAP provides the following dual values:

$$\pi_i = \max \left[\left(\sum_{j \leq m} \pi_{ij} \right) - \Delta, 0 \right], \quad (16)$$

$$\pi_{ij} = \max [p_{jk} - c_{ijk}, 0], \quad (17)$$

$$\pi^B = \sum_{j \leq m} \pi_j^B + \Delta, \quad (18)$$

$$\pi_j^B = \max_{k \leq |\Theta_j|} \{v_{jk} - p_{jk}, 0\}. \quad (19)$$

Each dual variable now has a very natural economic interpretation: π_i is the maximal payoff to seller i across all bundles at the prices; π_{ij} is the maximal payoff to the seller across all levels of attribute j (with the possibility of an opt out); π^B is the maximal payoff to the buyer; and π_j^B is the maximal payoff to the buyer for attribute j with the possibility of an opt out. The interesting CS conditions that relate to seller i are

$$\pi_{ij} > 0 \Rightarrow \sum_{k \leq |\Theta_j|} x_{ijk} = x_i, \quad \forall i, \forall j \quad (20)$$

$$\pi_i > 0 \Rightarrow x_i = 1, \quad \forall i \quad (21)$$

$$x_i > 0 \Rightarrow \pi_i = \sum_{j \leq m} \pi_{ij} - \Delta, \quad \forall i \quad (22)$$

$$x_{ijk} > 0 \Rightarrow \pi_{ij} = p_{jk} - c_{ijk}, \quad \forall i, \forall j, \forall k. \quad (23)$$

Note that $\pi_i > 0$ implies that the total discount-adjusted profit of a seller at the current prices is nonnegative, while $\pi_{ij} > 0$ implies that the total profit on attribute j is nonnegative considering only the level prices on that attribute.

Given the interpretation of dual variables in (19), these state that every seller with positive payoff for some bundle at the current prices must be a winner by CS Condition (21), and that the bundle selected in the primal solution must be exactly the bundle that maximizes the payoff of the winning seller by CS Conditions (20), (22), and (23). The interesting CS conditions that relate to the buyer, are

$$\pi_j^B > 0 \Rightarrow \sum_{k \leq |\Theta_j|} x_{jk}^B = \sum_{i \leq n} y_i, \quad \forall j \quad (24)$$

$$\pi^B > 0 \Rightarrow \sum_{i \leq n} y_i = 1 \quad (25)$$

$$y_i > 0 \Rightarrow \pi^B = \sum_{j \leq m} \pi_j^B + \Delta, \quad \forall i \quad (26)$$

$$x_{jk}^B > 0 \Rightarrow \pi_j^B = v_{jk} - p_{jk}, \quad \forall j, \forall k. \quad (27)$$

Given the interpretation of dual variables in (19), these conditions state that the bundle selected in the primal solution must be exactly the bundle that maximizes the payoff of the buyer at the current prices. We call these “linear + discount” CE prices.

PROPOSITION 5. *Linear + discount CE prices always exist in the MAP with PI, and these prices support the efficient allocation.*

To characterize maximal CE prices, we consider a restricted-dual formulation. Let W denote the attributes for which a nonnull level is selected in the efficient outcome, with i^* to index the winning seller, and k_j^* to index the efficient level on attribute $j \in W$. Then, the restricted dual is formulated to compute

prices p_{jk} , Δ that maximize $\sum_{j \in W} p_{jk_j^*} - \Delta$, while maintaining CS conditions. The most important CS conditions are provided with

$$c_{i^*jk_j^*} \leq p_{jk_j^*} \leq v_{jk_j^*}, \quad \forall j \in W \quad (28)$$

$$\max_{i \neq i^*} \sum_j \pi_{ij} \leq \Delta \leq \sum_j \pi_{i^*j}. \quad (29)$$

