An Inverse-Optimization-Based Auction Mechanism to Support a Multiattribute RFQ Process

Damian R. Beil • Lawrence M. Wein
University of Michigan Business School, Ann Arbor, Michigan 48109-1234
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142
dbeil@umich.edu • lwein@mit.edu

We consider a manufacturer who uses a reverse, or procurement, auction to determine which supplier will be awarded a contract. Each bid consists of a price and a set of nonprice attributes (e.g., quality, lead time). The manufacturer is assumed to know the parametric form of the suppliers' cost functions (in terms of the nonprice attributes), but has no prior information on the parameter values. We construct a multiround open-ascending auction mechanism, where the manufacturer announces a slightly different scoring rule (i.e., a function that ranks the bids in terms of the price and nonprice attributes) in each round. Via inverse optimization, the manufacturer uses the bids from the first several rounds to learn the suppliers' cost functions, and then in the final round chooses a scoring rule that attempts to maximize his own utility. Under the assumption that suppliers submit their myopic best-response bids in the last round, and do not distort their bids in the earlier rounds (i.e., they choose their minimum-cost bid to achieve any given score), our mechanism, indeed, maximizes the manufacturer's utility within the open-ascending format. We also discuss several enhancements that improve the robustness of our mechanism with respect to the model's informational and behavioral assumptions.
(*Inverse Optimization; Multiattribute Auctions; Mechanism Design*)

1. Introduction

Although the market for online business-to-business auctions is enormous (estimated at $746 billion in 2004 by Kafka et al. 2000), the price-only auctions that dominate the current e-commerce landscape severely hinder the range of products that can be auctioned over the Internet. In particular, within the industrial procurement setting, many low-cost standardized items are being transacted by current online procurement (or reverse) auctions, while high-value complex items are still being procured via the traditional request for quotes (RFQ) process. An RFQ process allows the sale to be determined by a variety of attributes, involving not only price, but quality, lead time, contract terms, supplier reputation, and incumbent switching costs. It also lets the manufacturer reveal his preferences and permits the suppliers to compete on their own specialized dimensions. Consequently, e-marketplaces are currently being developed to partially automate the RFQ process, i.e, to create an eRFQ process (see Kafka et al. 2000 for examples).

This paper was stimulated by about eight hours of discussions (during the fall of 2000 and winter of 2001) with the Chief Technology Officer (CTO) of Frictionless Commerce, who was seeking help with designing a multiattribute eRFQ mechanism. The CTO described the company's multiattribute procurement software (we are not at liberty to discuss its details) and the perceived needs and preferences

0025-1909/03/4911/1529
1526-5501 electronic ISSN

of their customers (i.e., the manufacturers who own their software) and the supplier companies (i.e., the potential bidders) with respect to various aspects of both traditional and electronic RFQ processes. This information helped to guide our eRFQ design and our assumptions about supplier behavior. After receiving a rough first draft of this paper, the CTO shared his ideas with several customers, and their impressions are briefly summarized in §4.

The appropriate mathematical context for this setting is the *multiattribute*, or multidimensional, auction; consequently, we often refer to the manufacturer as the auctioneer (or bid taker) and the suppliers as bidders. There are two primary objectives in the auction theory literature: (1) revenue maximization (on the part of the auctioneer), and (2) (allocative) efficiency. In multiattribute auctions, it is appropriate to speak in terms of utility maximization rather than revenue maximization. Efficiency is often the goal in public-sector auctions, whereas utility maximization is typically strived for in private-sector auctions. An efficient auction mechanism maximizes the total surplus, without concerning itself with how this surplus is divided among the bidders and the auctioneer. In complex auctions, these two objectives are often in conflict (e.g., Bikhchandani 1999), because the auctioneer can usually increase his utility by either withholding items for sale or by allocating items to those who do not value them the most. We assume a utility-maximizing auctioneer, but will discuss efficiency in §2.5.

To engage bidders in a multiattribute auction, an auctioneer needs to provide the bidders with some information pertaining to how he values the nonprice attributes. While several rather obtuse approaches are possible (e.g., the auctioneer could provide shadow prices from a mathematical program without revealing the mathematical program), the predominant approach in RFQ practice—and the one favored by most bidders because of its straightforward nature—is for the auctioneer to announce a *scoring rule* in terms of the bid price and various attributes. This scoring rule may or may not be identical to the auctioneer's true utility function, indeed, this is the crux of the strategic problem from the auctioneer's viewpoint.

The standard framework in which to study these incentive issues is mechanism design (Myerson 1981) with multidimensional types (which arises here from the multiple attributes). Although some results have been derived under special assumptions on the preference and distributions of types (see the survey paper by Rochet and Stole 2001), mechanism design with multidimensional types is difficult because of the nature of the incentive-compatibility constraints.

There are two key papers on multiattribute auctions with scoring rules, one addressing each objective. Milgrom (2000a) recently has shown that efficiency is achieved if the auctioneer announces his true utility function as the scoring rule, and conducts a Vickrey (i.e., second-price, sealed-bid) auction based on the resulting scores. The other important paper—and the one most relevant for our purposes—is by Che (1993), who considers a two-dimensional price-quality procurement auction in a sealed-bid setting. He assumes that the quality costs of his two symmetric suppliers are a function of a single parameter that is private information and independent across suppliers (Branco 1997 generalizes this research to the case of correlated supplier costs). Che (1993) assumes a positive cross partial of cost with respect to quality and type, implying that the producer with the lower cost for producing a given quantity is also the producer with the lower marginal cost of producing quality (this is not the case in our Figures 2 and 3 in §2.6). He shows that to maximize utility, the auctioneer—who only knows the probability distribution of the cost parameter—announces a scoring rule that understates the value of quality, so as to limit the informational rents collected by the low-cost suppliers.

Our paper takes advantage of the multiround nature of eRFQ processes by enabling the manufacturer to *learn* the suppliers' cost functions. More specifically, we propose a forward- and inverse-optimization-based approach (*inverse optimization* deduces some of the unobserved parameters of an optimization problem from the observed solution) that allows the manufacturer, via several changes in the announced scoring rule, to learn the suppliers' cost functions and then determine a scoring rule that maximizes his utility within the open-ascending auction format. We choose this format because it has many characteristics that make it superior to sealed-bid auctions from a practical standpoint (see Cramton 1998 for details) and because it is the most common type of procurement auction on the Internet.

While we relax several of Che's (1993) assumptions (two symmetric bidders with a positive cross partial of cost, although by the last round our auctioneer has perfect information), we do so at the expense of incentive compatibility. In particular, we assume that the bidders in the last round of the auction submit their myopic best-response bids; the reasonableness of this assumption is discussed in §2.3.

RFQ processes are typically less structured than auctions. Although changing the scoring rule during the course of a traditional auction would be perceived by many bidders as unfair, this practice is not uncommon in RFQ processes. The scoring rule may be changed throughout the course of an RFQ process for a variety of reasons, e.g., the buyer may learn from supplier presentations that the importance of certain attributes has been misestimated, or the buyer may want a certain supplier to be awarded the contract and needs to alter the scoring rule so as to enable this supplier to attain the highest score. Consequently, several commercial eRFQ software packages allow changes in the scoring rule throughout the course of the process. Not only can the scoring rule change over time, but neither the suppliers' bids nor the manufacturer's scoring rule needs to be binding. Nonetheless, RFQ processes maintain a certain amount of structure, because reputations are diminished by too much noncommittal behavior. We believe that our basic mechanism—by learning the suppliers' cost information and strategically setting the scoring rule—has the potential to increase a manufacturer's utility in an eRFQ setting, even if the RFQ process is less structured than assumed in our model.

The auction mechanism is described in §2. Section 3 discusses several practical considerations, and concluding remarks are offered in §4.

2. The Mechanism

2.1. Notation

We assume that the auctioneer is buying a single item. Although this item may represent six months of production for a subassembly, we assume that it is sold to a single bidder, and we model it as a single item. While many RFQ processes in practice are for multiple items, a significant fraction of RFQs are for a single item. Because our notation requires up to five subscripts, we use mnemonic subscripts, where $a = 1, \dots, A$ indexes the attributes, s indexes the S suppliers, p indexes the P cost parameters per attribute, and $r = 1, \dots, P+1$ indexes the rounds of the auction; the fifth subscript is introduced in §3.2. To limit the notational complexity, we frequently suppress subscripts that are not crucial to the immediate discussion, e.g., certain variables sometimes appear with two subscripts, and other times with three subscripts. We hope that these inconsistencies are more than offset by improved readability.

