An Inverse-Optimization-Based Auction Mechanism

To Support a Multi-Attribute RFQ Process

Paper 143

Damian R. Beil
Lawrence M. Wein

December 2001

AN INVERSE-OPTIMIZATION-BASED AUCTION MECHANISM TO SUPPORT A MULTI-ATTRIBUTE RFQ PROCESS

Damian R. Beil and Lawrence M. Wein

Operations Research Center, MIT, Cambridge, MA 02142, beil@mit.edu

Sloan School of Management, MIT, Cambridge, MA 02142, lwein@mit.edu

Abstract

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 non-price attributes (e.g., quality, lead time). The manufacturer is assumed to know the parametric form of the suppliers' cost functions (in terms of the non-price attributes), but has no prior information on the parameter values. We construct a multi-round 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 non-price 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.

December 6, 2001

1. INTRODUCTION

Although the market for online business-to-business auctions is enormous (estimated at \$746B in 2004 by Kafka *et al.* 2000), the price-only auctions that dominate the current eCommerce 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, eMarketplaces are currently being developed to partially automate the RFQ process; i.e, to create an eRFQ process (see Kafka *et al.* 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 multi-attribute eRFQ mechanism. The CTO described the company's multi-attribute procurement software (we are not at liberty to discuss its details) and the perceived needs and preferences 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 its ideas with several customers, and their impressions are briefly summarized in §4.

The appropriate mathematical context for this setting is the *multi-attribute*, 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, revenue maximization (on the part of the auctioneer) and (allocative) efficiency; in multi-attribute 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 it the most.

To engage bidders in a multi-attribute auction, an auctioneer needs to provide the bidders with some information pertaining to how he values the non-price 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.

There are two key papers on multi-attribute auctions with scoring rules, one addressing each objective. Milgrom (2000a) has recently 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 is by Che (1993), who considers a two-dimensional price-quality procurement auction in a sealed-bid setting. He assumes that the suppliers' quality costs 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). 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.

This paper is motivated by two opinions regarding this literature: (i) utility maximization is the appropriate objective in many industrial procurement settings, and (ii) Che's results, while elegant and insightful, are not practically implementable. More specifically, efficiency may be an appropriate objective if the goal (as is the case of Perfect.com in Milgrom 2000a) is to develop software to support an entire eMarketplace, which needs to attract both auctioneers and bidders. Moreover, Milgrom (2000a) and Wise and Morrison (2000), among others, warn that utility-maximizing auctions may chase suppliers away from the marketplace in the long run. Nonetheless, it is our view that for a large manufacturer that either develops its own procurement auction software in-house or buys it from an external vendor, utility maximization on the part of the auctioneer is the appropriate objective, at least over the short term and medium term. Also, in our utility-maximizing mechanism, 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. Moving to our second point, even if Che's utility-maximizing results could be extended to multiple attributes, we believe his key informational assumption – that the auctioneer knows the probability distribution of the bidders' cost for non-price attributes – is not amenable to practical implementation. That is, we do not think it is possible for a manufacturer to have a precise *a priori* estimate of the suppliers' cost functions and to reliably summarize them as a probability distribution over one or several parameters. Consequently, any attempt to strategically change the scoring rule based on this estimated probability distribution would have a reasonable likelihood of backfiring (i.e., of resulting in lower utility for the manufacturer), particularly because the optimal scoring rule depends only on the top few bidders' cost functions, which requires a precise estimate of the tail of the probability distribution.

Consequently, our goal in this research is to employ Che's strategic ideas while simultaneously enabling the manufacturer to *learn* the suppliers' cost functions. This paper proposes 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.

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 non-committal 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. The majority, but not all, RFQ processes in practice are for a single item. Because our notation requires up to five subscripts, we use mnemonic subscripts, where $a = 1, \ldots, A$ indexes the attributes, s indexes the S suppliers, p indexes the P cost (and utility) parameters per attribute, and $r = 1, \ldots, P+1$ 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 multi-attribute 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 non-price attributes are endogenous, and defer a discussion of exogenous attributes to §3.5. Each bid is of the form $(p, x_1, \ldots, x_A)$, where p denotes price and x_a is the magnitude of non-price attribute a for $a = 1, \ldots, A$. To simplify the presentation and analysis, we assume that the non-price 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 $\sum_{a=1}^{A} c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP})$, 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; this assumption is easily relaxed, as described at the end of §2.4. 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 $(\theta_{as1}, \ldots, \theta_{asP})$. This issue is revisited in §3.2, where we propose a procedure that chooses among alternative functional forms.

The auctioneer's true utility function and scoring rules are assumed to be additive across non-price attributes, and (except for the scoring rule in round $P + 1$) to consist of exactly P parameters for each non-price attribute. While the separability across attributes of the cost and utility functions and scoring rules makes the problem more tractable, existing multi-attribute software packages also make these assumptions, and nonseparable functions would likely be too arduous for industrial implementation. The true utility function is given by $\sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP}) - p$, 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, or utility, parameters per attribute. The scoring rule in round $r = 1, \ldots, P$ is denoted by $\sum_{a=1}^{A} v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr}) - p$. Notice that, for simplicity, we assume that the scoring rules prior to the final round are identical in form to the true utility function for a given attribute a, but may employ different parameter values. We relax this assumption for the scoring rule in the final round (as discussed in §2.5), and denote the generic (unparameterized) scoring rule in round $P+1$ as $\sum_{a=1}^{A} f_a(x_a) - p$. However, for fixed a and r, the scoring vector $(\phi_{a1r}, \ldots, \phi_{aPr})$ and the function f_a must be such that the scoring rule is increasing and concave in x_a. We also assume that $\frac{\partial c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial x_a} < \frac{\partial v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr})}{\partial x_a}$ as $x_a \to 0^+ \ \ \forall s$, and $\frac{\partial v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr})}{\partial x_a} \to 0$ as $x_a \to \infty$, to guarantee that the solution to the optimization problem in (1)-(2) possesses a finite positive solution; the same is assumed for the generic scoring rule in round $P+1$, i.e., $\frac{\partial c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial x_a} < \frac{df_a(x_a)}{dx_a}$ as $x_a \to 0^+ \ \ \forall s$, and $\frac{df_a(x_a)}{dx_a} \to 0$ as $x_a \to \infty$. One final, more technical assumption on only the scoring vectors $(\phi_{a1r}, \ldots, \phi_{aPr})$, $r = 1, \ldots, P$, 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 $r = 1, \ldots, P$, the auctioneer announces the scoring rule $\sum_{a=1}^{A} v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr}) - p$ to all suppliers. Suppliers then submit bids of the form $(p, x_1, \ldots, x_A)$ in an open-ascending manner; the same process is repeated in the final round when scoring rule $\sum_{a=1}^{A} f_a(x_a) - p$ is announced. 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 in order to proceed to the next round (at this point in the paper, we do not rule out the possibility that a supplier simply re-submits 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: how the suppliers bid given the current scoring rule and current best score, how the auctioneer estimates the suppliers' cost functions given their bids, and 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 was ending after his bid. More specifically, if in round $r \leq P$ 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}):

$$\max_{p,x_1,\ldots,x_A} \quad p - \sum_{a=1}^{A} c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP}) \tag{1}$$

$$\text{subject to} \quad \sum_{a=1}^{A} v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr}) - p = S + \epsilon. \tag{2}$$

Under the same circumstances in the final round, supplier s would solve (1) subject to

$$\sum_{a=1}^{A} f_a(x_a) - p = S + \epsilon. \tag{3}$$

In either case, 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.* 1999, Parkes and Ungar 2000, Gallien and Wein 2000, although the first two studies refer to MBR as "straightforward bidding"), and asserts a middle ground with regards to the bidders' rationality. On the one hand, they are assumed to be sophisticated enough to formulate and solve (1)-(2) or (1), (3). 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 $r \leq P$ bid $(p, x_1, \ldots, x_A)$ that generates the arbitrary score $\tilde{S}$, we say that it is undistorted if this bid is the solution to (1)-(2) with the right side of (2) replaced by $\tilde{S}$. Were the bid in round $P+1$, an undistorted bid would solve (1), (3) with $\tilde{S}$ replacing the right side of (3). A supplier submitting undistorted bids may withhold *absolute* information about his cost parameters by not bidding aggressively (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* $1, \ldots, P$, *and are MBR in round* $P+1$.

Our earlier assumptions about the cost function and scoring rules imply that, for $r \leq P$, 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

$$\frac{\partial v_a(x_a, \phi_{a1r}, \ldots, \phi_{aPr})}{\partial x_a} = \frac{\partial c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial x_a} \quad \text{for} \quad a = 1, \ldots, A. \tag{4}$$

Because the first-order condition (4) is independent of the right side of (2), it follows that all undistorted bids satisfy (4). Hence, all bids by a given supplier s in a given round r have the same magnitude for their non-price attributes; these undistorted bids differ only in their price. This observation provides a simple way to empirically validate or invalidate the undistorted-bid assumption.

Analogously, the solution to (1), (3) in round $P+1$ must satisfy

$$\frac{df_a(x_a)}{dx_a} = \frac{\partial c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial x_a} \quad \text{for} \quad a = 1, \ldots, A. \tag{5}$$

If we let x_a^* denote the solution to (5), then the corresponding bid price in round $P+1$ is

$$p^* = \sum_{a=1}^{A} f_a(x_a^*) - S - \epsilon. \tag{6}$$

If $p^* \geq \sum_{a=1}^{A} c_{as}(x_a^*, \theta_{as1}, \ldots, \theta_{asP})$ then $(p^*, x_1^*, \ldots, x_A^*)$ is the MBR bid in round $P+1$; 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 (4) with $r = 1, \ldots, P$, in terms of the P unknown cost parameters. If for fixed attribute a and supplier s the scoring vectors $(\phi_{a1r}, \ldots, \phi_{aPr})$ 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, $\frac{\partial v_a(x_{asr}^*, \phi_{a1r}, \ldots, \phi_{aPr})}{\partial x_{asr}} \neq \frac{\partial v_a(x_{as\bar{r}}^*, \phi_{a1\bar{r}}, \ldots, \phi_{as\bar{r}})}{\partial x_{as\bar{r}}}$, for $\bar{r} = 1, \ldots, r-1$; in practice, this could be enforced by simply perturbing (perhaps several times if needed) scoring vectors 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 $P + 1$. With the true cost functions in hand, the auctioneer can determine his optimal scoring rule for round $P+1$. Let us suppose for now that the auctioneer chooses the generic scoring rule f_a for attribute a; to avoid introducing more terminology, we refer to $f_1, \ldots, f_A$, collectively and individually, as scoring rules, although the actual scoring rule for the auction in the final round is $\sum_{a=1}^{A} f_a(x_a) - p$. By the MBR assumption, supplier s will submit bids that solve (1), (3). In §2.1 we imposed sufficient, but not necessary, conditions on $f_1, \ldots, f_A$ for problem (1), (3) to possess a unique finite positive solution; if $f_1, \ldots, f_A$ satisfies these conditions, we refer to $f_1, \ldots, f_A$ as *feasible*.