Constraints (28) provide CS conditions on each efficient attribute, while Constraints (29) ensure that only the winning seller has positive utility. Attributes $j \notin W$ and levels $k \neq k_j^*$ are priced to provide CE conditions and maximize the payment to the winning seller. Set $\Delta^* = \max_{i \neq i^*} \sum_j \pi_{ij}$, and set the other prices $p_{jk} = \max\{0, v_{jk} - v_{jk_j^*} + p_{jk_j^*}\}$, for $j \in W$, $k \neq k_j^*$, and any $v_{jk} \leq p_{jk} \leq \min_{i \neq i^*} c_{ijk}$ for $j \notin W$. Given this assignment, the problem reduces to solving for $p_{jk_j^*}$ to maximize $\sum_{j \in W} p_{jk_j^*} - (\max_{i \neq i^*} \sum_j \pi_{ij})$, with the prices, p_{jk} , on other attribute levels defined appropriately.

PROPOSITION 6. *The payments in the one-sided VCG auction are implemented in the maximal linear + discount CE prices in the case of PI.*

Let $\Pi_{ij} = \max_k [v_{jk} - c_{ijk}, 0]$, i.e., the maximal allocative surplus on attribute j from seller i . Let \tilde{i} denote the second-best seller. Consider CE prices, $p_{jk_j^*} = v_{jk_j^*} - z_j$, for some $z_j > 0$ on attribute $j \in W$. To characterize maximal CE prices, we consider whether an increase in z_j can also cause a corresponding decrease in $\max_{i \neq i^*} \sum_j \pi_{ij}$, so that the total payment remains constant.

LEMMA 5. *The space of maximal (linear + discount) CE prices are characterized by prices $p_{jk_j^*} = v_{jk_j^*} - z_j$ for $j \in W$ and $\Delta = \sum_j \pi_{ij}$, with $z_j \leq \Pi_{\tilde{i}j}$ and $\sum_j \pi_{ij} \geq \max_{i \neq i^*} \sum_j \pi_{ij}$.*

Note, for attributes $j \notin W$, we can set any price $v_{jk} \leq p_{jk} \leq \min_{i \neq i^*} c_{ijk}$.

4.4.1. Primal-Dual Analysis. We demonstrate convergence of Auction AD to maximal CE prices for truthful MBR strategies. First, we show that the auction implements a primal-dual algorithm for MAP, terminating with a feasible primal and dual satisfying CS Conditions (20) to (27). We then demonstrate that the conditions in Lemma 5 are satisfied, and thus maximal CE prices.

Call an attribute “seller-efficient” for seller i in the MAP when there is some level for that attribute for which the buyer’s value is greater than the seller’s cost. Specifically, we refer to the level that is surplus maximizing for the seller and the buyer as the *seller-efficient* level for this attribute. To simplify the presentation, we assume that there are at least two sellers that are seller-efficient on each attribute. All properties carry over to the more general case.

To construct a feasible primal solution in each round, let x_i and y_i correspond to the provisional winner, with x_{ijk} and x_{jk}^B defined to correspond with the levels that are selected in the provisional outcome. Given ask prices (p_{jk}^t, Δ^t) , construct dual prices $p_{jk} = p_{jk}^t$, $\Delta_i = \Delta^t$ for all i , with variables π^B , π_{ij} , π_i , and π_j^B implicitly defined to satisfy Equation (19).

LEMMA 6. Auction AD maintains CS Conditions (20), (22), (23), (24), (25), and (26) between the provisional allocation and ask prices in each round, when sellers follow MBR strategies, with PI, and as the minimal bid increment goes to zero.

LEMMA 7. Auction AD terminates with a final allocation and prices that satisfy CS Condition (21) for MBR seller strategies, with PI, and as the minimal bid increment goes to zero.

CS Condition (27) requires that the provisional allocation maximizes the buyer's payoff with respect to the prices in the auction. In other words, the provisional allocation must be the best allocation across all possible allocations at the final prices (not just across the allocations that are consistent with the bid from the winning seller). To establish this, we need that the winner will bid on a most-preferred level for all attributes that are priced below value on termination.

LEMMA 8. A seller bids on a monotonically increasing set of most-preferred levels on every attribute in each round while active in Auction AD, and when PI holds.