For multiattribute auctions, it is important to distinguish between attributes that are endogenous (i.e., bidder controllable), such as lead time and quality, versus attributes that are exogenous, such as a bidder's reputation at the time of the auction. For expositional purposes, we assume that all nonprice attributes are endogenous, and defer a discussion of exogenous attributes to §3.5. Each bid is of the form $(p, x_1, \dots, x_A)$, where p denotes price and x_a is the magnitude of nonprice attribute a for $a = 1, \dots, A$. To simplify the presentation and analysis, we assume that the nonprice attributes are continuous, nonnegative variables (thus the domains of the cost, utility and scoring functions are the nonnegative real numbers) and that larger values of x_a are more desirable from the auctioneer's point of view and more costly from the suppliers' point of view. Hence, attributes such as tolerance or lead time, which are desirable and costly when low in magnitude, need to be defined relative to a worst-case upper bound. Supplier s's cost function is additive across attributes and is given by $\sum_{a=1}^{A} c_{as}(x_a; \vec{\theta}_{as})$, where c_{as} is increasing, convex (convexity and concavity are strict in this paper), and twice continuously differentiable in x_a. For ease of presentation, we assume that c_{as} is a function of exactly $P > 1$ cost parameters for all a and s (we describe at the end of §2.4 how this assumption, $\vec{\theta}_{as} \in \mathbb{R}^P$, is easily relaxed), and in the sequel suppress the term $\vec{\theta}_{as}$ when writing c_{as}. In practice, we expect that P would typically equal two or three, and our example in §2.6 considers a three-parameter cost function c_{as} of the form $\theta_{as1}x_a + \theta_{as2}x_a^3 + \theta_{as3}x_a^{15}$. We make the crucial assumption that the auctioneer knows the form of the suppliers' cost functions, but does not have any information about the parameter values $\vec{\theta}_{as}$. Although an auctioneer is unlikely to know the form of the suppliers' cost functions in practice, in §3.2, we propose a procedure that chooses among alternative functional forms. Also, while the auctioneer may have some prior information about the parameters of the suppliers' cost functions (e.g., upper or lower bounds), we do not consider a Bayesian framework and, as explained in §2.4, we use the suppliers' bids in the first P rounds of the auction to estimate the parameter values.

The auctioneer's true utility function and scoring rules are assumed to be additive across nonprice attributes. While the separability across attributes of the cost and utility functions and scoring rules makes the problem more tractable and may not hold in practice, most existing multiattribute software packages also make these assumptions (one exception is CombineNets ClearBox 2.0, www.combinenet.com), and the CTO we dealt with felt that nonseparable functions would be too arduous for industrial implementation. The true utility function is given by $\sum_{a=1}^{A} v_a(x_a) - p$, where v_a (mnemonic for value) is increasing, concave, and twice continuously differentiable in x_a. Each round of the auction is characterized by a different scoring rule, and the number of rounds is exactly one more than the number of cost parameters per attribute. The scoring rule in round $r = 1, \ldots, P+1$ is denoted by $\sum_{a=1}^{A} \hat{v}_{ar}(x_a) - p$. However, for fixed a and r, the scoring rule $\hat{v}_{ar}$ must be increasing and concave in x_a. We also assume that

$$\frac{\partial c_{as}(x_a)}{\partial x_a} < \frac{\partial \hat{v}_{ar}(x_a)}{\partial x_a} \qquad \text{as } x_a \to 0^+ \quad \forall s,$$

and

$$\frac{\partial \hat{v}_{ar}(x_a)}{\partial x_a} \to 0 \qquad \text{as } x_a \to \infty,$$

to guarantee that the solution to the optimization problem in (1)–(2) possesses a finite positive solution. One final, more technical assumption on only the scoring rules in rounds $r = 1, \ldots, P$, is delayed for expository purposes until §2.4.

2.2. Basic Outline

The auctioneer initially informs the bidders that the auction will consist of $P+1$ rounds; our use of multiple rounds is not unlike traditional RFQ processes, which typically consist of multiple stages. At the beginning of each round $r = 1, \ldots, P+1$, the auctioneer announces the scoring rule $\sum_{a=1}^{A} \hat{v}_{ar}(x_a) - p$ to all suppliers. Suppliers then submit bids of the form $(p, x_1, \ldots, x_A)$ in an open-ascending manner. We envision this mechanism taking place electronically (e.g., over the Internet). Within each round, suppliers have ample opportunity to bid, and each supplier is forced to make a new bid during each round to proceed to the next round (at this point in the paper, we do not rule out the possibility that a supplier simply resubmits an earlier bid); other activity rules and transition rules are briefly discussed in §3.4. The auctioneer ranks the bids according to the current scoring rule and displays the ranked scores, but does not reveal the bidders' identities or detailed bids. In contrast to a traditional open-ascending auction, submitted bids need not exceed the current best bid. However, there is a minimum bid increment (with respect to the scoring rule) to take the lead (thereby speeding up the auction, at least in the final round), and we assume that the highest bidder at the end of round $P+1$ wins the contract, at his proposed bid.

The analysis that provides the basis for our mechanism consists of three main parts, which are described in the next three subsections: (1) how the suppliers bid given the current scoring rule and current best score, (2) how the auctioneer estimates the suppliers' cost functions given their bids, and (3) how the auctioneer determines an optimal scoring rule once he learns the suppliers' cost functions.

2.3. Supplier Behavior

To specify supplier behavior in our model, we need to describe the notion of *myopic best-response* (MBR) bids and introduce the weaker condition of *undistorted bids*. A supplier using MBR chooses his next bid to maximize his current profit, assuming no other suppliers change their bids, i.e., he behaves as if the auction were ending after his bid. More specifically, if in round r the current top score is S (we use S to denote the score and the number of suppliers, but this should cause no confusion) and the minimum bid increment is ϵ, then, supplier s solves the following optimization problem (note that the subscripts s and r are suppressed in the decision variables x_{asr} and p_{sr}):

$$\max_{p,x_1,\dots,x_A} \quad p-\sum_{a=1}^{A} c_{as}(x_a) \tag{1}$$

$$\text{subject to} \quad \sum_{a=1}^{A} \hat{v}_{ar}(x_a)-p=S+\epsilon. \tag{2}$$

If the optimal objective function value in (1) is nonnegative, then the corresponding optimal solution is the MBR bid. If the optimal objective function value is negative, then the MBR is to not submit a new bid.

The MBR assumption has been used in a variety of recent auction studies (e.g., Demange et al. 1986, Wellman et al. 2001, Parkes and Ungar 2000, Gallien and Wein 2000, although the first two studies refer to MBR as "straightforward bidding"), and asserts a middle ground in regard to the bidders' rationality. On the one hand, they are assumed to be sophisticated enough to formulate and solve (1)–(2). On the other hand, a more astute bidder would formulate prior distributions on the other bidders' cost functions and the auctioneer's utility function, and would account for the fact that these other players would be solving their own game-theoretic problems. While we view the MBR as a reasonable and tractable compromise to a difficult modeling question, it is important to point out that there may be bidders who do not even use an optimization-based mental model for their bidding (this was the opinion of the CTO we spoke with), and others that, upon repeated exposures to this mechanism, may realize that the first few rounds of bidding are for the purposes of learning on the part of the auctioneer, and may place bids that withhold, or even intentionally distort, their cost information. In §3, we discuss ways that our mechanism can be enhanced to mitigate this danger.

Before specifying supplier behavior, we introduce a weaker condition that we call undistorted bids. If a supplier submits a round r bid $(p, x_1, \dots, x_A)$ that generates the arbitrary score $\widetilde{S}$, we say that it is undistorted if this bid is the solution to (1)–(2) with the right side of (2) replaced by $\widetilde{S}$. A supplier submitting undistorted bids may withhold *absolute* information about his cost parameters by not aggressively bidding (i.e., even though the MBR bid may allow him to take the lead, he nonetheless may choose to withhold this bid for strategic reasons), but does not withhold *relative* information about his cost parameters. We believe that bid distortion (i.e., for a given score, purposely choosing a suboptimal bid to deceive the auctioneer about the relative contributions to the supplier's cost function) is considerably more sophisticated and risky than the strategy of withholding absolute information about the cost function, and is much less likely to be engaged in than the latter strategy, particularly if activity rules are in place (see §3.4). Moreover, because the suppliers know the number of rounds in the auction, they are unlikely to withhold absolute cost information in the final round of bidding, because doing so could cause them to lose the auction.

These arguments motivate our main assumption about supplier behavior: *suppliers' bids are undistorted in rounds* $1, \dots, P$, *and are MBR in round* $P+1$.

Our earlier assumptions about the cost function and scoring rules imply that the solution to (1)–(2) can be found by solving (2) for p, substituting for p into (1) to get an unconstrained optimization problem, and solving the first-order conditions

$$\frac{\partial \hat{v}_{ar}(x_a)}{\partial x_a}=\frac{\partial c_{as}(x_a)}{\partial x_a} \qquad \text{for } a=1,\dots,A. \tag{3}$$

Because the first-order condition (3) is independent of the right side of (2), which is due to the quasi linearity of (1), it follows that all undistorted bids satisfy (3). Hence, all bids by a given supplier s in a given round r have the same magnitude for their nonprice attributes; these undistorted bids differ only in their price. This observation provides a simple way to empirically validate or invalidate the undistorted-bid assumption.

If we let x_a^* denote the solution to (3), then the corresponding bid price is

$$p^*=\sum_{a=1}^{A} \hat{v}_{ar}(x_a^*)-S-\epsilon. \tag{4}$$

If $p^* \geq \sum_{a=1}^{A} c_{as}(x_a^*)$, then $(p^*, x_1^*, \dots, x_A^*)$ is the MBR bid, otherwise, the MBR is to not submit a new bid.

2.4. Cost Estimation

Because bids are undistorted and each supplier is forced to submit at least one new bid in each round, at the end of round P the auctioneer possesses, for each attribute a and each supplier s, the P equations given by (3) with $r = 1, \ldots, P$, in terms of the P unknown cost parameters. If for fixed attribute a and supplier s the scoring rules $\hat{v}_{ar}$ induce a different x_a^* for each round $r = 1, \ldots, P$, these P equations can be solved to obtain supplier s's true cost parameters for attribute a. This requirement is satisfied as long as, for fixed a and s,

$$\frac{\partial \hat{v}_{ar}(x_{asr}^*)}{\partial x_{asr}} \neq \frac{\partial \hat{v}_{a\bar{r}}(x_{as\bar{r}}^*)}{\partial x_{as\bar{r}}}, \qquad \text{for } \bar{r} = 1, \ldots, r-1.$$

In practice, this could be enforced by simply perturbing (perhaps several times if needed) scoring rules that would otherwise fail the condition. Per the above, at the end of P rounds, the auctioneer knows the suppliers' cost functions. In §3.3, we discuss how the first P scoring rules might be determined in practice.