Notice that supplier s will drop out of round $P + 1$'s open-ascending competition no later than when his profit $p - \sum_{a=1}^{A} c_{as}(x_a^*, \theta_{as1}, \ldots, \theta_{asP})$ equals zero (where x_a^* is the solution to (5)), which occurs at the *maximum drop-out score*

$$S_s = \sum_{a=1}^{A} f_a(x_a^*) - \sum_{a=1}^{A} c_{as}(x_a^*, \theta_{as1}, \ldots, \theta_{asP}). \tag{7}$$

Our analysis below, culminating in (13), depends only on the top two bidders, and we index the suppliers so that their maximum drop-out scores in round $P+1$ satisfy $S_1 \geq S_2 \geq \ldots \geq S_S$; note that this ranking depends upon our choice of generic scoring rule. To guarantee that these two suppliers can actually bid, we impose the constraint

$$S_2 > \epsilon \tag{8}$$

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 $f_1, \ldots, f_A$ to satisfy

$$S_1 > S_2 + \epsilon. \tag{9}$$

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

$$\max_{p, x_1, \ldots, x_A} \quad p - \sum_{a=1}^{A} c_{a1}(x_a, \theta_{a11}, \ldots, \theta_{a1P}) \tag{10}$$

$$\text{subject to} \quad \sum_{a=1}^{A} f_a(x_a) - p = S_2. \tag{11}$$

By (6) and (7), this solution is x^*_{a1} (we now include the supplier subscript in the bids), which solves (5) for $s = 1$, and

$$\begin{aligned} p^*_1 &= \sum_{a=1}^{A} f_a(x^*_{a1}) - S_2 \\ &= \sum_{a=1}^{A} f_a(x^*_{a1}) - \sum_{a=1}^{A} f_a(x^*_{a2}) + \sum_{a=1}^{A} c_{a2}(x^*_{a2}, \theta_{a21}, \ldots, \theta_{a2P}). \end{aligned} \tag{12}$$

Recall that the auctioneer's true utility function is $\sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP}) - p$. Using equation (12), the auctioneer's *optimal* scoring rule in round $P+1$ is the feasible scoring rule that solves

$$\max_{f_a} \quad \sum_{a=1}^{A} v_a(x_{a1}^*, \psi_{a1}, \ldots, \psi_{aP}) - \sum_{a=1}^{A} f_a(x_{a1}^*) + \sum_{a=1}^{A} f_a(x_{a2}^*) - \sum_{a=1}^{A} c_{a2}(x_{a2}^*, \theta_{a21}, \ldots, \theta_{a2P}), \tag{13}$$

subject to constraints (7)-(9). It is possible that the auctioneer's true valuation function violates equation (9); 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 (13) subject to (7)-(9) 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 (9) – we consider v to be optimal.

We re-emphasize that although equations (7)-(9), (13) 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 (7)-(9), (13) for all $2\binom{S}{2}$ ordered pairs of suppliers, and the ordered pair that generates the highest utility in (13) 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 $S_1 \geq S_2 \geq \ldots \geq S_S$.

To structure the presentation, we call a bid $(p, x_1, \ldots, x_A)$ *enforceable* if there exists a feasible scoring rule $f_1, \ldots, f_A$ satisfying (8)-(9) 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 omit the parameters from the suppliers' cost functions and the auctioneer's true valuation function, and 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.

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 $c_j(\vec{z}) < T_i(\vec{z}) + \pi$).*

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 non-price 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 f 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 1 and 2 in §2.6 for a two-dimensional (e.g., price, quality) problem. Near $q = 0.86$, the prices that are enforceable in Figure 2 are lower than those of Figure 1; the tangent lines (one-dimensional hyperplanes) to c_1 near $q = 0.86$ are much closer to c_2 in Figure 2, 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 (7)-(9), (13). 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.*

Note that M_s is supplier s's maximum drop-out 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. Since 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 $M_i - \epsilon$; since 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 (7)-(9), (13) from above. The function v is considered an optimal scoring rule and i and s are both taken to be optimal suppliers.

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 its Corollary 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

$$M_s = \max_{\vec{x}} \left\{ \sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP}) - \sum_{a=1}^{A} c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP}) \right\}. \tag{14}$$

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

$$\max_{\vec{x},\pi} \quad \sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP}) - \sum_{a=1}^{A} c_{ai}(x_a, \theta_{ai1}, \ldots, \theta_{aiP}) - \pi \tag{15}$$

$$\text{subject to} \quad \sum_{a=1}^{A} c_{as}(x_a, \theta_{as1}, \ldots, \theta_{asP}) > \sum_{a=1}^{A} c_{ai}(x_a, \theta_{ai1}, \ldots, \theta_{aiP}) + \epsilon, \; s \neq i, \tag{16}$$

$$\pi \geq \epsilon, \tag{17}$$

$$\pi \geq \min_{s \neq i} \min_{\vec{z}} \left\{ \sum_{a=1}^{A} c_{as}(z_a, \theta_{as1}, \ldots, \theta_{asP}) - T_i(\vec{z}) \right\}, \tag{18}$$

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 (15)-(18) by finding a closed-form solution (equation (48) in §A3) to the innermost minimization in (18). We will call any supplier

j who achieves the minimization in (18) in the optimal solution (and thereby enables the optimal solution to be enforced) a "best competitor" to supplier i.

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 (15)-(17), (48). There are three cases to consider. If equation (17) is tight at the optimal solution, then the scoring rule is the function f constructed in Claim 5 of the "⇒" direction proof of Proposition 1 in §A1.2. For fixed dimension a, this optimal scoring function f_a has the form

$$f_a(x_a) = \begin{cases} \omega_{a1} x_a^{\omega_{a2}} & \text{if} \quad x_a \leq z_{a1}; \\ \omega_{a3}(z_{a1} - x_a)^{\omega_{a4}} + \omega_{a5} x_a + \omega_{a6} & \text{if} \quad z_{a1} < x_a \leq z_{a2}; \\ \omega_{a7} x_a^{\omega_{a8}} & \text{if} \quad x_a > z_{a2}, \end{cases} \tag{19}$$

where three of the ten 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

$$\left(\sum_{a=1}^{A} c_{ai}(x_a^*, \theta_{ai1}, \ldots, \theta_{aiP}) + \pi^*, \ \vec{x}^* \right) \tag{20}$$

in a two-supplier auction between i and j. If equation (17) 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 f 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 (8). If v satisfies (8) then the optimal scoring rule is $\lambda^* f + (1 - \lambda^*)v$, where λ^* is given in (46) at the end of §A1.2. This scoring rule, which has 8+P parameters (seven from f, P from v, plus the parameter λ^*) enforces (20). Otherwise, if (17) is non-binding but v violates (8), then the optimal scoring rule is $\lambda^* f + (1 - \lambda^*)g$, where f is given in (19), λ^* is given in (46), 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 to enforce (20): the function g has a form similar to (19), but with the middle case repeated, for a total of ten non-redundant parameters.

2.6. A Numerical Example. We illustrate our mechanism with a simple numerical example that has $S = 2$ suppliers, $A = 1$ non-price 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., $\theta_{11} = 0$, $\theta_{12} = 1$, $\theta_{13} = 4$, $\theta_{11} = 3$, $\theta_{12} = 0$, and $\theta_{13} = 1$. The true value function is $\psi_1 q^{\psi_2} + \psi_3 q$, where $\psi_1 = 5$, $\psi_2 = 0.9$, and $\psi_3 = 0$; see Figure 1 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 (i.e., $\phi_{11} = 2$, $\phi_{21} = 1/2$, $\phi_{12} = 3.5$, $\phi_{22} = 0.7$, $\phi_{13} = 4$, $\phi_{23} = 0.8$, and $\phi_{3r} = 0$ for $r = 1, 2, 3$.). These scoring rules were chosen arbitrarily, but do satisfy the requirements of §2.1 and the requirement of §2.4 related to non-redundant 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 (4), which for round r is

$$\theta_{s1} + 3\theta_{s2}(q^*_{sr})^2 + 15\theta_{s3}(q^*_{sr})^{14} = \phi_{2r}\phi_{1r}(q^*_{sr})^{\phi_{2r}-1} + \phi_{3r}. \tag{21}$$

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, the equations (21) for rounds 1-3 yield a linear system with three unknowns; for supplier 2, the system is

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

Solving (22), 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 $\lambda f + (1-\lambda)v$ for the final round of bidding (v satisfies (8)). 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 (15)-(17), (48) to find q^*, π^*, the optimal quality level and profit at which supplier 1 wins the auction. Since $(c_2^{'})^{-1}(x) = \left(\frac{x-3}{15}\right)^{1/14}$, and $c_2^{'}(0) = 3$ and $c_1^{'}(0.7593) = 3$, we solve

$$
\begin{aligned}
\max_{q \geq 0, \pi} \quad & \psi_1 q^{\psi_2} + \psi_3 q - \theta_{11} q - \theta_{12} q^3 - \theta_{13} q^{15} - \pi \\
\text{subject to} \quad & \theta_{11} q + \theta_{12} q^3 + \theta_{13} q^{15} \; > \; \theta_{21} q + \theta_{22} q^3 + \theta_{23} q^{15} + \epsilon, \\
& \pi \geq \epsilon, \\
& \pi \geq \theta_{21}\hat{q} - \theta_{22}\hat{q}^3 - \theta_{23}\hat{q}^{15} - [\theta_{11} q + \theta_{12} q^3 + \theta_{13} q^{15}] \\
& \quad + \hat{q}[\theta_{11} + 3\theta_{12} q^2 + 15\theta_{13} q^{14}], \\
& \hat{q} = \begin{cases} 0 - q & \text{if} \;\; q \leq 0.7593 \\ \left(\frac{\theta_{11} + \theta_{12}\cdot 3q^2 + \theta_{13}\cdot 15 q^{14}}{15} - 3\right)^{1/14} - q & \text{if} \;\; q > 0.7593 \end{cases}.
\end{aligned}
$$

The minimum enforceable prices (per Corollary 1) are shown in Figure 1. 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 1. (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's) quality level is 0.01, the auctioneer's utility is $\psi_1(q^*)^{\psi_2} + \psi_3 q^* - \theta_{11} q^* - \theta_{12}(q^*)^3 - \theta_{a3}(q^*)^{15} - \pi^* - \epsilon = 2.3909$, the winning supplier's profit is $\pi^* + \epsilon = 0.6327$, and the total surplus is 3.0236.