LEMMA 9. All sellers for which an attribute is seller-efficient bid on a most-preferred level in Auction AD once the attribute is in play and when PI holds.

LEMMA 10. Auction AD maintains CS Condition (27) between the provisional allocation and ask prices on all attributes that are in play when sellers follow MBR bidding strategies, with PI, and as the minimal bid increment goes to zero.

LEMMA 11. Sellers in Auction AD drop out of the auction in order of increasing allocative surplus in the MAP restricted to that seller alone when PI holds.

LEMMA 12. Auction AD terminates with maximal CE prices when sellers follow MBR bidding strategies, and when PI holds.

4.5. Seller Dominance

In this section, we consider a separable special case of MAP with PI. We show that the MAP separates across attributes when the efficient seller dominates the second-best seller on all attributes, and the second-best seller in turn dominates all other sellers. In this case, the MAP can be solved with a simple iterative

auction, and linear CE prices exist without the aggregate discount term. Thus, this characterization makes explicit the role of the aggregate discount in Auction AD. The discount exists to support the efficient allocation and VCG payment when the efficient seller is not the best seller across all attributes, and provides a compact alternative to nonanonymous prices.

DEFINITION 5. Seller i is said to dominate seller i' , written $i \gg i'$, if

$$\begin{aligned} & \max\left[0, \max_{k \in \Theta_j} v_{jk} - c_{ijk}\right] \\ & \geq \max\left[0, \max_{k \in \Theta_j} v_{jk} - c_{i'jk}\right], \quad \forall j. \end{aligned} \quad (30)$$

That is, seller i dominates seller i' if the maximal allocative surplus from seller i is better than from seller i' on all attributes. We define *full dominance* to refer to an auction in which the first-best seller dominates the second-best seller, who in turn dominates all other sellers.

PROPOSITION 7. Linear and maximal CE prices (with a zero discount term) exist in the MAP problem if, and only if, there is both PI and full dominance.

In this case, the allocation problem is separable across attributes, and a simple auction with a separate price trajectory for each attribute is efficient and terminates with the one-sided VCG outcome. Sellers submit independent bids for each attribute, and the winner is determined separately for each attribute, with prices p_{jk}^t adjusted on that attribute to ϵ below the bid price of losing bidders. No explicit coordination across the attributes is required, because at the end of the auction the efficient seller wins for every attribute and the second-best seller sets the winning price for every attribute.

4.6. Number of Rounds to Terminate in Auction AD

The price space in Auction AD is much smaller than in Auction NLD, and the auction converges in a smaller (worst-case) number of rounds.

Let m denote the number of attributes, $l = \max_j |\Theta_j|$ the number of attribute levels, and $W_{\max} = \max_i [\max_j \max_k p_{jk}^1 - c_{ijk}]$, where p_{jk}^1 is the initial price on level k of attribute j . Auction AD converges in rounds polynomial in l , m , W_{\max} , and $1/\epsilon$.

PROPOSITION 8. Auction AD converges in $O(lmW_{\max}/\epsilon)$ rounds, with minimal bid increment ϵ , MBR strategies, and PI.

Counter to this worst-case analysis, it is interesting (and somewhat surprising) that the computational analysis in the next section suggests that the average number of rounds is actually *larger* in Auction AD than in Auction NLD, even though the auction has better information-revelation properties.

5. A Computational Analysis of Information Revelation

In this section, we provide computational results to demonstrate the useful preference-elicitation properties of Auction AD with linear + discount prices in the special case of MAP with PI. We compare the information revelation that is required to compute the efficient outcome in Auction AD with the information revelation that is required to compute the efficient outcome in Auction NLD. These computational experiments are provided to illustrate how the use of an iterative scheme mitigates the informational complexity associated with eliciting the complete cost and value functions from the sellers and buyers as compared to standard one-sided VCG auction.

The results are presented for a simple simulation model in which we generate distributions over values and costs that satisfy PI. We introduce a simple metric to measure the information revelation in each auction in terms of the residual uncertainty about the buyer values and seller costs at the end of the auction. The metric measures the space of possible values and costs that are ex post consistent with the MBR strategies followed by participants when the auction terminates.