Note that if the number of parameters per attribute, P, varied by attribute and supplier, then we would learn the parameter values for attribute a and supplier s at the end of round P_{as} and, hence, the number of rounds required to learn all suppliers' cost functions is $\max_{a,s} P_{as}$.

2.5. The Optimal Scoring Function in Round $P+1$

Before turning to the optimal scoring rule in round $P+1$, we briefly discuss an alternative approach for the auctioneer, now that he knows the suppliers' cost functions: In lieu of an auction in round $P+1$, simply cut a deal with the "best supplier" (defined as supplier 1 four paragraphs below) and extract all the rents. We discussed this approach with the CTO of Frictionless Commerce, who felt (as a result of numerous conversations with customers, i.e., auctioneers) that it was not a viable option (in the private sector or the military) because many of these suppliers engage in several RFQs per year with a given auctioneer. To continue to attract suppliers back to these RFQ processes, it is vitally important that the RFQ process be viewed by the suppliers as fair (e.g., Frictionless Commerce insists on providing the same type of information to each supplier). More specifically, the CTO felt that extracting all the rents in a side arrangement would cause the best suppliers to disengage from the process entirely (i.e., look elsewhere for business), and would cause the weaker suppliers to aggressively withhold their cost information in the first P rounds of subsequent auctions.

The CTO's comments illustrate that auctions are not conducted in a vacuum, e.g., the bidders have an outside option in the form of competition with other auctioneers (procurers), and repetition causes reputation effects to come into play. These unmodeled assumptions serve as restraints upon the surplus the auctioneer can extract. However, to model these additional assumptions is an ambitious undertaking that is not attempted here. An interesting approach to these issues (perfect.com in Milgrom 2000a) is to develop efficiency-maximizing software (although these efficiency arguments assume full efficiency in the after markets) to support an entire e-marketplace, which needs to attract both auctioneers and bidders. Indeed, some of the CTO's comments are consistent with Milgrom (2000a) and Wise and Morrison (2000), among others, who warn that utility-maximizing auctions may chase suppliers away from the marketplace in the long run.

To motivate our utility-maximizing approach in round $P+1$, we note one of its important features: the utility achieved by the auctioneer is dictated by the second-best bidder. To see the significance of this observation, suppose that after the auctioneer learns the suppliers' cost function, the auctioneer announces a scoring rule and identifies the winning bidder and offers this bidder a price. If the winning bidder has some bargaining power (which is implied by the unmodeled assumptions in the previous paragraph) and is unhappy with this price, he can renege or make a counteroffer (reneging at the start of the auction is less credible than at the end when the bidder knows he is the winner). If he has all the bargaining power, the best counteroffer the bidder can quote to the auctioneer is exactly the auctioneer's outside option, which is given by the second-best bidder. Recognizing this possibility, the auctioneer tacks on a final round of the auction that will give the winning bidder what he would get anyway if he were to bargain with the auctioneer. Moreover, the winning bidder still earns the amount by which he exceeds his most able competitor, which is not unlike the outcome in many RFQs, contracts, and auctions. Hence, the fine line that our approach is trying to walk is to maximize the auctioneer's utility while generating the perception that the suppliers' disappointing profits are a result of a highly competitive marketplace.

We now turn to our analysis. With the true cost functions in hand, the auctioneer can determine his optimal scoring rule for round $P+1$. For readability, the round-subscript $P+1$ is suppressed from the scoring rule and bids in this section. Furthermore, to avoid introducing more terminology, we refer to $\hat{v}_1, \ldots, \hat{v}_A$, collectively and individually, as scoring rules, although the actual scoring rule for the auction in the final round is $\sum_{a=1}^{A} \hat{v}_a(x_a) - p$. By the MBR assumption, supplier s will submit bids that solve (1), (2). In §2.1, we imposed sufficient but not necessary, conditions on $\hat{v}_1, \ldots, \hat{v}_A$ for problem (1), (2) to possess a unique finite positive solution; if $\hat{v}_1, \ldots, \hat{v}_A$ satisfies these conditions, we refer to $\hat{v}_1, \ldots, \hat{v}_A$ as *feasible*.

Notice that supplier s will drop out of round $P+1$'s open-ascending competition no later than when his profit $p - \sum_{a=1}^{A} c_{as}(x_a^*)$ equals zero (where x_a^* is the solution to (3)), which occurs at the *maximum dropout score*

$$S_s = \sum_{a=1}^{A} \hat{v}_a(x_a^*) - \sum_{a=1}^{A} c_{as}(x_a^*). \tag{5}$$

Our analysis below, culminating in (11), only depends on the top two bidders, and we index the suppliers so that their maximum dropout scores in round $P+1$ satisfy $S_1 \geq S_2 \geq \cdots \geq S_S$. Note that this ranking depends on our choice of scoring rule. To guarantee that these two suppliers can actually bid, we impose the constraint

$$S_2 > \epsilon \tag{6}$$

on the scoring rule for round $P+1$. To keep our analysis simple, we ignore the effect of the minimum bid increment on the detailed sequence of bids, and exclude the possibility that $S_1 \leq S_2 + \epsilon$. That is, we require the scoring rule $\hat{v}_1, \ldots, \hat{v}_A$ to satisfy

$$S_1 > S_2 + \epsilon. \tag{7}$$

Depending on the detailed sequence of bids, supplier 1's winning score may lie anywhere in the interval $(S_2 - \epsilon, S_2 + \epsilon]$. To make the analysis cleaner, we assume that the winning bidder wins with a score equal to S_2. This assumption miscalculates the auctioneer's final utility by at most ϵ, which is dwarfed by the magnitude of the bids. With this assumption, the winning score in the open-ascending auction will be submitted by supplier 1 and will equal S_2; supplier 1's winning bid is the solution to

$$\max_{p, x_1, \ldots, x_A} \quad p - \sum_{a=1}^{A} c_{a1}(x_a) \tag{8}$$

$$\text{subject to} \quad \sum_{a=1}^{A} \hat{v}_a(x_a) - p = S_2. \tag{9}$$

By (4) and (5), this solution is x_{a1}^* (we now include the supplier subscript in the bids), which solves (3) for $s = 1$, and

$$\begin{aligned} p_1^* &= \sum_{a=1}^{A} \hat{v}_a(x_{a1}^*) - S_2 \\ &= \sum_{a=1}^{A} \hat{v}_a(x_{a1}^*) - \sum_{a=1}^{A} \hat{v}_a(x_{a2}^*) + \sum_{a=1}^{A} c_{a2}(x_{a2}^*). \end{aligned} \tag{10}$$

Recall that the auctioneer's true utility function is $\sum_{a=1}^{A} v_a(x_a) - p$. Using Equation (10), the auctioneer's *optimal* scoring rule in round $P+1$ is the feasible scoring rule that solves

$$\begin{aligned} \max_{\hat{v}_a} \quad & \sum_{a=1}^{A} v_a(x_{a1}^*) - \sum_{a=1}^{A} \hat{v}_a(x_{a1}^*) \\ & + \sum_{a=1}^{A} \hat{v}_a(x_{a2}^*) - \sum_{a=1}^{A} c_{a2}(x_{a2}^*), \end{aligned} \tag{11}$$

subject to constraints (5)–(7). It is possible that the auctioneer's true valuation function violates Equation (7); in spite of—or rather, because of—such "near ties," the auctioneer extracts nearly all the surplus in the auction by revealing his true valuation function as the scoring rule. In this case, solving (11) subject to (5)–(7) can do no better than simply announcing the true valuation function as the scoring rule (see the Proof of Proposition 2), and—despite its violation of (7)—we consider v to be optimal.

We reemphasize that although Equations (5)–(7) and (11) depend on only the top two of the S suppliers, the rankings of the suppliers in this optimization problem is a function of the decision variables (i.e., scoring rule). A brute force approach to the problem is to solve (5)–(7) and (11) for all $2\binom{S}{2}$ ordered pairs of suppliers, and the ordered pair that generates the highest utility in (11) provides the optimal scoring rule in round $P+1$. With the aid of Propositions 1 and 2 below, we derive a more efficient approach to this problem. In the discussion below, we do not refer to a specific scoring rule and, therefore, drop the assumption that the suppliers are ordered such that $S_1 \geq S_2 \geq \cdots \geq S_S$.

To structure the presentation, we call a bid $(p, x_1, \ldots, x_A)$ *enforceable* if there exists a feasible scoring rule $\hat{v}_1, \ldots, \hat{v}_A$, satisfying (6)–(7) that causes the auction to be won with attribute levels $x_1, \ldots, x_A$ at price p. We say that such a rule *enforces* bid $(p, x_1, \ldots, x_A)$ and utility $v(x_1, \ldots, x_A) - p$, where v denotes the auctioneer's true valuation function. For ease of presentation, we let $c_s(\vec{x}) = \sum_{a=1}^{A} c_{as}(x_a)$ and $v(\vec{x}) = \sum_{a=1}^{A} v_a(x_a)$, where $\vec{x} = (x_1, \ldots, x_A)$. Proposition 1 and all subsequent nonobvious results are proved in the online appendix at mansci.pubs.informs.org/ecompanion.html.