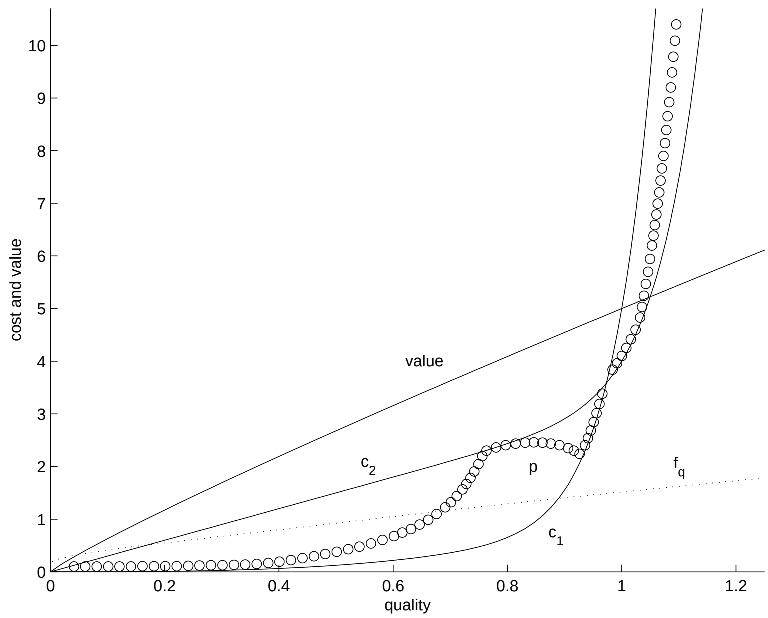

Figure 1: True value, supplier 1's cost c_1, supplier 2's cost c_2, minimum enforceable price p, and the optimal scoring rule $(\cdots)$ versus quality.

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 2. Running the optimization (15)-(17), (48) 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 2), the payoff to the auctioneer is 3.2627, which is roughly 35% greater than in the previous example. Referring to Figure 2, 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 1) or overstate (Figure 2) the true valuation of quality. In summary, our scoring rule can operate in a fundamentally different manner than Che's rule. Che's optimal scoring rule

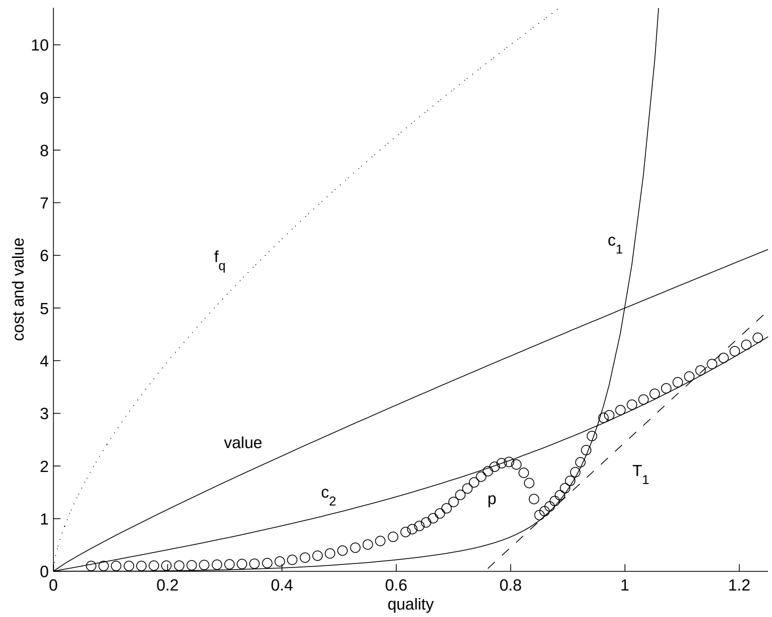

Figure 2: True value, supplier 1's cost c_1, supplier 2's cost c_2, minimum enforceable price p, and optimal scoring rule $(\cdots)$ versus quality for the high-competition example. T_1, the line tangent to supplier 1's cost curve at $q = 0.8640$, intersects the cost curve of supplier 2.

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 assume that the suppliers submit their MBR bids. This scenario adheres to equations (7)-(9), (13), but with $\psi_1 q^{\psi_2} + \psi_3 q$ in place of f; we consider the first, "low-competition" example. Supplier s's maximum drop out score is M_s, so supplier 1 wins the auction bidding quality $q^* = 0.8008$ (the solution to the right hand side of (14) 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: (i) the suppliers do not distort their bids in rounds $1, \ldots, P$ and they bid their MBR in round $P + 1$; and (ii) 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: the determination of the optimal enforceable bid may require solving a rather difficult mathematical program; the resulting scoring rule may be too complex ($8 + P$ parameters per attribute if v satisfies (9), 18 otherwise) for practical implementation; and the scoring rule may force the losing supplier to submit bids with negligible non-price attribute levels. To deal with the first problem, Step 2 of our method can be replaced by restricting (15)-(17), (48) to $\vec{x} = \vec{\hat{x}}$, where $\vec{\hat{x}}$ maximizes the right side of (14) 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 $\vec{\tilde{x}}$. 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 (19)) can perhaps be finessed in practice by providing the scoring rule in graphical form (one graph per non-price 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. For this purpose, a natural parameterization is that of the true value function; in this case, 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

$$\max_{\phi_{ap}} \quad \sum_{a=1}^{A} v_a(x^*_{ai}, \psi_{a1}, \dots, \psi_{aP}) - \sum_{a=1}^{A} v_a(x^*_{ai}, \phi_{a1}, \dots, \phi_{aP}) + \sum_{a=1}^{A} v_a(x^*_{aj}, \phi_{a1}, \dots, \phi_{aP}) - \sum_{a=1}^{A} c_{aj}(x^*_{aj}, \theta_{aj1}, \dots, \theta_{ajP}) \tag{23}$$

$$\text{subject to} \quad S_s = \sum_{a=1}^{A} v_a(x^*_{a}, \phi_{a1}, \dots, \phi_{aP}) - \sum_{a=1}^{A} c_{as}(x^*_{a}, \theta_{as1}, \dots, \theta_{asP}), \quad s = i, j, \tag{24}$$

(8)-(9), and the scoring vector constraints at the end of §2.1 (the scoring vector constraint mentioned later in §2.4 is superfluous here). Since 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 (8)-(9), (23)-(24) 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 $\vec{\hat{x}}$, where it is precisely ϵ units higher. While ignoring these constraints may increase the auctioneer's utility over 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 non-concave 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 non-price 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 non-price attributes.

3.2. Cost Estimation. Each time a supplier submits a bid, he generates a new version of equation (4) that the auctioneer can use to estimate the suppliers' cost parameters. As noted earlier, if these bids are undistorted, then the vector of non-price 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 (4) 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 $(\theta_{as1}, \ldots, \theta_{asP})$ 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 non-price 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 non-strategic 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) cost parameter estimates for the two highest-bidding suppliers in the previous round. 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 (4), we also assume that suppliers submitted their MBR bids in round $r-1$. In particular, this assumption implies that the second-highest bidder's final score in round $r-1$ is his drop-out score. If we denote this bidder as supplier 2 and his bid as

$(p_2^*, x_{12}^*, \ldots, x_{A2}^*)$, then this assumption generates the additional equation

$$p_2^* = \sum_{a=1}^{A} c_{a2}(x_{a2}^*, \theta_{a21}, \ldots, \theta_{a2P}). \tag{25}$$

If all bids are indeed undistorted, and supplier 2 submitted MBR bids in round $r-1$, then equation (25) can be combined with the $r-1$ first-order equations from the earlier rounds to solve uniquely for the second-highest bidder'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 (25) and the additional constraints $\sum_{b=1}^{B_r} w_{rb} = 1$, $w_{rb} \geq 0$.

We only have censored information, namely $p_1^* \geq \sum_{a=1}^{A} c_{a1}(x_{a1}^*, \theta_{a11}, \ldots, \theta_{a1P})$, for the highest bidder in round $r-1$. That is, we do not know this bidder's drop-out score. Consequently, in the third step of our procedure we assume that the current first-place bidder has a drop-out score that is a fixed percentage (e.g., 5% or 10%) higher than the second-highest bidder's observed drop-out score; this fixed percentage should include the estimated gap between the top two bidders and the perceived amount by which the second-place bidder is "holding back" (i.e., his true drop-out score minus his observed drop-out score). This estimated drop-out score provides the additional equation (see (7)) to estimate the current first-place 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 FCC auctions) to prevent bidders from delaying their bid submissions until the very 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 $\sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP}) + e - p$.

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 shifted vertically by $-e_s$ units; this shift is allocated to the costs over individual attributes by taking c_{as} to be shifted $-\lambda_{as}e_s$ units vertically, 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 multi-attribute procurement auction from the auctioneer's perspective is the lack of knowledge about the suppliers' cost functions for endogenous non-price 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 taken place. 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 modeled probabilistically. 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., page 6 of Laseter 1998), 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 drop-out score (see equation (7)) 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 non-price 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 1.5 to 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 non-smooth cost and utility functions.

ACKNOWLEDGMENT

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

References

[1] Ahuja, R. K., J. B. Orlin. 2001. Inverse optimization. *Operations Research* **49**, 771-783.

[2] Bikhchandani, S. 1999. Auctions of heterogeneous objects. *Games and Economic Be-*

havior **26**, 193-220.

[3] Branco, F. 1997. The design of multidimensional auctions. *RAND Journal of Economics* **28**, 63-81.

[4] Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND Journal of Economics* **24**, 668-680.

[5] Cramton, P. 1998. Ascending auctions. *European Economic Review* **42**, 745-756.

[6] Demange, G., D. Gale, M. Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy* **94**, 863-872.

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

[8] Gallien, J., L. M. Wein. 2000. Design and analysis of a smart market for industrial procurement. Sloan School of Management, MIT, Cambridge, MA.

[9] Gittens, J. C. 1989. *Multi-armed Bandit Allocation Indices*. John Wiley and Sons, New York.

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

[11] Kelly, F., R. Steinberg. 2000. A combinatorial auction with multiple winners for universal service. *Management Science* **46**, 586-596.

[12] Laseter, T. M. 1998. *Balanced Sourcing*. Jossey-Bass, Inc., San Francisco, CA.