5.1. Measuring Information Revelation

To measure residual uncertainty about agent preferences when an auction terminates, we assume an additive form for the cost curves (value curves) for each attribute type. Introducing seller weights $w_{ij} \geq 0$ ($w_j^B \geq 0$) on attribute j , we can write the cost (value) function of a seller (buyer) as

$$c_i(\theta) = \sum_{j \leq m} w_{ij} c_{ij}(\theta_j) \quad (31)$$

$$v(\theta) = \sum_{j \leq m} w_j^B v_j(\theta_j). \quad (32)$$

Let $w_i = (w_{i1}, \dots, w_{im})$ denote the weight vector for seller i , and $w^B = (w_1^B, \dots, w_m^B)$ denote the weight vector for the buyer. We normalize the weights so that $\sum_j w_{ij} = 1$ and $0 \leq w_{ij} \leq 1$ for all attributes (similarly for the buyer). With this, we can encapsulate what is not known about a seller's preferences or a buyer's preferences in a space of feasible *weights* for each attribute type. Notice that the uncertainty is represented by the unit simplex irrespective of the form of the functions $v_j(\theta_j)$.

Every time a participant responds to prices we can add new constraints to the weight space. Let $p^t(\theta)$ denote the ask prices on attribute bundles θ in round t of the auction. In Auction AD these prices are defined in terms of the linear + discount price structure. Suppose that seller i bids on bundles θ^*

at these prices. Given a truthful MBR strategy, this implies the following constraints on her weights:

$$\begin{aligned} p^t(\theta^*) - \sum_{j \leq m} w_{ij} c_{ij}(\theta_j^*) + \epsilon \\ \geq p^t(\theta') - \sum_{j \leq m} w_{ij} c_{ij}(\theta_j'), \quad \forall \theta'. \end{aligned} \quad (33)$$

Note that MBR implies that the payoff to the seller for her bid is maximal across all bundles given the current prices, and that Constraints (33) are linear in the space of weights. Additional revealed-preference information in each iteration reduces the volume of the polytope that is used to represent the uncertainty in the weights. The *residual volume* is a measure of the information that has *not* been revealed by a participant; i.e., large residual volumes indicate that there is still considerable uncertainty about preferences. We define the *normalized residual volume*, $\text{Vol}(C)$, given a set of constraints C on weights $w \in [0, 1]^m$ as

$$\text{Vol}(C) = \left(\int_{w \in [0, 1]^m} f(w) dw \right)^{1/m}, \quad (34)$$

where $f(w) = 1$ when weights w satisfy constraints C , and $f(w) = 0$ otherwise. We take the m th root to normalize for the number of attributes (m) and provide a measure of the average residual per-attribute uncertainty. A normalized volume of one represents complete uncertainty, while a normalized volume of zero represents complete certainty and exact information.

We adopt a similar method to measure the information revelation from the buyer, via the revealed-preference information in the solution to the winner-determination problem and in the price updates in each round. This information revelation on the buy side provides a measure of the preference-elicitation cost that a buyer would face if we introduced dynamic bidding for buyers as well as sellers (as discussed in §3.4).

Algorithmically, we maintain a list of constraints on weight space for each seller and the buyer, introducing new constraints in each round. The normalized residual volume given current constraints is estimated using a simple *Monte Carlo* algorithm, in which we generate nS uniform random weight vectors and test whether the sample is within the feasible weight space region as defined by the constraints. Let $x(nS)$ denote the number of samples that are within the region. We approximate the normalized residual volume as $(x(nS)/nS)^{1/m}$.