Proposition 1 (Enforceability). *Let supplier i be the low-cost supplier at $\vec{x} = (x_1, \ldots, x_A)$, i.e., $c_i(\vec{x}) < c_s(\vec{x})$, $s \neq i$. Let T_i be the hyperplane tangent to supplier i's cost surface c_i at $\vec{x}$, and let supplier i's profit π satisfy $\pi > \epsilon$, where ϵ is the minimum bid increment. Then, $(c_i(\vec{x})+\pi, \vec{x})$ is enforceable if, and only if, $c_s(\vec{x}) > c_i(\vec{x})+\pi$ for all $s \neq i$, and for some $j \neq i$, $T_i + \pi$ intersects supplier j's cost surface c_j (i.e., if there exists a $\vec{z}$ such that $c_j(\vec{z}) < T_i(\vec{z}) + \pi$).*

Intuition for Proposition 1 is provided in Figure 1. The four graphs of Figure 1 depict the geometry of enforcing the rightmost circle in graph (a). Graph (b) shows a tangent of c_2 that permits a half rainbow-shaped function to pass through both it and $T_1 + \pi$ in graph (c); this function, when shifted to intersect the origin, acts as an enforcing scoring rule in (d). Enforcement of the rightmost circle in (a) is achieved because of perfect competition between this bid and a dropout bid for supplier 2, namely the leftmost circle in (a). Generally speaking, for two given bids, graphs (d) and (c) visually equate perfect induced competition with the existence of a half rainbow-shaped curve passing through the bids' corresponding point-slope pairs. This geometric observation governs how the auctioneer may induce competition through choice of scoring rule, leading to the necessary and sufficient conditions for enforcement in Proposition 1, which recasts the half rainbow requirement in terms of tangent hyperplanes. The general Proof of Proposition 1 requires an attribute-by-attribute construction of an enforcing scoring rule as done for a single attribute in Figure 1, and proper modification of this rule to account for the presence of other (i.e., beyond the two pictured in Figure 1) suppliers.

Referring to graph (a) in Figure 1, Proposition 1 would hold even if π is lowered until $T_1 + \pi$ barely intersects c_2. This observation leads to the following corollary.

Corollary 1 (Minimum Prices). *Let supplier i be the low-cost supplier at $\vec{x}$ such that $c_i(\vec{x}) + \epsilon < c_s(\vec{x})$ for all $s \neq i$. If for some $j \neq i$, $T_i + \epsilon$ intersects supplier j's cost surface c_j, then enforceable prices at $\vec{x}$ approach $c_i(\vec{x}) + \epsilon$ from above, otherwise, enforceable prices at $\vec{x}$ approach $c_i(\vec{x}) + \min_{\vec{z}, s\neq i}\{c_s(\vec{z}) - T_i(\vec{z})\}$ from above.*

When applying Corollary 1 to cases in which $(p+\delta, \vec{x})$ is enforceable for $\delta \to 0^+$, we will ignore the arbitrarily small δ, and for practical purposes, consider $(p, \vec{x})$ enforceable. In words, Corollary 1 states that supplier i's profit, π, will be ϵ if $T_i + \epsilon$ intersects some c_j, and is $\min_{\vec{z}, s\neq i}\{c_s(\vec{z}) - T_i(\vec{z})\}$ otherwise. Hence, this result allows us to put the price tag $c_i(\vec{x}) + \pi$ on any A-tuple of nonprice attribute levels, which quantifies how effectively competition can or cannot be exploited via a properly chosen scoring rule. This result transforms the problem from a "what if" we choose scoring rule $\hat{v}$ approach, to an "informed shopper" approach of utility maximization given prices $c_i(\vec{x}) + \pi$. Notice that the prices themselves are not market prices in the traditional sense, but rather the prices of idiosyncratic markets distorted for hypercompetition (lowering supplier i's profit π).

Illustrations of enforceability with two suppliers are provided in Figures 2 and 3 in §2.6 for a

Figure 1 Graphical Intuition Behind the Proof of Proposition 1 for a Two-Supplier, Two-Dimensional (Price and Quality) Auction

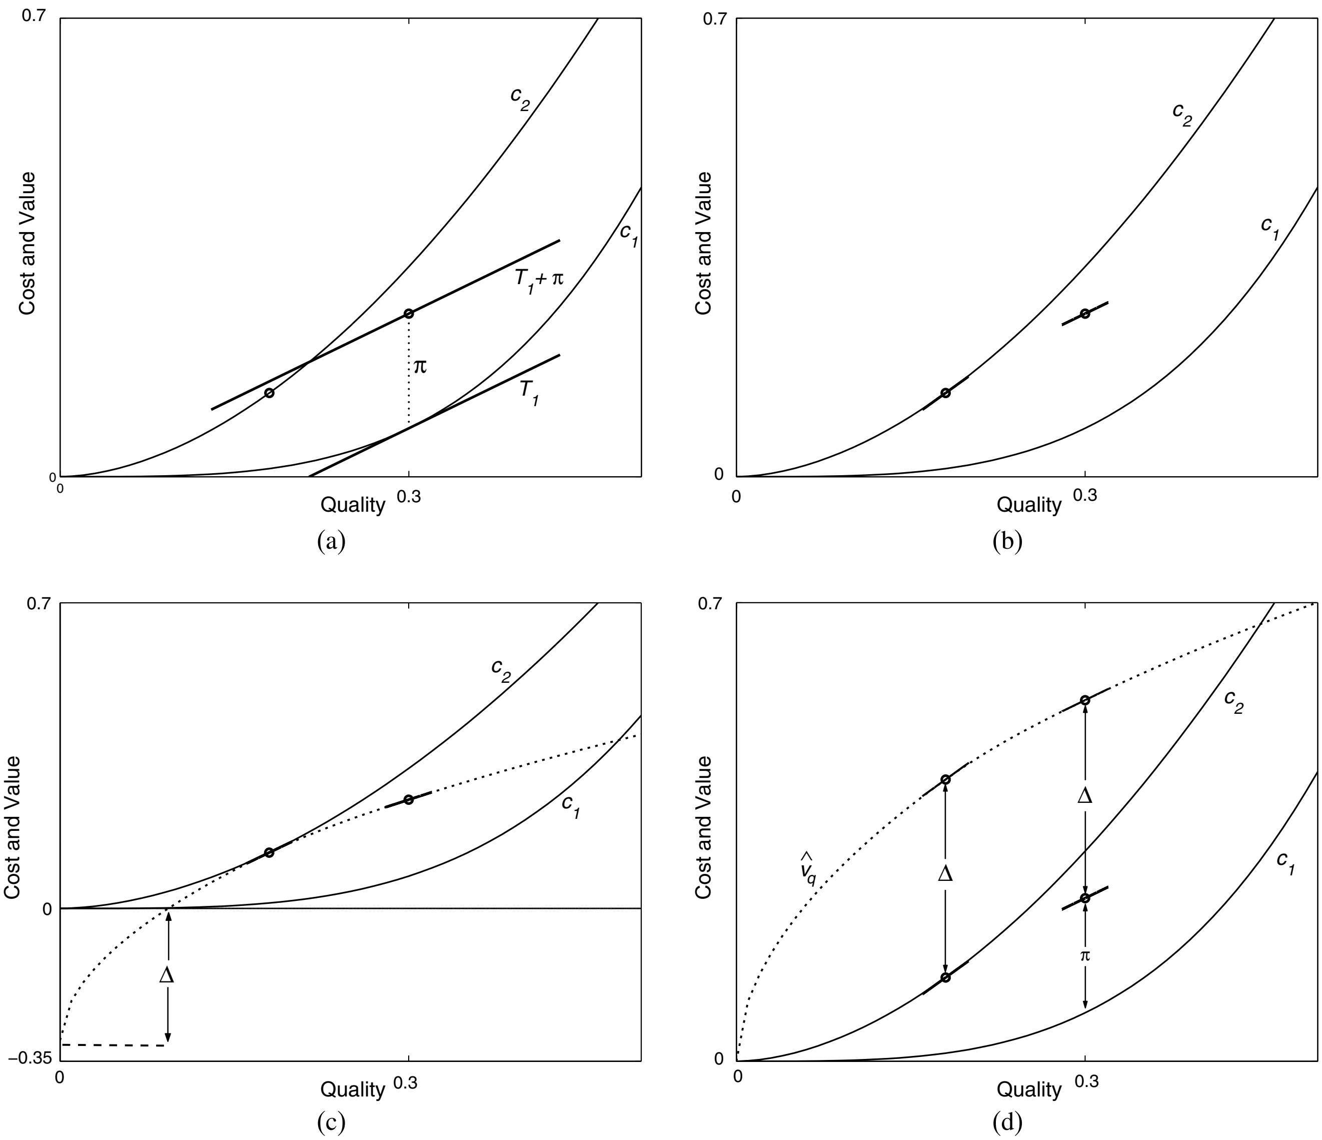

two-dimensional (price, quality) problem. Near $q = 0.86$, the prices that are enforceable in Figure 3 are lower than those of Figure 2; the tangent lines (one-dimensional hyperplanes) to c_1 near $q = 0.86$ are much closer to c_2 in Figure 3, and by Corollary 1, this permits prices much closer to supplier 1's true cost curve.

We now broaden our view and present a result that incorporates the auctioneer's utility maximization problem in (5)–(7) and (11). If supplier s wins the auction under an optimal scoring rule, we say that supplier s is *optimal*.