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

[14] Milgrom, P. 2000b. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* **108**, 245-272.

[15] Parkes, D. C., L. H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *Proceedings 17th National Conference on Artificial Intelligence (AAAI-00)*, 74-81.

[16] Tarantola, A. 1987. *Inverse Problem Theory*. Elsevier, Amsterdam.

[17] Stanley, E. D., D. P. Honig, L. Gainen. 1954. Linear programming in bid evaluation. *Naval Research Logistics Quarterly* **1**, 48-54.

[18] Wellman, M. P., W. E. Walsh, P. R. Wurman, J. K. MacKie-Mason. 1999. Auction protocols for decentralized scheduling. To appear in *Games and Economic Behavior*.

[19] Wise, R., D. Morrison. 2000. Beyond the exchange: The future of B2B. *Harvard Business Review*, November-December, 86-96.

ONLINE APPENDIX

A1. Proof of Proposition 1. We begin by proving the "$\Leftarrow$" direction, which is easier.

A1.1. Proposition 1 "$\Leftarrow$" Direction. Suppose $f_1, \ldots, f_A$ enforces $(c_i(\vec{x}) + \pi, \vec{x})$. For convenience, let $f(\vec{x}) = \sum_{a=1}^{A} f_a(x_a)$. Clearly, the assumption that supplier i is the winner of the auction implies that $c_j(\vec{x}) > c_i(\vec{x}) + \pi$. We now show that $T_i + \pi$ intersects c_j.

Let $\vec{z}$ denote supplier j's bid induced by f. The assumption that supplier i wins the auction with profit π implies that

$$f(\vec{x}) - c_i(\vec{x}) - (f(\vec{z}) - c_j(\vec{z})) = \pi,$$

$$\Rightarrow f(\vec{x}) - f(\vec{z}) = \pi + c_i(\vec{x}) - c_j(\vec{z}),$$

$$\Rightarrow \nabla f(\vec{x})\prime(\vec{x} - \vec{z}) < \pi + c_i(\vec{x}) - c_j(\vec{z}) \qquad \text{since } f \text{ is concave},$$

$$\Rightarrow \nabla c_i(\vec{x})\prime(\vec{x} - \vec{z}) < \pi + c_i(\vec{x}) - c_j(\vec{z}) \qquad \text{since supplier } i \text{ bids } \vec{x} \text{ iff } \nabla c_i(\vec{x}) = \nabla f(\vec{x}),$$

$$\Rightarrow c_j(\vec{z}) < \pi + c_i(\vec{x}) + \nabla c_i(\vec{x})\prime(\vec{z} - \vec{x}) = T_i(\vec{z}) + \pi.$$

Since c_j is convex, increasing and lies below $T_i + \pi$ at $\vec{z}$, c_j must eventually cross the hyperplane $T_i + \pi$.

We next prove the other direction of Proposition 1, which is more difficult.

A1.2. Proof of Proposition 1 "$\Rightarrow$" Direction. The proof proceeds in two stages. First, we prove Proposition 1 for the two-supplier case, where $s = 1, 2$. We index the assumptions as (A1): $c_2(\vec{x}) > c_1(\vec{x}) + \pi$; and (A2): the hyperplane $T_1 + \pi$ intersects c_2. Given (A1) and (A2), we need to find a feasible scoring rule $f_1, \ldots, f_A$ such that $S_2 > \epsilon$ and supplier 1 wins the auction with bid $(c_1(\vec{x}) + \pi, \vec{x})$ (which also implies $S_1 > \epsilon$). Such an f is found via Claims 1-5.

In the second stage, the two-supplier results are extended to S suppliers. Claim 7 (proof of the "$\Rightarrow$" direction of Proposition 1) shows that $(c_i(\vec{x}) + \pi, \vec{x})$ can be enforced by taking a convex combination of the scoring rule that enforces $(c_i(\vec{x})+\pi, \vec{x})$ in the two-supplier case (with i and j in the roles of 1 and 2), and a second scoring rule constructed in Claim 6.

Stage 1. Claim 1 contains the main construction result for the two-dimensional (e.g., price and a non-price attribute), two supplier case; Claim 1 will be our main tool, as the construction of f for the multi-attribute case uses an attribute-by-attribute construction procedure. This procedure is provided in Claim 5, which builds upon the groundwork laid by Claims 1-4.

Claim 1. *For $i = 1,\ 2$, let $g_i : \mathbb{R}^+ \to \mathbb{R}$ be convex, increasing smooth functions. Let l_1 be the line tangent to g_1 at $z_1 > 0$, and let $\alpha \geq 0$. If $g_2(z_1) > g_1(z_1) + \alpha$ and $l_1 + \alpha$ intersects g_2, then there exists $z_2 > 0$ and a concave, increasing smooth function $\tilde{f}$ such that $\tilde{f}'(z) \to 0$ as $z \to \infty$, for $i = 1,\ 2$ $\tilde{f}'(z) > g_i(z)$ as $z \to 0^+$, and*

$$\tilde{f}(z_1) = g_1(z_1) + \alpha + h, \quad \tilde{f}(z_2) = g_2(z_2) + h, \tag{26}$$

$$\tilde{f}'(z_1) = g'(z_1), \quad \textit{and} \quad \tilde{f}'(z_2) = g_2'(z_2), \tag{27}$$

where h can be chosen arbitrarily large.

Proof. We first find a z_2 such that

$$g_2(z_2) + g_2'(z_2)[z_1 - z_2] > g_1(z_1) + \alpha, \quad \text{and} \tag{28}$$

$$g_2(z_2) < l_1(z_2) + \alpha. \tag{29}$$

Since g_2 is smooth, convex, and increasing and lies above $l_1 + \alpha$ at z_1, g_2 intersects $l_1 + \alpha$ either to the right or left of z_1 (but not both).

(c1): g_2 intersects $l_1 + \alpha$ to the left of z_1. Let $\bar{z}$ be the closest point to z_1 at which g_2 intersects $l_1 + \alpha$. At $\bar{z}$, g_2 must cross $l_1 + \alpha$ from below; hence, $g_2'(\bar{z}) > g_1'(z_1) =$ slope of $l_1 + \alpha$. Since $(\bar{z}, g_2(\bar{z}))$ and $(z_1, g_1(z_1) + \alpha)$ are points on the line $l_1 + \alpha$, we can write the point slope formula for $l_1+\alpha$ at the point $(\bar{z}, g_2(\bar{z}))$ and evaluate at z_1, yielding the equation

$$g_2(\bar{z}) + g_1'(z_1)[z_1 - \bar{z}] = g_1(z_1) + \alpha.$$

Replacing $g_1'(z_1)$ by the larger value $g_2'(\bar{z})$ implies

$$g_2(\bar{z}) + g_2'(\bar{z})[z_1 - \bar{z}] > g_1(z_1) + \alpha, \tag{30}$$

since $z_1 > \bar{z}$. Noting that the left side of (30) can be viewed as a continuous function of $\bar{z}$, we can find a $\eta_1 > 0$ such that (30) holds with z in place of $\bar{z}$ as long as $|\bar{z} - z| < \eta_1$. Since g_2 is itself continuous and approaches $l_1 + \alpha$ from below at $\bar{z}$, there exists a $\eta_2 > 0$ such that $g_2(z) < l_1(z) + \alpha$ provided that $z < \bar{z}$ and $|\bar{z} - z| < \eta_2$. An appropriate value for z_2 is then found by taking

$$z_2 = \bar{z} - \eta, \tag{31}$$

provided $0 < \eta < \min\{\eta_1, \eta_2\}$.

(c2): g_2 intersects $l_1 + \alpha$ to the right of z_1. Let $\bar{z}$ be the closest point to z_1 at which g_2 intersect $l_1 + \alpha$. At $\bar{z}$ g_2 must cross $l_1 + \alpha$ from above; the remaining arguments to find z_2 are straightforward analogues to those of (c1).

Now that we have shown that a z_2 satisfying (28)-(29) exists, we now construct $\tilde{f}$, beginning with the case $z_2 < z_1$; the complement case is essentially the same, and is omitted for brevity.

The proof for the case $z_2 < z_1$ synthesizes $\tilde{f}$ from three concave, increasing functions $\tilde{f}_1$, $\tilde{f}_2$, and $\tilde{f}_3$ (note that the subscripts on f in this proof do not correspond to attributes), defined over respective intervals $[0, z_2]$, $[z_2, z_1]$, and $[z_1, \infty)$. For each $\tilde{f}_i$ we enforce conditions at z_1 and z_2 toward satisfying (26)-(27):

$$\tilde{f}_i(z_2) = g_2(z_2) + h, \quad \text{and} \quad \tilde{f}_i'(z_2) = g_2'(z_2), \quad \text{for} \quad i = 1, 2, \quad \text{and} \tag{32}$$

$$\tilde{f}_i(z_1) = g_1(z_1) + \alpha + h, \quad \text{and} \quad \tilde{f}_i'(z_1) = g_1'(z_1), \quad \text{for} \quad i = 2, 3, \tag{33}$$

as well as choose the functions' forms to ensure that the other conditions of the Claim's statement are satisfied. We begin by constructing $\tilde{f}_1$.

We set $\tilde{f}_1(z) = \gamma_{11} z^{\gamma_{12}}$. Solving (32) with $i = 1$ for $\gamma_{11},\ \gamma_{12}$ yields

$$\gamma_{11} = [g_2(z_2) + h] z_2^{-\gamma_{12}}, \quad \text{and} \quad \gamma_{12} = \frac{g_2'(z_2) z_2}{g_2(z_2) + h}.$$

Over $[0, z_2]$ $\tilde{f}_1$ is increasing, since $\gamma_{11} > 0$. Choosing $h > g_2'(z_2)z_2 - g_2(z_2)$ ensures that $\gamma_{12} < 1$, and thereby the concavity of $\tilde{f}_1$ and $\tilde{f}_1'(z) \to \infty$ as $z \to 0^+$.