5.2. Experimental Details

We consider a distribution on PI preferences that is parameterized by the number of bidders n , the number of levels l on each attribute, the number of attributes m , and two constants, $\alpha^S > 0$ and $\alpha^B > 0$. For

each seller, i , we randomly select weight $w_{ij} \sim U(0, 1)$ and then normalize so that $\sum_j w_{ij} = 1$. Then, to generate a marginal cost function c_{ij} for each attribute, we randomly select l values from $U(0, \alpha^S \cdot l)$ and sort these values to get the value of $c_{ij}(\theta_j)$ evaluated at each level θ_j of attribute j . We define a weight vector for the buyer in the same way, and then generate a marginal valuation function v_j^B for each attribute by selecting l values $U(0, \alpha^B \cdot l)$. We choose $\alpha^B \geq \alpha^S$ to model the idea that the value of the buyer is greater than the cost of the typical seller.

In our computational experiments, we assume that both the buyer and the sellers follow straightforward (truthful) bidding strategies. The buyer reports her true valuation function to the auction, and the sellers follow truthful MBR strategies. By default, we set the number of sellers, the number of attributes, and the number of attribute levels to 4; we set $\alpha^S = 30$ and $\alpha^B = 40$; and we adjust the minimal bid increment to achieve an allocative efficiency of at least 98%. We also remove from the simulation any instances for which there is no competition and any instances in which there is some attribute not supplied in the efficient outcome. All experimental results are averaged over 10 trials, and we performed 800 Monte Carlo samples in each round to track the information revelation in each trial. We checked that our results are robust to performing larger numbers of Monte Carlo samples.

5.3. Results

Figure 3 compares the information-revelation properties of Auction AD and Auction NLD on a problem with four sellers, four attributes, and four levels per attribute. We plot the efficiency at termination,

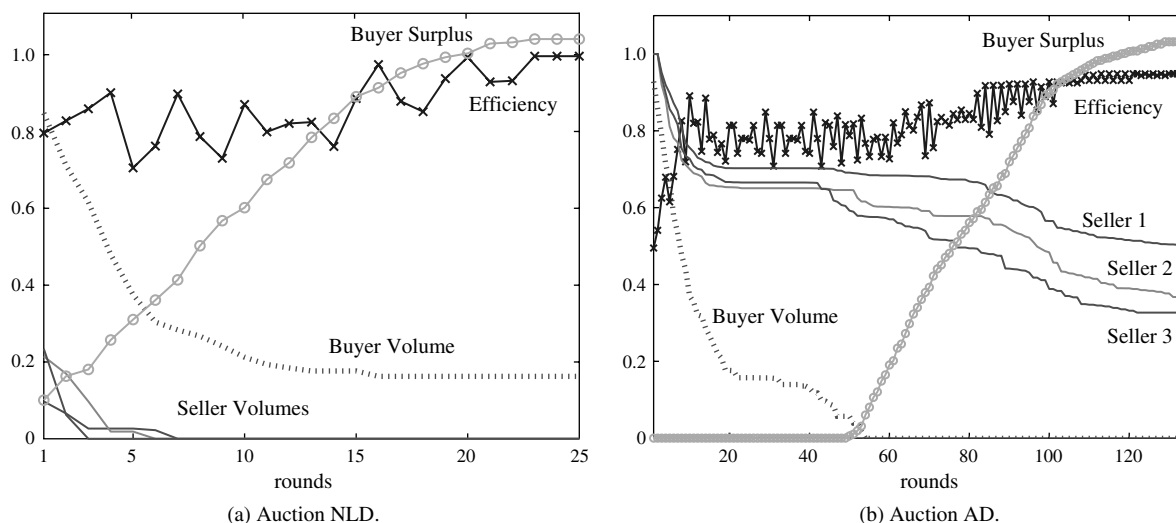
and the ratio of the buyer’s payoff at the end of the auction to the buyer’s payoff in the one-sided VCG auction. The slight overshoot from the payoff in the VCG auction is due to the error introduced because of the minimal bid increment. We also plot the normalized residual volume for three different sellers—the *efficient* seller, the *second-best* seller, and a *random* seller—along with the normalized residual volume for the buyer. Recall that the normalized residual volume provides a measure of the information that a participant has *not* revealed about her preferences.

The most interesting effect of moving from Auction NLD to Auction AD is on information revelation by the sellers. Notice that the sellers reveal complete information about costs in Auction NLD, but are able to retain between 30%–50% of their information in Auction AD. This saving does come at some cost to the buyer, who retains around 20% of her information in Auction NLD, but must reveal all of her information in Auction AD.