Proposition 2. (Restricting Search Over Suppliers). *For suppliers $s = 1, \ldots, S$, let $M_s = \max_{\vec{x}}\{v(\vec{x}) - c_s(\vec{x})\}$, where v is the auctioneer's true valuation function and c_s is supplier s's cost surface. Suppose* (*without loss of generality*) *that $M_i \geq M_s$, for all s. Then, supplier i is optimal.*

Figure 2 True Value, Supplier 1's Cost c_1, Supplier 2's Cost c_2, Minimum Enforceable Price p, and the Optimal Scoring Rule $(\cdots)$ vs. Quality

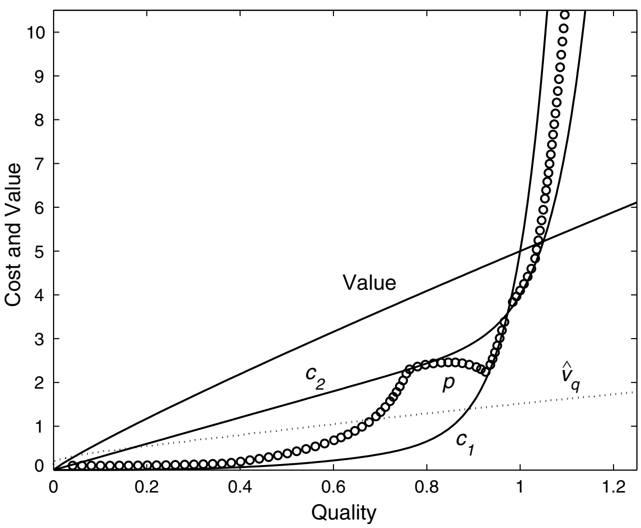

Note that M_s is supplier s's maximum dropout score if the auctioneer reveals his true valuation in the scoring rule. To describe the main intuition behind the Proof of Proposition 2, we note that the maximization problem that defines M_i has its optimal solution at some $\vec{x}^*$, at which v and c_i's tangent hyperplanes (call them T_v and T_i) are parallel. By Corollary 1, supplier i's profit is $\pi = \min_{\vec{z},\, s\neq i}\{c_s(\vec{z}) - T_i(\vec{z})\}$ if $T_i + \epsilon$ does not intersect any c_s, $s \neq i$, and $\pi = \epsilon$ if for some $s \neq i$, $T_i + \epsilon$ intersects c_s. In the former case, the auctioneer can enforce $(c_i(\vec{x}^*) + \pi, \vec{x}^*)$ and receive $M_i - \pi$ in utility. Because T_v and $T_i + \pi$ sandwich v and c_s (for $s \neq i$), respectively, and are $M_i - \pi$ units apart, this utility bounds from above any utility possible with supplier $s \neq i$ winning. In the latter case where for some $s \neq i$, $T_i + \epsilon$ intersects c_s, we can enforce price $c_i(\vec{x}^*) + \epsilon$ at $\vec{x}^*$, in which case, the auctioneer walks away with utility $M_i - \epsilon$. Because any winning supplier must receive profit of at least ϵ, the result follows. In the above, we tacitly assume that $M_i - \epsilon > M_s$ for all $s \neq i$; if not, we have the trivial case in which announcing v leads to a "near tie." To see this, note that the payoff if $s \neq i$ wins (where $M_i - \epsilon \leq M_s$) is M_s, while the payoff if i wins is no smaller than $M_i - (M_i - M_s)$. In both cases, the auctioneer's utility is at least $M_i - \epsilon$, which bounds any solution to (5)–(7) and (11) from above. The function v is considered an optimal scoring rule and i and s are both taken to be optimal suppliers.

Figure 3 True Value, Supplier 1's Cost c_1, Supplier 2's Cost c_2, Minimum Enforceable Price p, and Optimal Scoring Rule $(\cdots)$ vs. Quality for the High-Competition Example

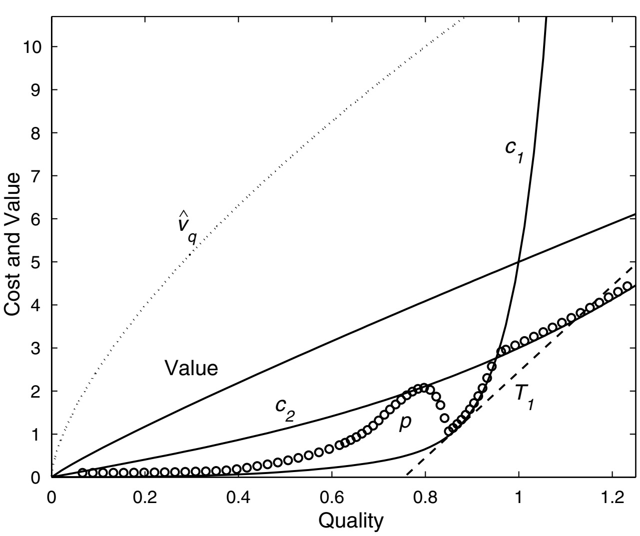

Note. T_1, the line tangent to Supplier 1's cost curve at $q = 0.8640$, intersects the cost curve of Supplier 2.

In the remainder of this subsection, we apply these results to construct a three-step method for finding the optimal scoring rule in round $P+1$. Proposition 1 and Corollary 1 reduce the problem to one of utility maximization given prices, but the prices are determined with respect to low-cost supplier i and competing supplier j. The crucial idea behind the method's first step is that Proposition 2 allows us to fix the identity of the low-cost supplier when computing prices.

Step 1 (Choose an Optimal Supplier). For $s = 1, \ldots, S$, let

$$M_s = \max_{\vec{x}} \left\{ \sum_{a=1}^{A} v_a(x_a) - \sum_{a=1}^{A} c_{as}(x_a) \right\}. \tag{12}$$

Set $i = \arg\max_{s=1,\ldots,S} M_s$; i is the optimal supplier. If there exists an $s \neq i$ such that $M_s \geq M_i - \epsilon$, announce the true valuation function in the scoring rule and exit the three-step method. Otherwise, proceed to Step 2.

Step 2 (Find a Best Competitor). Maximize utility given price, which by Corollary 1 is given by

$$\max_{\vec{x},\,\pi} \quad \sum_{a=1}^{A} v_a(x_a) - \sum_{a=1}^{A} c_{ai}(x_a) - \pi \tag{13}$$

$$\text{subject to} \quad \sum_{a=1}^{A} c_{as}(x_a) > \sum_{a=1}^{A} c_{ai}(x_a) + \epsilon, \quad s \neq i, \tag{14}$$

$$\pi \geq \epsilon, \tag{15}$$

$$\pi \geq \min_{s\neq i} \min_{\vec{z}} \left\{ \sum_{a=1}^{A} c_{as}(z_a) - T_i(\vec{z}) \right\}, \tag{16}$$

where T_i is the hyperplane tangent to supplier i's cost surface at $\vec{x}$ and π is supplier i's variable profit. In §A3, we simplify (13)–(16) by finding a closed-form solution (Equation (46) in §A3) to the innermost minimization in (16). We will call any supplier j who achieves the minimization in (16) in the optimal solution (and thereby enables the optimal solution to be enforced) a "best competitor" to supplier i. (Sections A1–A4 and Equations (24)–(47) are located in the online appendix at mansci.pubs.informs.org/ecompanion.html.)

Step 3 (Choose Optimal Scoring Rule). The derivation of this scoring rule is given in §A4. Let $\vec{x}^*$, π^* denote an optimal solution to (13)–(15) and (46). There are three cases to consider. If Equation (46) is tight at the optimal solution, then the scoring rule is the function $\hat{v}$ constructed in Claim 5 of the "$\Rightarrow$" direction Proof of Proposition 1 in §A1.2. For fixed dimension a, this optimal scoring function $\hat{v}_a$ has the form

$$\hat{v}_a(x_a) = \begin{cases} \omega_{a1} x_a^{\omega_{a2}} & \text{if } x_a \leq z_{a1}, \\ \omega_{a3}(z_{a1}-x_a)^{\omega_{a4}} + \omega_{a5}x_a + \omega_{a6} & \text{if } z_{a1} < x_a \leq z_{a2}, \\ \omega_{a7} x_a^{\omega_{a8}} & \text{if } x_a > z_{a2}, \end{cases} \tag{17}$$

where three of the 10 parameters $\omega_{a1}, \ldots, \omega_{a8}, z_{a1}, z_{a2}$ are actually redundant, but are included here to preserve readability. This function enforces optimal bid

$$\left(\sum_{a=1}^{A} c_{ai}(x_a^*) + \pi^*, \vec{x}^* \right) \tag{18}$$

in a two-supplier auction between i and j. If Equation (46) is not tight at the optimal solution, then—as explained in §A1.2—the optimal scoring rule must buffer against supplier i actually losing when $\hat{v}$ is announced to all S suppliers. The choice of optimal scoring rule in this case depends on whether the true valuation v, if used as a scoring rule, satisfies or violates constraint (6). If v satisfies (6), then the optimal scoring rule is $\lambda^*\hat{v} + (1-\lambda^*)v$, where λ^* is given in (44) at the end of §A1.2. This scoring rule enforces (18), using for each attribute the number of parameters in v_a plus 8 (7 from $\hat{v}_a$, plus the parameter λ^*). Otherwise, if (46) is nonbinding but v violates (6), then the optimal scoring rule is $\lambda^*\hat{v} + (1-\lambda^*)g$, where $\hat{v}$ is given in (17), λ^* is given in (44), and g is defined in §A1.2. The optimal scoring function in this pathological case—in which, at most, one supplier can bid below the true valuation—requires 18 parameters per attribute to enforce (18). The function g_a has a form similar to (17), but with the middle case repeated, for a total of 10 nonredundant parameters.