We find $\tilde{f}_2$ by first finding what $\tilde{f}_2$ would be in a new, shifted and rotated coordinate space. We then take the resulting function, $\bar{f}_2$, and apply a series of reflections, translations and rotations to produce the desired $\tilde{f}_2$. The new coordinate space we find convenient is that in which we view $(z_1, g_1(z_1) + \alpha + h)$ as the origin. We take $l_1 + \alpha + h$ as the horizontal axis, the line perpendicular to $l_1 + \alpha + h$ at $(z_1, g_1(z_1) + \alpha + h)$ as the vertical axis, and as the positive quadrant everything below the former and to the left of the latter. The appropriate equations for $\bar{f}_2$ are

$$\bar{f}_2(z_1 - z_2) = l_1(z_1) + \alpha - g_2(z_2), \quad \text{and}$$

$$\begin{aligned} \bar{f}_2'(z_1 - z_2) &= g_2'(z_2) - \text{ slope of } (l_1 + \alpha + h), \\ &= g_2'(z_2) - g_1'(z_1). \end{aligned}$$

If we let $\bar{f}_2(\bar{z}) = \bar{\gamma}_{21} \bar{z}^{\bar{\gamma}_{22}}$ and solve the above two equations for $\bar{\gamma}_{21},\ \bar{\gamma}_{22}$, we get

$$\bar{\gamma}_{21} = [l_1(z_2) + \alpha - g_2(z_2)][z_1 - z_2]^{-\gamma_{22}}, \quad \text{and}$$

$$\bar{\gamma}_{22} = [g_2'(z_2) - g_1'(z_1)][l_1(z_2) + \alpha - g_2(z_2)]^{-1}[z_1 - z_2],$$

where $l_1(z_2) = g_1(z_1) + g_1'(z_1)[z_2 - z_1]$. The step-by-step transformations to yield $\tilde{f}_2$ from $\bar{f}_2$ (to translate the new coordinate space into traditional Euclidean coordinate space) are: 1. Reflect about the origin: $-\bar{f}_2(-\bar{z})$; 2. Shift z_1 units to the right and $g_1(z_1) + \alpha + h$ units up: $-\bar{f}_2(-(\bar{z}-z_1))+g_1(z_1)+\alpha+h$; 3. Rotate counter-clockwise by $\sin^{-1} g_1'(z_1)$ degrees, pivoting about the point $(z_1, g_1(z_1)+\alpha+h)$:

$$-\bar{f}_2(-(\bar{z}-z_1))+g_1(z_1)+\alpha+h+g_1'(z_1)[\bar{z}-z_1]. \tag{34}$$

Equation (34) is a function in the traditional Euclidean space; if we replace $\bar{z}$ by z and set the result equal to $\tilde{f}_2$ we have

$$\tilde{f}_2(z) = -\bar{\gamma}_{21}(-z+z_1)^{\bar{\gamma}_{22}} + g_1(z_1)+\alpha+h+g_1'(z_1)[z-z_1];$$

it is straightforward to check that the equations (32)-(33) with $i=2$ are satisfied. Since $\tilde{f}_2'(z) = \bar{\gamma}_{22}\bar{\gamma}_{21}(-z+z_1)^{\gamma_{22}-1} + g_1'(z_1)$, $\tilde{f}_2''(z) = -\bar{\gamma}_{22}(\bar{\gamma}_{22}-1)\bar{\gamma}_{21}(-z+z_1)^{\bar{\gamma}_{22}-2}$, for verifying that in $[z_2, z_1]$ $\tilde{f}_2$ is concave, increasing, it suffices to show that $\bar{\gamma}_{21} > 0$ and $\bar{\gamma}_{22} > 1$. The first inequality is true by (29) and our assumption that $z_2 < z_1$. A series of equivalences prove that the second inequality holds:

$$\begin{aligned}
\bar{\gamma}_{22} = [g_2'(z_2)-g_1'(z_1)][l_1(z_2)+\alpha-g_2(z_2)]^{-1}[z_1-z_2] \quad &> \quad 1, \\
\iff \quad [g_2'(z_2)-g_1'(z_1)][z_1-z_2] \quad &> \quad [l_1(z_2)+\alpha-g_2(z_2)] \text{ by (29)}, \\
\iff \quad g_2'(z_2)[z_1-z_2]-g_1'(z_1)[z_1-z_2] \quad &> \quad g_1(z_1)+g_1'(z_1)[z_2-z_1]+\alpha-g_2(z_2), \\
\iff \quad g_2(z_2)+g_2'(z_2)[z_1-z_2] \quad &> \quad g_1(z_1)+\alpha,
\end{aligned}$$

which is equation (28).

If we set $\tilde{f}_3(z) = \gamma_{31} z^{\gamma_{32}}$ and solve (33), we get

$$\gamma_{31} = [g_1(z_1)+\alpha+h]z_1^{-\gamma_{32}}, \quad \text{and} \quad \gamma_{32} = \frac{g_1'(z_1)z_1}{g_1(z_1)+\alpha+h}.$$

As for $\tilde{f}_1$, $\gamma_{31} > 0$ implies that $\tilde{f}_3$ increases over $[z_2, \infty)$. Concavity and $\tilde{f}_3'(z) \to 0$ as $z \to \infty$ is ensured if $\gamma_{32} < 1$, or equivalently, if $h > g_1'(z_1)z_1 - g_1(z_1) - \alpha$.

To complete the overall construction of $\tilde{f}$ we need only choose h such that $h > \max\{g_2'(z_2)z_2 - g_2(z_2), g_1'(z_1)z_1 - g_1(z_1) - \alpha\}$. Finally, note that the smoothness of the individual $\tilde{f}_i$'s and the requirements (32)-(33) ensure that $\tilde{f}$ is smooth.

Taking a step back, notice that in defining $\tilde{f}$ we used only z_1, z_2, $g_1(z_1)$, $g_2(z_2)$, $g_1'(z_1)$, $g_2'(z_2)$, and h – i.e., $\tilde{f}$ is a function of seven parameters. □

Claim 2. *Consider a two-supplier auction. Under assumptions (A1) and (A2), there exists a $\tilde{\pi} > \pi > \epsilon$ and a scoring rule $\tilde{f}$ that enforces supplier 1 winning at $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$.*

Proof. Let $d_a = c_{a2}(x_a) - c_{a1}(x_a),\ a = 1, \ldots, A$. We construct $\tilde{f}$ one dimension at a time. Assume without loss of generality that the attributes are ordered such that for $a = 1, \ldots, \hat{A}$, $c'_{a1}(x_a) \neq c'_{a2}(x_a)$. Notice that $\hat{A} \geq 1$; otherwise, $T_1 + \pi$ intersecting c_2 implies that

$$\begin{aligned} \pi &> \min_{\vec{v}}\{c_2(\vec{v}) - T_1(\vec{v})\}, \\ &= c_2(\vec{x}) - c_1(\vec{x}) \qquad \text{if } \hat{A} = 0, \end{aligned}$$

which contradicts assumption (A1).

For $a = 1, \ldots, \hat{A}$, for fixed a, consider the curve $\bar{c}_{a1} \triangleq c_{a1} + d_a - \delta_a$. Let $\bar{l}_{a1}$ be the line tangent to $\bar{c}_{a1}$ at x_a. For $\delta_a > 0$ small enough, $\bar{l}_{a1}$ intersects c_{a2}, and we can apply Claim 1 (with $g_1 = \bar{c}_{a1}$, $g_2 = c_{a2}$, and $\alpha = 0$) to find a feasible scoring rule $\tilde{f}_a$ such that, for a $z_a \in \mathbb{R}^+$,

$$\tilde{f}_a(x_a) = \bar{c}_{a1}(x_a) + h_a = c_{a1}(x_a) + d_a - \delta_a + h_a, \qquad \tilde{f}_a(z_a) = c_{a2}(z_a) + h_a,$$

$$\tilde{f}'_a(x_a) = \bar{c}'_{a1}(x_a) = c'_{a1}(x_a), \qquad \text{and} \qquad \tilde{f}'_a(z_a) = c'_{a2}(z_a),$$

where h_a can be made arbitrarily large. Now, if we take the arbitrarily large h_a from Claim 1 to be at least greater than $\max\{-d_a + \delta_a + \epsilon, \epsilon\}$ (to ensure that $\tilde{f}_a(x_a) > c_{a1}(x_a) + \epsilon$, and $\tilde{f}_a(z_a) > c_{a2}(z_a) + \epsilon$), $\tilde{f}_a$ yields a difference in dimension a maximum dropout scores of

$$\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(z_a) - c_{a2}(z_a)) = d_a - \delta_a.$$

For $a = \hat{A}+1, \ldots, A$, a fixed, if in auction a we announce a scoring rule $\tilde{f}_a(z) = \mu_{a1} z^{\mu_{a2}}$ such that

$$\tilde{f}_a(x_a) = \max\{c_{a1}(x_a), c_{a2}(x_a)\} + h_a, \qquad \text{and} \qquad \tilde{f}'_a(x_a) = c'_{a1}(x_a), \tag{35}$$

(h_a some positive real number greater than ϵ), then supplier 2's bid for dimension a is $z_a = x_a$ and

$$\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(z_a) - c_{a2}(z_a)) = c_{a2}(x_a) - c_{a1}(x_a) \equiv d_a.$$

Since $\tilde{f}_a$ is defined by two parameters (has two degrees of freedom) and has two equations in (35) to satisfy, such an $\tilde{f}_a$ can be found (i.e., appropriate values of μ_{a1} and μ_{a2} can be found).

After applying the above for $a = 1, \ldots, A$, let $\delta = \min\{\delta_1, \ldots, \delta_{\hat{A}}\}$. Since we can choose the δ_a's as small as we like, for easier bookkeeping set $\delta_a = \delta$ for $a = 1, \ldots, \hat{A}$. We then define

$$\begin{aligned} \tilde{\pi} \triangleq \tilde{f}(\vec{x}) - c_1(\vec{x}) - (\tilde{f}(\vec{z}) - c_2(\vec{z})) &= \sum_{a=1}^{A} \left[\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(z_a) - c_{a2}(z_a)) \right], \\ &= \sum_{a=1}^{A} d_a - \delta \hat{A}. \end{aligned} \tag{36}$$

Assumption (A1) implies that $c_2(\vec{x}) - c_1(\vec{x}) > \pi$, which re-written is $\sum_{a=1}^{A} d_a > \pi$. Hence, if δ is chosen sufficiently small (possible since the earlier δ arguments for Case 1 assumed only that δ was sufficiently small), equation (36) is strictly greater than π, meaning that $\tilde{f}$ enforces $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$ when announced to suppliers 1 and 2, where $\tilde{\pi} > \pi$.