Figure 4 compares the information-revelation properties between Auction NLD and Auction AD as the number of sellers are increased. We plot the normalized residual volume, averaged across the buyer and the same three sellers as in the initial set of experiments, and investigate the effect of the number of sellers on the information-revelation requirements. We see that Auction AD dominates Auction NLD for all numbers of sellers. It is also interesting to observe that increasing the number of sellers seems to *reduce* the final information revelation in both auctions, although the transients in Auction AD show the opposite trend.

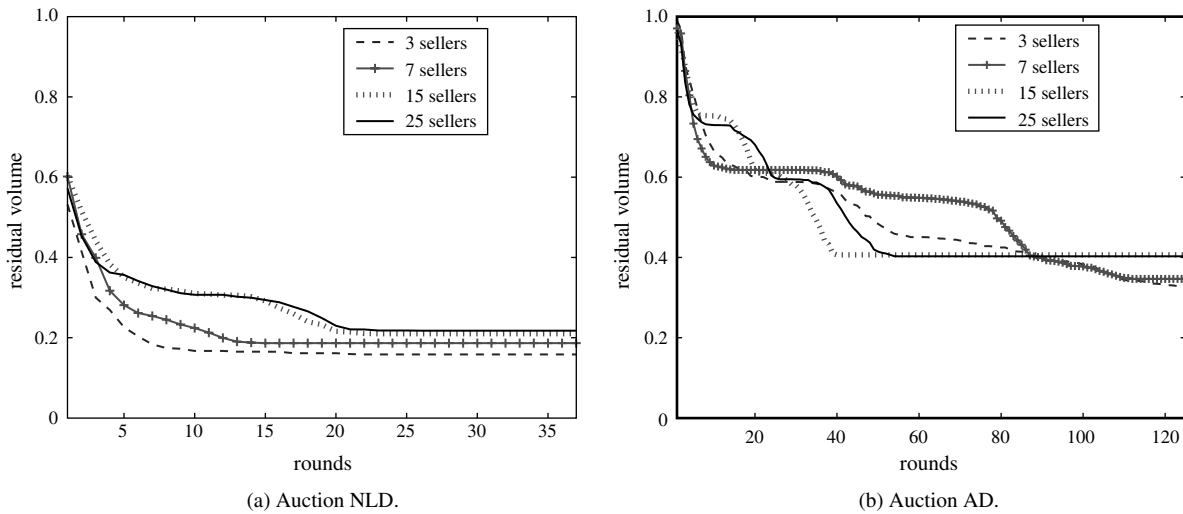
Figure 5 compares the information-revelation properties in Auction NLD and Auction AD as we increase

Figure 3 Multiattribute Auction Problem: Four Sellers, Four Attributes, Four Levels per Attribute