2.6. A Numerical Example

We illustrate our mechanism with a simple numerical example that has $S = 2$ suppliers, $A = 1$ nonprice attribute, which we call quality, and $P = 3$ cost parameters per attribute. The cost of quality is of the form $\theta_{s1}q + \theta_{s2}q^3 + \theta_{s3}q^{15}$, where we suppress the subscript a for the attribute and use q in place of x_1. More specifically, supplier 1's cost is $q^3 + 4q^{15}$ and supplier 2's cost is $3q + q^{15}$; i.e., $\theta_{11} = 0$, $\theta_{12} = 1$, $\theta_{13} = 4$, $\theta_{21} = 3$, $\theta_{22} = 0$, and $\theta_{23} = 1$. The true value function is $\psi_1 q^{\psi_2}$, where $\psi_1 = 5$ and $\psi_2 = 0.9$ (see Figure 2 for a plot of the cost and value functions). Our auction mechanism requires $P + 1 = 4$ rounds of bidding, the first three for learning and the last for optimizing. We assume that the auctioneer announces the scoring rule $2\sqrt{q}$ in round 1, $3.5q^{0.7}$ in round 2, and $4q^{0.8}$ in round 3. These scoring rules were arbitrarily chosen, but do satisfy the requirements of §2.1 and the requirement of §2.4 related to nonredundant bids (see below). Later in §3.3, we consider a systematic approach for determining these rules. We also assume that the minimum bid increment is $\epsilon = 0.1$. Supplier s's undistorted bids satisfy (3), which for round r is

$$\theta_{s1} + 3\theta_{s2}(q_{sr}^*)^2 + 15\theta_{s3}(q_{sr}^*)^{14} = \phi_{2r}\phi_{1r}(q_{sr}^*)^{\phi_{2r}-1}, \tag{19}$$

where the scoring rule in round r is denoted $\phi_{1r}q^{\phi_{2r}}$. In round 1, supplier 2's undistorted bids have quality $q_{21}^* = 0.6265$, and supplier 1's quality is $q_{11}^* = 0.1111$. Similarly, in round 2, supplier 2's quality is

$q^*_{22} = 0.7466$ and supplier 1's quality is $q^*_{12} = 0.5085$; in round 3, $q^*_{23} = 0.7714$ and $q^*_{13} = 0.7681$. For fixed s, Equation (19) for rounds 1–3 yield a linear system with three unknowns. For supplier 2, the system is

$$\begin{pmatrix} 1 & 1.1774 & 0.0215 \\ 1 & 1.6721 & 0.2506 \\ 1 & 1.7852 & 0.3963 \end{pmatrix} \begin{pmatrix} \hat{\theta}_{21} \\ \hat{\theta}_{22} \\ \hat{\theta}_{23} \end{pmatrix} = \begin{pmatrix} 1.2634 \\ 2.6745 \\ 3.3705 \end{pmatrix}. \tag{20}$$

Solving (20), the estimated cost values $\hat{\theta}_{21}$, $\hat{\theta}_{22}$, $\hat{\theta}_{23}$ coincide with the true values of θ_{21}, θ_{22}, θ_{23}. The equations for supplier 1 are omitted, but the analysis is identical.

With knowledge of these cost parameters, the auctioneer chooses the optimal scoring rule $\lambda\hat{v} + (1-\lambda)v$ for the final round of bidding (v satisfies (6)). In the first step, the auctioneer finds $M_2 = 1.6823$, which is smaller than $M_1 = 3.4375$, hence, supplier 1 is optimal. Because M_1 is more than ϵ units greater than M_2, we solve (13)–(15) and (46) to find q^*, π^*, the optimal quality level and profit at which supplier 1 wins the auction. Because $(c_2')^{-1}(x) = ((x-3)/15)^{1/14}$, and $c_2'(0) = 3$ and $c_1'(0.7593) = 3$, we solve

$$\max_{q\geq 0,\,\pi} \quad \psi_1 q^{\psi_2} - \theta_{11}q - \theta_{12}q^3 - \theta_{13}q^{15} - \pi$$

subject to

$$\theta_{21}q + \theta_{22}q^3 + \theta_{23}q^{15} > \theta_{11}q + \theta_{12}q^3 + \theta_{13}q^{15} + \epsilon,$$

$$\pi \geq \epsilon,$$

$$\pi \geq \theta_{21}\hat{q} + \theta_{22}\hat{q}^3 + \theta_{23}\hat{q}^{15} - \left[\theta_{11}q + \theta_{12}q^3 + \theta_{13}q^{15}\right] + \left[\hat{q} - q\right]\left[\theta_{11} + 3\theta_{12}q^2 + 15\theta_{13}q^{14}\right],$$

$$\hat{q} = \begin{cases} 0 & \text{if } q \leq 0.7593 \\ \left(\dfrac{\theta_{11} + \theta_{12}\cdot 3q^2 + \theta_{13}\cdot 15q^{14} - 3}{15}\right)^{1/14} & \text{if } q > 0.7593 \end{cases}.$$

The minimum enforceable prices (per Corollary 1) are shown in Figure 2. An exhaustive search (with a discretization grid of 0.0001) yields $q^* = 0.6238$ and $\pi^* = 0.5327$. In the third and final step, the construction in §A1.2 produces parameter values $\lambda^* = 1$, $w_1 = 0.4194$, $w_2 = 0.1304$, $w_3 = 47.7807$, $w_4 = 13.0341$, $w_5 = 1.2484$, $w_6 = 0.3000$, $w_7 = 1.5167$, $w_8 = 0.7219$, $z_1 = 0.0100$, and $z_2 = 0.6238$ (see Figure 2). (The rule constructed above enforces $(c_1(q^*) + \pi^* + \epsilon, q^*)$, though actually prices arbitrarily close to (but greater than) $c_1(q^*) + \pi^*$ can be enforced.) Under this scoring strategy, the losing supplier's (supplier 2) quality level is 0.01, the auctioneer's utility is $\psi_1(q^*)^{\psi_2} - \theta_{11}q^* - \theta_{12}(q^*)^3 - \theta_{a_{13}}(q^*)^{15} - \pi^* - \epsilon = 2.3909$, the winning supplier's profit is $\pi^* + \epsilon = 0.6327$, and the total surplus is 3.0236.

To illustrate the effects of inducible competition, we next consider a "high-competition" example in which the parameters for supplier 1's cost and the true value functions are unchanged, but we set $\theta_{21} = 2$, $\theta_{22} = 1$, and $\theta_{23} = 0$; the cost curves, the true value function, and the resulting minimum enforceable prices are shown in Figure 3. Running the optimization (13)–(15) and (46) under the new parameters for supplier 2 results in $q^* = 0.8482$ and $\pi^* = 0.1000 = \epsilon$. If we enforce $(c_1(q^*) + \pi^* + \epsilon, q^*)$ (see Figure 3), the payoff to the auctioneer is 3.1627, which is roughly 35% greater than in the previous example. Referring to Figure 3, notice that c_2 crosses the line tangent to c_1 at c_1's "elbow," which allows the auctioneer to enforce near-cost prices just past where the elbow starts.

We see that, depending on the situation, the optimal scoring rule can downplay (Figure 2) or overstate (Figure 3) the true valuation of quality. In summary, our scoring rule can operate in a fundamentally different manner than Che's (1993) rule. Che's optimal scoring rule understates the true value of quality to reduce the information rents received by the more cost-efficient suppliers. In contrast, the auctioneer in our mechanism knows the suppliers' cost functions before round $P+1$, and the optimal scoring rule might even exaggerate the value function. Some practical implications of such a scoring rule are discussed in §3.1.

To put our proposed mechanism into perspective, we also consider a "straw" mechanism, where the auctioneer announces his true utility function as the scoring rule, and assumes that the suppliers submit their MBR bids. This scenario adheres to Equations (5)–(7) and (11), but with $\psi_1 q^{\psi_2}$ in place of $\hat{v}$; we consider the first "low competition" example. supplier s's maximum dropout score is M_s, so supplier 1 wins the auction bidding quality $q^* = 0.8008$ (the solution to the right-hand side of (12) for $s = 1$) at price $c_1(q^*) + M_1 - M_2$. (By Step 1 of §2.5, we note that the winning supplier in our proposed mechanism will be the same as the winning supplier in the straw mechanism, but the winning bids will likely be different.) The auctioneer's utility is $v(q^*) - c_1(q^*) - (M_1 - M_2) = M_2 = 1.6823$, supplier 1's profit is 1.7552, and the total surplus is 3.4375. Hence, as expected (see Milgrom 2000a), the proposed mechanism leads to an increase in the auctioneer's utility, and a decrease in efficiency relative to the straw mechanism. However, an industrial data set would be required to assess the magnitude of utility enhancement that might be achieved in practice. Such an assessment is beyond the scope of this paper.

3. Practical Considerations

In §2, we made two restrictive assumptions: (1) the suppliers do not distort their bids in rounds $1, \ldots, P$ and they bid their MBR in round $P+1$, and (2) the auctioneer knows the form of the suppliers' cost functions, but not their parameter values. In practice, some bidders may distort their bids or not conform to MBR, and the auctioneer may not know the form of the suppliers' cost functions. In this section, we discuss several enhancements to the proposed mechanism, which either improve its robustness with respect to these two assumptions or address other practical concerns. We begin with a discussion of the scoring rule in the last round.