Notice that the $\tilde{f}$ we constructed is a feasible scoring rule: for $a = 1, \ldots, \hat{A}$ the feasibility of the scoring rule $\tilde{f}_a$ follows directly from the assumptions on Claim 1's $\tilde{f}$; and for $a = \hat{A}+1, \ldots, A$, the feasibility of scoring rule $\tilde{f}_a$ follows by simple analysis. For $\tilde{f}$ enforcing $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$, it remains only to check that $S_1, S_2 > \epsilon$, and $S_1 > S_2 + \epsilon$. The first condition holds by our choices of h_a; the second condition follows from $\pi > \epsilon$. This concludes the proof of Claim 2. □

Claim 3. *Consider a two-supplier auction. Suppose $\tilde{f}$ enforces $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$, $\tilde{\pi} > \pi > \epsilon$.*

Let $\hat{c}_{a1} = c_{a1} - w_a$*, where*

$$w_a = \sum_{\substack{i=1,\ldots,A \\ i \neq a}} \left[\tilde{f}_i(x_i) - c_{i1}(x_i) - (\tilde{f}_i(\tilde{z}_i) - c_{i2}(\tilde{z}_i)) \right], \tag{37}$$

and $\vec{\tilde{z}}$ *is supplier 2's bid induced by* $\tilde{f}$*. Let* $\hat{l}_{a1}$ *be the line tangent to* $\hat{c}_{a1}$ *at* x_a*. If we can pick an* $\alpha \geq 0$ *such that* $\hat{c}_{a1}(x_a) + \pi + \alpha < c_{a2}(x_a)$ *and* $\hat{l}_{a1} + \pi + \alpha$ *intersects* c_{a2}*, then we can find a feasible scoring rule for dimension* a*,* $\hat{f}_a$*, such that announcing* $\tilde{f}_1, \ldots, \tilde{f}_{a-1}, \hat{f}_a, \tilde{f}_{a+1}, \ldots, \tilde{f}_A$ *enforces* $(c_1(\vec{x}) + \pi + \alpha, \vec{x})$*.*

To better understand the idea behind Claim 3, notice that if we ran all dimensions except that for attribute a, then checked the difference between the bidding positions of supplier 1 and supplier 2, this difference would be w_a. (Equivalently, if we ran an $A-1$ dimension auction over attributes $1, \ldots, a-1, a+1, \ldots, A$, w_a would be the amount by which supplier 1 won ($w_a > 0$) or lost ($w_a \leq 0$) the auction.) In essence, the shift of c_{a1} to $\hat{c}_{a1}$ summarizes supplier 1's position just before dimension a is run. We now prove Claim 3.

Proof. Suppose that there exists an $\alpha \geq 0$ such that $\hat{c}_{a1}(x_a) + \pi + \alpha < c_{a2}(x_a)$ and $\hat{l}_{a1} + \pi + \alpha$ intersects c_{a2}. Then, using Claim 1 we can find a scoring rule $\hat{f}_a$ and a $z_a \in \mathbb{R}^+$ such that, by choosing the arbitrarily large h_a at least greater than $\max\{w_a + \epsilon, \epsilon\}$,

$$\begin{aligned} \hat{f}_a(x_a) = \hat{c}_{a1}(x_a) + \pi + \alpha + h_a &= c_{a1}(x_a) - w_a + \pi + \alpha + h_a, \\ &> c_{a1}(x_a) + \epsilon \qquad \text{by } \pi,\ \alpha \geq 0 \text{ and how we chose } h_a, \end{aligned}$$

$$\hat{f}_a(\hat{z}_a) = c_{a2}(\hat{z}_a) + h_a, \quad \hat{f}'_a(x_a) = \hat{c}'_{a1}(x_a) = c'_{a1}(x_a), \quad \text{and} \quad \hat{f}'_a(\hat{z}_a) = c'_{a2}(\hat{z}_a).$$

Announcing scoring rules $\tilde{f}_1, \ldots, \tilde{f}_{a-1}, \hat{f}_a, \tilde{f}_{a+1}, \ldots, \tilde{f}_A$ results in a post-auction profit for supplier 1 of

$$
\begin{aligned}
&w_a + \hat{f}_a(x_a) - c_{a1}(x_a) - (\hat{f}_a(\hat{z}_a) - c_{a2}(\hat{z}_a)) \\
&\qquad = w_a + c_{a1}(x_a) - w_a + \pi + \alpha + h_a - c_{a1}(x_a) - (c_{a2}(\hat{z}_a) + h_a - c_{a2}(\hat{z}_a)) \\
&\qquad = \pi + \alpha. \quad (38)
\end{aligned}
$$

To prove the claim, it remains only to check that $\tilde{f}_1, \ldots, \tilde{f}_{a-1}, \hat{f}_a, \tilde{f}_{a+1}, \ldots, \tilde{f}_A$ induces $S_1, S_2 > \epsilon$ and $S_1 > S_2 + \epsilon$. The former holds by how we chose h_a; the latter holds because $\tilde{\pi} > \pi > \epsilon$. □

Claim 4. *Consider a two-supplier auction. Suppose $\tilde{f}$ enforces $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$, $\tilde{\pi} > \pi > \epsilon$. Let $\hat{c}_{a1} = c_{a1} - w_a$, where w_a is as given in equation* (37)*. Then*

$$\hat{c}_{a1}(x_a) + \pi < c_{a2}(x_a).$$

Proof. First, note that since $\tilde{f}$ enforces $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$, we know that after running all dimensions supplier 1 wins the auction by $\tilde{\pi}$; in equations,

$$w_a + \tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a)) = \tilde{\pi}, \quad (39)$$

where $\vec{\tilde{z}}$ is the bid by supplier 2 induced by $\tilde{f}$. Furthermore, using the definition of $\vec{\tilde{z}}$,

$$
\begin{aligned}
&\tilde{f}_a(x_a) - c_{a2}(x_a) \le \max_y\{\tilde{f}_a(y) - c_{a2}(y)\} = \tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a), \\
&\Rightarrow \tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a)) \le c_{a2}(x_a) - c_{a1}(x_a). \quad (40)
\end{aligned}
$$

Then, combining equations (39) and (40), we have

$$\pi < \tilde{\pi} \le w_a + c_{a2}(x_a) - c_{a1}(x_a) = c_{a2}(x_a) - \hat{c}_{a1}(x_a),$$

which verifies Claim 4. □

Claim 5. *Consider a two-supplier auction. Suppose $\pi > \epsilon$ and assumptions (A1) and (A2) hold. Then there exists a scoring rule f that enforces $(c_1(\vec{x}) + \pi, \vec{x})$.*

Proof. First, use Claim 2 to construct a feasible scoring rule $\tilde{f}$ that enforces $(c_1(\vec{x}) + \tilde{\pi}, \vec{x})$, $\tilde{\pi} > \pi$. We carry out iterations on $\tilde{f}$ to construct a scoring rule f to enforce $(c_1(\vec{x}) + \pi, \vec{x})$. The general idea is to start with $\tilde{f}$, then dimension-by-dimension decrease the profit of supplier 1 (by revising $\tilde{f}_a$'s) until supplier 1 wins the auction with profit exactly equal to π, at which point we set $f \equiv \tilde{f}$.

Let $\vec{\tilde{z}}$ be supplier 2's bid induced by $\tilde{f}_1, \ldots, \tilde{f}_A$. Assume w.o.l.o.g. that the attributes are ordered such that $c'_{A1}(x_A) \neq c'_{A2}(x_A)$. Notice that we cannot have $c'_{a1}(x_a) = c'_{a2}(x_a)$ for all a, since if we did, then T_1 would be parallel to c_2's tangent hyperplane at $\vec{x}$, implying that by (A1) $T_1 + \tilde{\pi}$ cannot not intersect c_2, which contradicts (A2).

Set $a = 1$.

Procedure:

If $c'_{a1}(x_a) = c'_{a2}(x_a)$, set $a = a + 1$.

Otherwise, let $\hat{c}_{a1}$ and w_a be as defined in Claim 3, and let l_{a1} ($\hat{l}_{a1}$) be the line tangent to c_{a1} ($\hat{c}_{a1}$) at x_a. Also, let

$$D_a = \max_y \{l_{a1}(y) - c_{a2}(y)\},$$

and let

$$\hat{D}_a = D_a - w_a + \pi.$$

Now, if $\hat{D}_a > 0$, set $\alpha = 0$. The line $\hat{l}_{a1} + \pi + \alpha$ intersects c_{a2} (since $\hat{D}_a > 0$), and

$$\begin{aligned} \hat{c}_{a1}(x_a) + \pi + \alpha &= \hat{c}_{a1}(x_a) + \pi, \\ &< c_{a2}(x_a) \qquad \text{by Claim 4.} \end{aligned}$$

Hence, we can apply Claim 3, set $\tilde{f}_a = \hat{f}_a$, and by announcing $\tilde{f}$ enforce $(c_1(\vec{x}) + \pi, \vec{x})$, and we are done – terminate the procedure.

Otherwise, if $\hat{D}_a \leq 0$, we find $\eta_a > 0$ small enough that $\hat{c}_{a1}(x_a) + \pi - \hat{D}_a + \eta_a < c_{a2}(x_a)$.

This is possible since

$$\begin{aligned} c_{a2}(x_a) - (\hat{c}_{a1}(x_a) + \pi) &= c_{a2}(x_a) - (\hat{l}_{a1}(x_a) + \pi), \\ &> \min_y \left\{ c_{a2}(y) - (\hat{l}_{a1}(y) + \pi) \right\}, \\ &= -\hat{D}_a \qquad \text{since } \hat{D}_a \leq 0. \end{aligned}$$

The minimum in the second line will occur either at zero or with first order conditions. The inequality follows since x_a will satisfy neither: $c'_{a1}(x_a) \neq c'_{a2}(x_a)$ by assumption, and since f enforces $\vec{x}$, $\vec{x}$ is an interior solution to (1), (3), and so $x_a > 0$. For producing a contradiction if the procedure runs through all the indices without terminating (explained below), we take

$$\delta_a < \min \left\{ \eta_a, \quad A^{-1} \left(\sum_{i=1}^{A} D_i + \pi \right) \right\},$$

and set $\alpha = -\hat{D}_a + \delta_a$. To check that $\delta_a > 0$, note that assumption (A2) implies

$$\begin{aligned} 0 &< \max_{\vec{y}} \{ T_1(\vec{y}) + \pi - c_2(\vec{y}) \}, \\ &= \sum_{i=1}^{A} D_i + \pi. \end{aligned}$$

Since $\delta_a > 0$, the definitions of $\hat{D}_a$ and α imply that $\hat{l}_{a1} + \pi + \alpha$ intersects c_{a2}. Thus we can apply Claim 3 to construct an $\hat{f}_a$, then set $\tilde{f}_a = \hat{f}_a$, such that announcing $\tilde{f}$ enforces $(c_1(\vec{x}) + \pi - \hat{D}_a + \delta_a, \vec{x})$, which, re-written, is

$$(c_1(\vec{x}) + w_a - D_a + \delta_a, \vec{x}). \tag{41}$$

To prepare for the next step in the iterative procedure, let $\hat{z}_a$ be supplier 2's bid in dimension a induced by $\hat{f}_a$, set $\tilde{z}_a = \hat{z}_a$, and set $\tilde{\pi} = w_a - D_a + \delta_a$. Since

$$0 \geq \hat{D}_a = D_a - w_a + \pi \qquad \text{implies that} \qquad w_a - D_a + \delta_a > \pi,$$

we get that $\tilde{\pi} > \pi$.