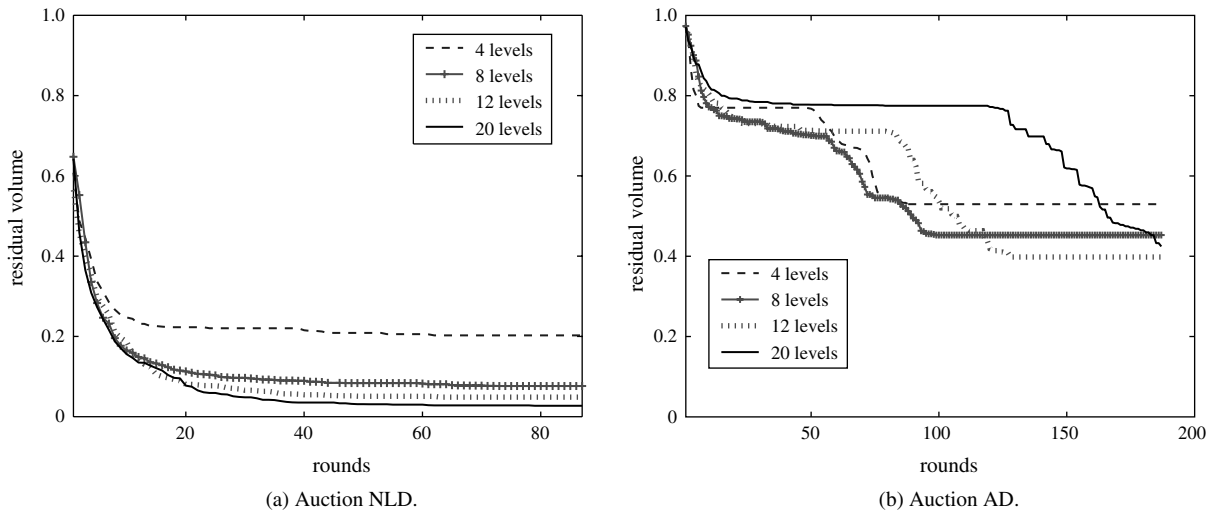
Notes. Buyer surplus is normalized to that of the surplus in the one-sided VCG auction. Normalized residual volume is plotted for the buyer, and for three of the four sellers (the *efficient* seller, the *second-best* seller, and a *random* seller). Results averaged over 10 trials.

Figure 4 Multiattribute Auction Problem: Three Attributes, Three Levels per Attribute



Notes. Normalized residual volume is averaged across the buyer, the *efficient* seller, the *second-best* seller, and another seller selected at random. Results are averaged over 10 trials.

Figure 5 Multiattribute Auction Problem: Two Attributes, Five Sellers



Notes. Normalized residual volume is averaged across the buyer, the *efficient* seller, the *second-best* seller, and another seller selected at random. Results are averaged over 10 trials.

the number of levels on each attribute. As before, we take the average of the normalized residual volume across, averaged across the buyer, and the same three sellers.⁴ Auction AD continues to dominate the performance of Auction NLD for all numbers of attribute levels. Also, we see the same dichotomy in that, while increasing the number of levels seems to *increase* the final information revelation in both auctions, the transients in Auction AD show the opposite trend.

⁴These experiments were performed with buyer valuation functions parameterized with $\alpha^B = 60$.

6. Conclusions

Multiattribute auctions can support the efficient procurement of configurable goods and services through the combined use of expressive bidding languages and competition across suppliers. Efficient markets are central to procurement activity where buyers and suppliers are engaged in long-term relationships. Allocative efficiency, rather than pure profit maximization for the buyer, is important to sustain the relationship in these strategic situations. In addition, due to power asymmetry (typically big buyers and small suppliers) and due to the cost of preference elicitation and concern about revealing value and cost

information, it is important that these protocols solve problems with minimal information revelation.

We proposed two models for iterative multiattribute procurement auctions. The auctions are price based and support incremental bidding from suppliers. Auction NLD solves general problems and employs a large price space, while Auction AD is optimized for the important special case of PI and employs a compact linear price space with an aggregate price discount. The auctions support straightforward bidding by sellers and bound the possible gain from manipulation to buyers. Computational results demonstrate that Auction AD allows for a significant reduction in information revelation over one-shot auctions, with the average seller retaining around 50% of her cost information at termination.

Looking to future work, in considering the preference-elicitation properties of multiattribute auctions, we are interested to continue the study of *accelerated* auctions that was initiated in Sunderam and Parkes (2003). The idea is to allow multiple virtual rounds between proxy agents and the auction, and only fall back and ask for additional bids from suppliers when no progress is possible within the auction. Another recent idea is to use computational learning theory to generate elicitation queries (Lahaie and Parkes 2004). We also find it interesting to explore the role of hybrid auctions (Porter et al. 2003, Ausubel and Milgrom 2004), with linear prices used in the early stage as a method to perform elicitation, coupled with a final one-shot stage. Finally, we identified an interesting tension between allowing adaptive buy-side strategies and providing incentives for straightforward bidding on the sell side. We are interested to understand how effective a proxy agent, able to constrain a buyer to follow a bidding strategy with consistent revealed-preference information, would be in mitigating this effect.