3.1. Scoring Rule in the Last Round

There are three potential problems with the analysis of §2.5: (1) the determination of the optimal enforceable bid may require solving a nonconvex math program, (2) the resulting scoring rule may be too complex (for each attribute, 8 plus the number of parameters in v_a if v satisfies (6), 18 otherwise) for practical implementation, and (3) the scoring rule may force the losing supplier to submit bids with negligible nonprice attribute levels. To deal with the first problem, Step 2 of our method can be replaced by restricting (13)–(15) and (46) to $\vec{x} = \vec{\hat{x}}$, where $\vec{\hat{x}}$ maximizes the right side of (12) for $s = i$. This alternative Step 2 searches for the lowest possible enforceable price at $\vec{\hat{x}}$, the bid level with the potential to yield the largest utility for the auctioneer. While this alternative Step 2 is not guaranteed to find the optimal scoring rule, it will do so if any supplier's cost surface intersects the hyperplane tangent to supplier i's cost surface at $\vec{\hat{x}}$. Furthermore, the Proof of Proposition 2 shows that, though not necessarily optimal, the rule generated using this alternative Step 2 yields an auctioneer's utility that is greater than or equal to the utility level from any auction in which supplier i is not the top bidder, i.e., even with this simplified approach, we are still sure to do better than is possible via full optimization over generic scoring rules in which supplier $s \neq i$ wins.

The complexity of the scoring rule (see (17)) can perhaps be finessed in practice by providing the scoring rule in graphical form (one graph per nonprice attribute), together with a calculation device that converts uncommitted bids into scores. An alternative approach is to employ a parametric scoring rule in Step 3. With the identity of suppliers i and j (a best competitor to i) in hand from the method's first two steps, the auctioneer's scoring rule selection problem becomes a maximization over the scoring rule parameters, which we denote $\vec{\phi}_1, \ldots, \vec{\phi}_A$:

$$\max_{\vec{\phi}_a} \quad \sum_{a=1}^{A} v_a(x^*_{ai}) - \sum_{a=1}^{A} \hat{v}_a(x^*_{ai}; \vec{\phi}_a) + \sum_{a=1}^{A} \hat{v}_a(x^*_{aj}; \vec{\phi}_a) - \sum_{a=1}^{A} c_{aj}(x^*_{aj}) \tag{21}$$

$$\text{subject to} \quad S_s = \sum_{a=1}^{A} \hat{v}_a(x^*_{as}; \vec{\phi}_a) - \sum_{a=1}^{A} c_{as}(x^*_{as}), \quad s = i, j, \tag{22}$$

(6)–(7), and the scoring rule constraints at the end of §2.1 (the scoring rule constraint mentioned later in §2.4 is superfluous here). Because i and j were selected with respect to a generic scoring rule, they are not guaranteed to be optimal for the parameterized case. However, this drawback—though difficult to quantify a priori—compensates for the need to solve $2\binom{S}{2}$ mathematical programs (a version of (6)–(7) and (21)–(22) for every ordered supplier pair), which may not be practical if S is large. In practice, some approach midway between these two extremes could be used; for instance, examining all pairs from the top 10% of candidates in Steps 1 and 2.

In addressing the third potential problem, we first note that our scoring rule restrictions in §2.1 and §2.4 (i.e., the scoring rules are strictly concave and satisfy the conditions on the derivatives at zero and infinity to ensure a unique, interior MBR bid response by suppliers) are not innocuous. In the absence of these constraints, we have constructed nonpathological cases in which the optimal scoring rule in the last round is convex or closely mimics Step 1's supplier i's cost surface everywhere except near $\hat{\vec{x}}$, where it is precisely ϵ units higher. While ignoring these constraints may increase the auctioneer's utility in the short run, in the longer run, cost mimicking can eliminate meaningful bids from the non-low-cost suppliers, compromising competition in—and, consequently, the credibility of—the auction. Furthermore, a nonconcave scoring rule may make transparent the strategic nature of our proposed mechanism. While we have been careful to avoid cost mimicking and convex scoring rules, we note that the optimal scoring rule in the last round can force the best competitor (i.e., supplier j in Step 2) to submit bids with negligible nonprice attribute levels, e.g., in our first numerical example, supplier 2's quality level is 0.01. This phenomenon may arouse bidder suspicion, and it may be shrewd for the auctioneer to either add a lower-bound constraint on the MBR attribute levels that result from his scoring rule or impose a reservation level for all nonprice attributes.

3.2. Cost Estimation

Each time a supplier submits a bid, he generates a new version of Equation (3) that the auctioneer can use to estimate the suppliers' cost parameters. As noted earlier, if these bids are undistorted, then the vector of nonprice attributes is identical for all bids within a given round, and after P rounds the P first-order equations determine the cost parameters for each attribute. However, if bidders intentionally (i.e., strategically) or unintentionally (e.g., lack of sophistication) distort their bids, then an inconsistent set of first-order conditions is generated. If we define

$$f_{asr}(x_a) = \frac{\partial v_a}{\partial x_a} - \frac{\partial c_{as}}{\partial x_a},$$

then the first-order condition in (3) is $f_{asr}(x_a) = 0$. Let us now consider a fixed attribute a and a fixed supplier s and suppress these subscripts, and suppose that B_r bids are submitted in round r. Note that B_r is likely to be a small number because bids in RFQ processes require more work on the bidders' part than in a traditional price-only auction. Then, for bid b in round r, let the first-order condition be given by $f_{rb}(x_{rb}) = 0$. If bid distortion occurs, we can estimate the unknown cost parameters by choosing $\vec{\theta}_{as}$ to minimize the weighted-average least-squares quantity, $\sum_{r=1}^{P}\sum_{b=1}^{B_r} w_{rb} f_{rb}^2(x_{rb})$, where the weights satisfy $\sum_{b=1}^{B_r} w_{rb} = 1$ for all r, and $w_{rb} \geq 0$. This quadratic program might also incorporate convexity constraints on the parameter values. Note that the weights $w_{rb} = 1/B_r$ for all r and b would minimize the variance for fitting the curve $f_{rb}(x_{rb}) = 0$ in the case where the true model was $f_{rb}(x_{rb}) = \epsilon_{rb}$, where ϵ_{rb} are iid normal, mean-zero random variables. However, if we believe that bids based on bad judgment or experimentation are less likely to occur as each round proceeds, then later bids within each round should be assigned higher weights. Finally, more complex methods, which employ Bayesian a priori estimates on the output data (e.g., ϵ_{rb}) and the parameter values have been developed in the geophysics field (Tarantola 1987).

This same weighted-average least-squares procedure can be used to choose among several alternative cost functions. For example, we can simultaneously compute first-order conditions for two cost functions, $\theta_{as1}x^{\theta_{as2}}$ and $\theta_{as1}x + \theta_{as2}x^2$, and use the option that leads to a more consistent sequence of first-order conditions as measured by $\sum_{r=1}^{P}\sum_{b=1}^{B_r} w_{rb} f_{rb}^2(x_{rb})$.

3.3. Scoring Rules in Earlier Rounds

Under the assumptions in §2, accurate parameter estimation will be achieved as long as the auctioneer announces P distinct scoring functions in the first P rounds that satisfy the conditions stated in §2.4 and at the end of §2.1. However, in practice, large fluctuations in the auctioneer's scoring rule across rounds may cause the bidders to suspect strategic behavior on the part of the auctioneer, which, in turn, may lead the bidders to counter with their own strategic behavior (e.g., intentional cost distortion). At the other extreme, a minuscule change in the auctioneer's scoring rule may not cause any change in the suppliers' bids, perhaps because the nonprice attributes—in contrast to our model's assumptions—take on only a discrete set of values. In our view, the auctioneer's goal in the early rounds is to make changes in the scoring rule that are as subtle as possible, while still generating new bids from the suppliers. Ideally, these changes are made in such a way that the bidders have the impression that the auctioneer is tweaking his scoring rule for nonstrategic reasons (e.g., an improved understanding of the relative importance of the attributes).

The choice of the initial scoring rule is the most difficult, because the auctioneer is assumed to possess no bidding information. As a general guideline, the auctioneer should use his experience and historical data to choose an initial scoring function that is close to the predicted final scoring rule.

We propose the following procedure for choosing the scoring rule in rounds $2, \ldots, P$, which makes effective use of the bidding information from the previous rounds. This procedure is described in four steps, the first three of which are devoted to finding (possibly inaccurate) estimates for the suppliers' cost parameters. First, to determine the scoring rule in round r, we assume that all cost functions have no more than r parameters, e.g., even if we plan to use three-parameter cost functions and four rounds of bidding, to determine the scoring rule in round 2, we assume that all cost functions have only two parameters. Second, in addition to the first-order (i.e., undistorted bid) conditions in (3), we also assume that suppliers submitted their MBR bids in round $r-1$. In particular, this assumption implies that, for all but the highest bidder, a bidder's final score in round $r-1$ is his dropout score. If we denote such a bidder as supplier s and his bid as $(p_s^*, x_{1s}^*, \ldots, x_{As}^*)$, then this assumption generates the additional equation

$$p_s^* = \sum_{a=1}^{A} c_{as}(x_{as}^*). \tag{23}$$

If all bids are, indeed, undistorted, and supplier s submitted MBR bids in round $r-1$, then Equation (23) can be combined with the $r-1$ first-order equations from the earlier rounds to uniquely solve for supplier s's true cost parameters. If, instead, some earlier bids are distorted (see §3.2), then we propose minimizing $\sum_{r=1}^{P}\sum_{b=1}^{B_r} w_{rb} f_{rb}^2(x_{rb})$, subject to (23) and the additional constraints $\sum_{b=1}^{B_r} w_{rb} = 1$, $w_{rb} \geq 0$.