Finally, to complete the groundwork for the proof of termination, by the definition of w_a, the fact that we have not changed $\tilde{f}_1, \ldots, \tilde{f}_{a-1}, \tilde{f}_{a+1}, \ldots, \tilde{f}_A$, and that $\tilde{f}$ now enforces (41), we have that

$$\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a)) = -D_a + \delta_a. \tag{42}$$

Setting $a = a + 1$, we can now repeat the same procedure with our updated $\tilde{f}$, $\vec{\tilde{z}}$, and $\tilde{\pi}$.

Proof of termination with the desired $\tilde{f}$: Suppose that the procedure iterates through indices $1, \ldots, A$ without terminating with an $\alpha = 0$. We show that this would produce a contradiction.

After the procedure iterates through all the indices, $\tilde{f}$ enforces $(c_{a1}(\vec{x})+w_A-D_A+\delta_A, \vec{x})$, with $w_A - D_A + \delta_A > \pi > 0$. Since no index is visited twice, each $\tilde{f}_a$ is updated exactly once, and: If $c'_{a1}(x_a) = c'_{a2}(x_a)$, then $z_a = x_a$ implies

$$\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a)) = c_{a2}(\tilde{z}_a) - c_{a1}(x_a) = -D_a, \tag{43}$$

since in this case the maximization that defines D_a will occur with first order conditions; and, if $c'_{a1}(x_a) \neq c'_{a2}(x_a)$, then equation (42) holds. Hence,

$$\begin{aligned} w_A - D_A + \delta_A &= \sum_{a=1}^{A-1} \left[\tilde{f}_a(x_a) - c_{a1}(x_a) - (\tilde{f}_a(\tilde{z}_a) - c_{a2}(\tilde{z}_a))\right] - D_A + \delta_A && \text{by (37),} \\ &\leq \sum_{a=1}^{A} -D_a + \frac{A}{A}\left(-\sum_{a=1}^{A} D_a + \pi\right) && \text{by (42)-(43) and our choice of } \delta_a, \\ &= \pi, \end{aligned}$$

which contradicts $w_A - D_A + \delta_A > \pi$. We conclude that the procedure must terminate with $\alpha = 0$ for some index, and an $\tilde{f}$ that enforces $(c_1(\vec{x}) + \pi, \vec{x})$. This proves Claim 5. □

Stage 2: Extension to S-supplier auction. To eliminate a variable, we assume without loss of generality that supplier $i = 1$ in the statement of Proposition 1.

Claim 6. *Let $c_1(\vec{x}) + \pi < c_s(\vec{x})$ $\forall s = 2, \ldots, S$, where $\pi > \epsilon$. There exists a feasible scoring rule f that satisfies* (8)-(9) *and causes supplier* 1 *to win the S-supplier auction with attribute bid $\vec{x}$ and profit at least π.*

Proof. Set $d_{as} = c_{as}(\vec{x}) - c_{a1}(\vec{x})$, $D_s = \sum_{a=1}^{A} d_{as}$, and $\hat{c}_{as} = c_{as} - d_{as} + D_s/A$. Suppose f is feasible with respect to the cost surface c_s (and hence $\hat{c}_s$). Then f induces interior attribute bids from the cost functions c_s and $\hat{c}_s$; furthermore, since the cost functions differ from each other by a constant, they both must yield the same attribute bid $\vec{z}_s$. Defining

$$\begin{aligned}
\hat{S}_{as} &\triangleq \max_z \{f_a(z) - \hat{c}_{as}(z)\}, \\
&= \max_z \{f_a(z) - c_{as}(z)\} + d_{as} - \frac{D_s}{A}, \\
&= f_a(z_{as}) - c_{as}(z_{as}) + d_{as} - \frac{D_s}{A}
\end{aligned}$$

implies that $\sum_{a=1}^{A} \hat{S}_{as} = S_s$. We construct a scoring rule f that is feasible for the c_s cost surfaces, with $\hat{S}_{as} < \hat{S}_{a1} - \pi/A \ \forall \ a, s$, and $\vec{z}_1 = \vec{x}$. The first two conditions imply that when f is announced to all S suppliers (with actual cost surfaces c_s), supplier 1 wins with profit at least π and (9) is satisfied; by also ensuring that $\hat{S}_{as} > \epsilon/A \ \forall \ s$, we guarantee that f satisfies (8). Once again, we construct f dimension-by-dimension. We present the construction for fixed a.

Notice that $\hat{c}_{a1} = c_{a1}$, and

$$\begin{aligned}
\hat{c}_{as}(\vec{x}) - \hat{c}_{a1}(\vec{x}) &= c_{as}(\vec{x}) - d_{as} + \frac{D_s}{A} - c_{a1}(\vec{x}), \\
&= c_{as}(\vec{x}) - (c_{as}(\vec{x}) - c_{a1}(\vec{x})) + \frac{D_s}{A} - c_{a1}(\vec{x}), \\
&= \frac{D_s}{A} > \frac{\pi}{A},
\end{aligned}$$

where the inequality follows because $D_s/A = (c_s(\vec{x}) - c_1(\vec{x}))/A$. Let $\hat{l}_{a1}$ denote the line tangent to $\hat{c}_{a1}$ at x_a. We consider two cases, (c1) and (c2).

(c1): $\hat{l}_{a1} + \pi/A$ does not intersect any curve $\hat{c}_{as}, \ s \neq 1$. In this case, $\hat{l}_{a1} + \pi/A$ lies below $\hat{c}_{as}, \ s \neq 1$. Choose $f_a(z) = \gamma_{a1} z^{\gamma_{a2}}$ such that $f_a(x_a) = \hat{c}_{a1}(x_a) + \pi/A + h_a$, $f'_a(x_a) = \hat{c}'_{a1}(x_a)$,

$h_a \in \mathbb{R}^+$. Solving for γ_{a1}, γ_{a2} yields

$$\gamma_{a1} = [\hat{c}_{a1}(x_a) + \frac{\pi}{A} + h_a] x_a^{-\gamma_{a2}}, \qquad \gamma_{a2} = \frac{\hat{c}'_{a1}(x_a) x_a}{\hat{c}_{a1}(x_a) + \frac{\pi}{A} + h_a}.$$

The function f_a is increasing, since $\gamma_{a1} > 0$. Choosing

$$h_a > \max\left\{\hat{c}'_{a1}(x_a)x_a - \hat{c}_{a1}(x_a) - \frac{\pi}{A},\ \max_s \{\hat{c}_{as}(x_a) - \hat{c}_{a1}(x_a)\}\right\}$$

ensures that $\gamma_{a1} < 1$ (and hence f_a is concave, $f'_a \to \infty$ as $z \to 0^+$, and $f'_a \to 0$ as $z \to \infty$), as well as ensuring that $\hat{S}_{as} > \pi/A > \epsilon/A$. Furthermore, for $s \neq 1$,

$$\begin{aligned} f_a(z) - \hat{c}_{as}(z) &< f_a(z) - \hat{l}_{a1}(z) - \frac{\pi}{A} \qquad \text{because } \hat{l}_{a1}(z) + \frac{\pi}{A} \text{ lies below } \hat{c}_{as}, \\ &< \hat{l}_{a1}(z) + \hat{S}_{a1} - \hat{l}_{a1}(z) - \frac{\pi}{A}, \\ &= \hat{S}_{a1} - \frac{\pi}{A}, \end{aligned}$$

where the second inequality follows because $\hat{l}_{a1}(z) + \hat{S}_{a1}$ is tangent to f_a at x_a, and f_a is concave. Hence, $\hat{S}_{as} = \max_z\{f_a(z) - \hat{c}_{as}(z)\} < \hat{S}_{a1} - \pi/A$ for $s \neq 1$.

(c2): $\hat{l}_{a1} + \pi/A$ intersects some curve $\hat{c}_{as}$, $s \neq 1$. Let $\Delta_a = \max_{z,s\neq 1}\{\hat{l}_{a1}(z) + \pi/A - \hat{c}_{as}(z)\}$. Set $\underline{z}_a < x_a$ such that $\hat{l}_{a1} + \pi/A$ intersects no curves $\hat{c}_{as},\ s \neq 1$, in $[\underline{z}_a, x_a]$, and set $\bar{z}_a > x_a$ such that $\hat{l}_{a1} + \pi/A$ intersects no curves $\hat{c}_{as},\ s \neq 1$, in $[x_a, \bar{z}_a]$.

In Claim 1 we showed that if $(z_1, g_1(z_1))$, $g'_1(z_1)$, and $(z_2, g_2(z_2))$, $g'_2(z_2)$ are a point-slope pair such that (28)-(29) hold, then we can find a function f such that (26)-(27) hold for $h > H,\ H \in \mathbb{R}^+$. Notice that if we set

$$z_1 = x_a, \qquad g_1(z_1) = \hat{c}_{a1}(x_a), \qquad g'_1(z_1) = \hat{c}'_{a1}(x_a),$$

$$z_2 = \underline{z}_a, \qquad g_2(z_2) = \hat{c}_{a1}(x_a) + \frac{\pi}{A} - \left(\Delta_a + \hat{c}'_{a1}(x_a)[x_a - \underline{z}_a]\right),$$

$$\text{and} \qquad g'_2(z_2) > \frac{\Delta_a}{x_a - \underline{z}_a} + \hat{c}'_{a1}(x_a), \qquad \text{then}$$

$$\begin{aligned}
g_2(z_2)+g_2'(z_2)[z_1-z_2] &= g_2(z_2)+g_2'(z_2)[x_a-\underline{z}_a], \\
&> g_2(z_2)+\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\underline{z}_a], \\
&= \hat{c}_{a1}(x_a)+\frac{\pi}{A}-(\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\underline{z}_a])+\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\underline{z}_a], \\
&= \hat{c}_{a1}(x_a)+\frac{\pi}{A}=g_1(z_1)+\frac{\pi}{A},
\end{aligned}$$

$$\begin{aligned}
\text{and} \qquad g_2(z_2) &= \hat{c}_{a1}(x_a)+\frac{\pi}{A}-\left(\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\underline{z}_a]\right), \\
&= \hat{c}_{a1}(x_a)+\hat{c}_{a1}'(x_a)[\underline{z}_a-x_a]+\frac{\pi}{A}-\Delta_a, \\
&< \hat{l}_{a1}(\underline{z}_a)+\frac{\pi}{A} \qquad \text{since } \Delta_a>0.
\end{aligned}$$