An electronic companion to this paper is available at <http://mansci.pubs.informs.org/ecompanion.html>.

Acknowledgments

The authors thank Aditya Sunderam for his assistance with the computational experiments reported in §5 of this paper. The first author gratefully acknowledges the support of an IBM Faculty support grant, and the useful suggestions made by participants in seminars at Caltech, the University of Connecticut, and INFORMS'02 in San Jose. Finally, the authors thank William Walsh, and the anonymous reviewers and the associate editor for their helpful comments on earlier drafts of this paper. This work is supported in part by NSF grant IIS-0238147.

References

Ausubel, L. M., P. Milgrom. 2004. Ascending proxy auctions. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chapter 3. MIT Press, Cambridge, MA.

Beil, D. R., L. M. Wein. 2003. An inverse-optimization-based auction mechanism to support a multiattribute RFQ process. *Management Sci.* **49** 1529–1545.

Bichler, M., J. Kalagnanam, H. S. Lee. 2003. RECO: Representation and evaluation of configurable offers. K. Hemant, Bhargava, Nong Ye, eds. *Computational Modeling and Problem Solving in the Networked World*. Kluwer Academic Publishers.

Bikhchandani, S., J. M. Ostroy. 2002. Ascending price Vickrey auctions. *Games Econom. Behavior* Forthcoming.

Bikhchandani, S., S. de Vries, J. Schummer, R. V. Vohra. 2001. Linear programming and Vickrey auctions. B. Dietrich, R. Vohra, eds. *Mathematics of the Internet: E-Auction and Markets. IMA Volumes in Mathematics and Its Applications*. Springer-Verlag, New York, 75–116.

Branco, F. 1997. The design of multidimensional auctions. *RAND J. Econom.* **28** 63–81.

Che, Y. K. 1993. Design competition through multidimensional auctions. *RAND J. Econom.* **24** 668–680.

Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11** 17–33.

Cramton, P. 1998. Ascending auctions. *Eur. Econom. Rev.* **42** 745–756.

de Vries, S., J. Schummer, R. V. Vohra. 2003. On ascending Vickrey auctions for heterogeneous objects. Technical report, MEDS, Kellogg School, Northwestern University, Evanston, IL.

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

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

Jennings, N., P. Faratin, A. Lomuscio, S. Parsons, C. Sierra, M. Wooldridge. 2001. Automated negotiation: Prospects, methods and challenges. *Internat. J. Group Decision Negotiation* **10**(2) 199–215.

Keeney, R. L., H. Raiffa. 1993. *Decision Making with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, Cambridge, UK.

Kraus, S. 1997. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence* **94** 79–97.

Krishna, V., M. Perry. 2000. Efficient mechanism design. Technical report, Pennsylvania State University, University Park, PA.

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

Lahaie, S. M., D. C. Parkes. 2004. Applying learning algorithms to preference elicitation. *Proc. ACM Conf. Electronic Commerce*. ACM Press, New York, 180–188.

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

Mishra, D., D. C. Parkes. 2004. Ascending price Vickrey auctions using primal-dual algorithms. Technical report, Harvard University, Boston, MA.

Myerson, R. B., M. A. Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *J. Econom. Theory* **28** 265–281.

Parkes, D. C., L. H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *Proc. 17th National Conf. Artificial Intelligence (AAAI-00)*, MIT Press, Cambridge, MA, 74–81.

Porter, D., S. Rassenti, V. Smith. 2003. Combinatorial auction design. *Proc. National Acad. Sci.* **100**(19) 11153–11157.

Sunderam, A. V., D. C. Parkes. 2003. Preference elicitation in proxied multiattribute auctions. *Fourth ACM Conf. Electronic Commerce (EC'03)* Cambridge, MA. ACM Press, New York.

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

Vulkan, N., N. R. Jennings. 2000. Efficient mechanisms for the supply of services in multi-agent environments. *Decision Support Systems* **28** 5–19.

Copyright 2005, by INFORMS, all rights reserved. Copyright of Management Science is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.