We only have censored information, namely $p_1^* \geq \sum_{a=1}^{A} c_{a1}(x_{a1}^*)$, for the highest bidder in round $r-1$. That is, we do not know this bidder's dropout score. Consequently, in the third step of our procedure, we assume that the current first-best-bidder has a dropout score that is a fixed percentage (e.g., 5% or 10%) higher than the second-best-bidder's observed dropout score. This fixed percentage should include the estimated gap between the top two bidders and the perceived amount by which the second-best-bidder is "holding back" (i.e., his true dropout score minus his observed dropout score). This estimated dropout score provides the additional equation (see (5)) to estimate the current first-best-bidder's cost parameters. In the final step of our procedure, we substitute these (possibly inaccurate) parameter estimates into the three-step (optimal scoring rule determination) method of §2.5, and find the proposed scoring rule for round r.

3.4. Activity and Transition Rules

In the absence of activity rules, bidders are unlikely to bid aggressively in the earlier rounds of our mechanism. Activity rules are often imposed in auctions (see, e.g., Kelly and Steinberg 2000 and Milgrom 2000b for details pertaining to Federal Communications Commission auctions) to prevent bidders from delaying their bid submissions until the end of the auction. Similarly, weak bidding suppliers are often weeded out throughout the course of a RFQ process. We propose that after each round of the auction, the auctioneer allows only a subset of suppliers to proceed to the next round. The criteria could be based on either the number of suppliers (e.g., only five suppliers compete in round 2 and only three suppliers compete in round 3) or on their scores (e.g., only suppliers with scores within a fixed percentage of the current leader may proceed to the next round). Our mechanism also requires a rule for transitioning to the next round. This rule could be time based (e.g., each round lasts a certain number of days) and/or activity based (e.g., a round terminates after a certain number of bid-free days). Finally, to encourage competition down the homestretch, we propose using activity-generated overtime periods in the final round.

3.5. Exogenous Attributes

In addition to bid price and endogenous attributes such as quality and lead time, exogenous attributes such as a supplier's reputation and his past history with the manufacturer typically play a vital role in the allocation decision. These factors are easily incorporated into our model. Let e_s represent the auctioneer's total utility derived from supplier s's exogenous attributes. For an incumbent supplier, this utility might incorporate the fixed cost to switch to a different supplier. Then, the auctioneer's true utility function (with the supplier notation suppressed) becomes $\sum_{a=1}^{A} v_a(x_a) + e - p$.

We recommend that the auctioneer reveals to supplier s his truthful exogenous value e_s, but not the other suppliers' exogenous values. While we have not attempted to prove that truthful revelation is optimal on the auctioneer's part, withholding all information about e_s would be unsatisfactory to the suppliers because they would only possess a partial scoring function. Moreover, a large portion of the exogenous value is likely to be based on standardized supplier ratings, which are readily available in many industries.

Under the assumption that the true e_s is revealed to supplier s, the analysis extends in a straightforward manner. The supplier's cost surface c_s is simply vertically shifted by $-e_s$ units. This shift is allocated to the costs over individual attributes by taking c_{as} to be vertically shifted $-\lambda_{as} e_s$ units, where $\lambda_{as} \geq 0$ and $\sum_{a=1}^{A} \lambda_{as} = 1$. It is even possible for a supplier's e_s value to change during the course of the eRFQ process, e.g., by delivering an unexpectedly impressive presentation or by providing perks such as tickets to sporting or cultural events. Although by strategically assigning exogenous attribute levels, the auctioneer can contrive to enforce any A-tuple at ϵ profit, generating competition in this way is likely to be much more obvious to suppliers than relying on scoring rules as described in §2.5.

4. Concluding Remarks

One of the most difficult aspects of running a multiattribute procurement auction from the auctioneer's perspective is the lack of knowledge about the suppliers' cost functions for endogenous nonprice attributes. We develop an auction mechanism that is in the same spirit as other dynamic strategies for problems with imperfect information (e.g., Gittens 1989), which first focuses on learning the relevant information and then switches to an optimization mode after sufficient learning has occurred. In our model, we use inverse-optimization techniques to learn the suppliers' cost functions. Although optimization-based (or smart-market) mechanisms have been in use for nearly a half century (Stanley et al. 1954), existing studies have all used forward optimization, and this paper appears to be the first to use an inverse-optimization-based approach. Considering the ease with which individualized data can be collected on the Internet and the fact that an auction's outcome often depends on the parameter values of only a few bidders, individualized learning using MBR and inverse optimization may be more fruitful than "collective" learning, where the bidding population's parameters are probabilistically modeled. This inverse-optimization-based approach may be applicable in other types of auctions, and perhaps other settings of learning in games (Fudenberg and Levine 1998). Moreover, consultants have argued that understanding the cost drivers of purchased items is the most fundamental capability of an effective sourcing strategy (e.g., Laseter 1998, p. 6), and this approach could even be used solely for cost estimation purposes. Although our analysis uses elementary techniques, considerable theory has been developed in recent years for various aspects of inverse optimization in mathematical programming (e.g., Ahuja and Orlin 2001 and references therein) that may be useful for more complex problems.

Aside from transaction cost savings, the prospect for competition is what makes a procurement auction compelling for the buyer. Our (largely geometric) analysis in §2 shows how the auctioneer, via the choice of the scoring rule in the last round, can manipulate the rules of the competition so as to maximize his own utility within the open-ascending auction format. In particular, it is optimal to first identify the winning supplier, which is the one with the largest dropout score (see Equation (5)) if the auctioneer revealed his true valuation in the announced scoring rule, and then to identify his best competitor, i.e., the one that will minimize the winning supplier's profit, thereby leaving more utility for the auctioneer. Ideally, all suppliers exit the auction with a renewed sense of respect and fear for the opposing suppliers, rather than feeling as if they have been manipulated by the auctioneer.

While our numerical examples consider only one nonprice attribute, our analysis has the potential to provide nonobvious insights about which attributes and competing suppliers provide the most fruitful focus of competition.

Section 3 discusses several important practical issues that bring this mechanism closer to practice. The CTO of Frictionless Commerce shared the basic ideas of our mechanism with several key customers. While they found the ideas intriguing, they did not seem ready to use it (even assuming it had undergone successful human testing prior to release) for two reasons. First, several customers (i.e., manufacturers) did not feel comfortable with the complexity of the mechanism. In this regard, we agree that the mechanism is much less transparent than the efficiency-maximizing mechanism. Second, although traditional RFQ processes allow the changing of the scoring rule as the process proceeds, one customer (from the private sector) felt that an equity issue would arise if activity rules (see §3.4) were in place. In his words, "a supplier would get upset if he was thrown out of the auction when the score was based on apples, but would have done better when the score was later changed to oranges." This suggests that activity rules must strike a delicate balance between the perception of fairness and the mitigation of strategic behavior. More generally, the CTO thought that major changes in the scoring rule of a private-sector auction would likely fuel the perception that the manufacturer was "beating up" the vendors, which suggests that the discussion of the early-round scoring rules in §3.3 is of particular importance. In summary, the CTO conjectured (in the winter of 2001) that the market would not be ready for this type of mechanism for another 1–2 years. He also thought that it might make sense to first attempt to implement the cost-estimation portion of our analysis as an extra software feature that allows the auctioneer to gain valuable information (e.g., estimating cost parameters, assessing the magnitude of bid distortion) without committing to a new eRFQ mechanism. Finally, before this mechanism could be practically implemented, it would need to incorporate discrete-valued attributes and nonsmooth cost and utility functions.

Acknowledgments
The authors thank Rob Guttman, former Chief Technology Officer of Frictionless Commerce Incorporated, for invaluable input during the modeling phase of our project. They also thank Opher Baron, Ani Dasgupta, Jeremie Gallien, the associate editor, and referees for helpful comments on an earlier draft of this manuscript.

References
Ahuja, R. K., J. B. Orlin. 2001. Inverse optimization. *Oper. Res.* **49** 771–783.

Bikhchandani, S. 1999. Auctions of heterogeneous objects. *Games Econom. Behavior* **26** 193–220.

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.

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

Demange, G., D. Gale, M. Sotomayor. 1986. Multi-item auctions. *J. Political Econom.* **94** 863–872.

Fudenberg, D., D. K. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.

Gallien, J., L. M. Wein. 2000. Design and analysis of a smart market for industrial procurement. *Management Sci.* Forthcoming.

Gittens, J. C. 1989. *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, New York.

Kafka, S. J., B. D. Temkin, L. Wegner. 2000. B2B auctions go beyond price. Forrester Research, Inc., Cambridge, MA.

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

Laseter, T. M. 1998. *Balanced Sourcing*. Jossey-Bass, San Francisco, CA.

Milgrom, P. 2000a. An economist's vision of the B-to-B marketplace. Executive white paper, www.perfect.com.

——. 2000b. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* **108** 245–272.

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

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

Rochet, J.-C., L. A. Stole. 2001. The economics of multidimensional screening. Graduate School of Business, University of Chicago, Chicago, IL.

Stanley, E. D., D. P. Honig, L. Gainen. 1954. Linear programming in bid evaluation. *Naval Res. Logist. Quart.* **1** 48–54.

Tarantola, A. 1987. *Inverse Problem Theory*. Elsevier, Amsterdam, The Netherlands.

Wellman, M. P., W. E. Walsh, P. R. Wurman, J. K. MacKie-Mason. 2001. Auction protocols for decentralized scheduling. *Games Econom. Behavior* **35** 271–303.

Wise, R., D. Morrison. 2000. Beyond the exchange: The future of B2B. *Harvard Bus. Rev.* (November–December) 86–96.

Accepted by Arthur Geoffrion and Ramayya Krishnan; received January 2002. This paper was with the authors 12 months for 3 revisions.