That is, (28)-(29) hold (with π/A in the role of α and $\hat{l}_{a1}$ in the role of l_1), and we can find an increasing, concave, smooth function $\underline{f}_a$ such that $\underline{f}_a'(z_a)\to\infty$ as $z_a\to 0^+$,

$$\underline{f}_a(\underline{z}_a)=\hat{c}_{a1}(x_a)+\frac{\pi}{A}-(\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\underline{z}_a])+\underline{h}_a,$$

$$\underline{f}_a(x_a)=\hat{c}_{a1}(x_a)+\frac{\pi}{A}+\underline{h}_a, \qquad \text{and} \qquad \underline{f}_a'(x_a)=\hat{c}_{a1}'(x_a),$$

for $\underline{h}_a>\underline{H}_a$, $H_a\in\mathbb{R}^+$. If we leave z_1, $g_1(z_1)$, and $g_1'(z_1)$ the same but take $z_2=\bar{z}_a$, $g_2(z_2)=\hat{c}_{a1}(x_a)+\pi/A-(\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\bar{z}_a])$, and $g_2'(z_2)<\frac{\Delta_a}{x_a-\bar{z}_a}+\hat{c}_{a1}'(x_a)$, then arguments similar to those above show that (28)-(29) hold (again with π/A in the role of α and $\hat{l}_{a1}$ in the role of l_1), and hence we can find an increasing, concave, smooth function $\bar{f}_a$ such that $\bar{f}_a'(z_a)\to 0$ as $z_a\to\infty$,

$$\bar{f}_a(x_a)=\hat{c}_{a1}(x_a)+\frac{\pi}{A}+\bar{h}_a,$$

$$\bar{f}_a(\bar{z}_a)=\hat{c}_{a1}(x_a)+\frac{\pi}{A}-(\Delta_a+\hat{c}_{a1}'(x_a)[x_a-\bar{z}_a])+\bar{h}_a, \qquad \text{and} \qquad \bar{f}_a'(x_a)=\hat{c}_{a1}'(x_a),$$

for $\bar{h}_a>\bar{H}_a$, $\bar{H}_a$ some positive real number. Setting

$$\underline{h}_a=\bar{h}_a=h_a>\max\left\{\underline{H}_a,\ \bar{H}_a,\ \max_s\{\hat{c}_{as}(x_a)-\hat{c}_{a1}(x_a)\}\right\}$$

and

$$f_a(z)=\begin{cases}\underline{f}_a(z) & \text{if } z\le x_a,\\ \bar{f}_a(z) & \text{if } z>x_a\end{cases}$$

ensures that f_a is smooth, concave, increasing, $f_a' \to \infty$ as $z \to 0^+$, $f_a' \to 0$ as $z \to \infty$, and $\hat{S}_{as} > \epsilon/A$.

It remains to show that $\hat{S}_{a1} - \pi/A > \hat{S}_{as}\ \forall\ s \neq 1$. Suppose not, that is, suppose $\hat{S}_{a1} - \pi/A \leq \hat{S}_{as}$ for some $s \neq 1$. Let z_{as} denote supplier s's attribute bid in dimension a induced by f. We first show that

$$\hat{c}_{as}(z_{as}) < \hat{l}_{a1}(z_{as}) + \frac{\pi}{A}. \tag{44}$$

Suppose $\hat{c}_{as}(z_{as}) \geq \hat{l}_{a1}(z_{as}) + \pi/A$. Then

$$\begin{aligned}
\hat{S}_{as} = f_a(z_{as}) - \hat{c}_{as}(z_{as}) &\leq f_a(z_{as}) - \hat{l}_{a1}(z_{as}) - \frac{\pi}{A}, \\
&< \hat{l}_{a1}(z_{as}) + \hat{S}_{a1} - \hat{l}_{a1}(z_{as}) - \frac{\pi}{A}, \\
&= \hat{S}_{a1} - \frac{\pi}{A},
\end{aligned}$$

which is a contradiction. For the second inequality, we have made use of the fact that $\hat{l}_{a1} + \hat{S}_{a1}$ is tangent to f_a at x_a, and f_a is concave. We now use (44) along with the definition of f_a to derive a contradiction if $\hat{S}_{a1} - \pi/A \leq \hat{S}_{as}$. We treat the case $z_{as} \leq x_a$; the complement case is treated similarly, and is omitted for brevity. Let l_{af} denote the line tangent to f_a at $\underline{z}_a$.

$$\begin{aligned}
f_a(z_{as}) - \hat{c}_{as}(z_{as}) &< f_a(z_{as}) - (\hat{l}_{a1}(z_{as}) + \frac{\pi}{A} - \Delta_a) \qquad \text{by the definition of } \Delta_a, \\
&< l_{af}(z_{as}) - \hat{l}_{a1}(z_{as}) - \frac{\pi}{A} + \Delta_a, \\
&< l_{af}(\underline{z}_a) - \hat{l}_{a1}(\underline{z}_a) - \frac{\pi}{A} + \Delta_a, \\
&= \hat{c}_{a1}(x_a) + \frac{\pi}{A} - (\Delta_a + \hat{c}_{a1}'(x_a)[x_a - \underline{z}_a]) + h_a \\
&\qquad - (\hat{c}_{a1}(x_a) + \hat{c}_{a1}'(x_a)[\underline{z}_a - x_a]) - \frac{\pi}{A} + \Delta_a, \\
&= h_a, \\
&= f_a(x_a) - (\hat{c}_{a1}(x_a) + \frac{\pi}{A}) = \hat{S}_{a1} - \frac{\pi}{A},
\end{aligned}$$

which contradicts $\hat{S}_{a1} - \pi/A \leq \hat{S}_{as}$. The third inequality follows because $l_{af} - \hat{l}_{a1}$ is an increasing function (since l_{af} is steeper than $\hat{l}_{a1}$), and $z_{as} < \underline{z}_a$ by (44) together with the assumption that $z_{as} \leq x_a$. Taking a step back, notice that in defining f_a in case 1 we used only x_a, $c_{a1}(x_a)$ and $c'_{a1}(x_a)$ (three parameters), and in case 2 we used these three in addition to $\underline{z}$, $g_2(\underline{z})$, $g'_2(\underline{z})$ (when defining $\underline{f}$), $\bar{z}$, $g_2(\bar{z})$, $g'_2(\bar{z})$ (when defining $\bar{f}$), and h_a – i.e., f is a function of ten non-redundant parameters. $\square$

Claim 7. *Let $c_1(\vec{x}) + \pi < c_s(\vec{x})$ $\forall s = 2, \ldots, S$, where $\pi > \epsilon$. If there exists a $j \neq 1$ such that $T_1 + \pi$ intersects c_j, then $(c_1(\vec{x}) + \pi, \vec{x})$ is enforceable.*

Proof. Let f denote the scoring rule for the two-supplier auction constructed in Claim 5 such that $S_1(f) - \pi = S_j(f)$, where $S_s(f)$ denotes supplier s's maximum dropout score if f is announced. Since for all a $f'_a(z) \to \infty$ as $z \to 0^+$ and $f'_a(z) \to 0$ as $z \to \infty$, f is feasible (induces positive, finite attribute bids) in the S-supplier auction.

(c1): $S_1(f) - \pi \geq S_s(f)$ $\forall$ $s \neq 1$. In this case, announcing f enforces $(c_1(\vec{x}) + \pi, \vec{x})$ and we are done.

(c2): $\exists$ $k \neq 1$ such that $S_1(f) - \pi < S_k(f)$. Let g be the feasible scoring rule satisfying (8)-(9) constructed in Claim 6 such that $S_1(g) - \pi > S_s(g)$ for all $s \neq 1$. In the remainder of our proof, we will show that there exists a convex combination of f and g that enforces $(c_1(\vec{x}) + \pi, \vec{x})$. Notice that since f and g are feasible, the same is true for convex combinations of f and g. Define

$$S_s(\lambda) \triangleq \sum_{a=1}^{A} (\lambda f_a + (1 - \lambda) g_a)(x_{as}(\lambda)) - \sum_{a=1}^{A} c_{as}(x_{as}(\lambda)), \tag{45}$$

where $\vec{x}_s(\lambda)$ is the attribute bid of supplier s when $\lambda f + (1 - \lambda) g$ is announced.

Notice that since $\nabla f(\vec{x}) = \nabla g(\vec{x}) = \nabla c_1(\vec{x})$, supplier 1's attribute bid $\vec{x}_1(\lambda)$ will equal $\vec{x}$ for all $\lambda \in [0, 1]$. This implies that $S_1(\lambda) = \lambda S_1(f) + (1 - \lambda) S_1(g) > \pi + \epsilon$ for all $\lambda \in [0, 1]$, because $S_1(f), S_1(g) > \pi + \epsilon$ since f and g satisfy (8). By these observations, if there exists λ^* for which $S_1(\lambda^*) - \pi \geq S_s(\lambda^*)$ holds for all $s \neq 1$, and holds for some supplier $s \neq 1$ with equality, then $\lambda^* f + (1 - \lambda^*) g$ is a feasible scoring rule that satisfies (8)-(9) and enforces

$(c_1(\vec{x}) + \pi, \vec{x})$ in the S-supplier auction. To find such a λ^*, we first show that $S_s(\lambda)$ is a continuous function of λ.

We begin by showing that $x_{as}(\lambda)$ is continuous in λ, where for convenience we drop the subscript s in the general proof that follows. We first suppose that $x_a(0) \leq x_a(1)$, and show that $x_a(0) \leq x_a(\lambda) \leq x_a(1)$ for all $\lambda \in [0,1]$ (the proof in the case $x_a(1) \leq x_a(0)$ is essentially the same, with the roles of $x_a(0)$ and $x_a(1)$ reversed). We prove $x_a(0) \leq x_a(\lambda)$ for all $\lambda \in [0,1]$; the other inequality follows similarly.

By the definition of $x_a(0)$, $z_a < x_a(0)$ implies $g_a'(z_a) > c_a'(z_a)$, and similarly the definition of $x_a(1)$ implies that for $z_a < x_a(1)$, $f_a'(z_a) > c_a'(z_a)$. Since $x_a(0) \leq x_a(1)$, we have that for $z_a < x_a(0)$, $g_a'(z_a) > c_a'(z_a)$ and $f_a'(z_a) > c_a'(z_a)$, and hence $\lambda f_a'(z_a) + (1-\lambda) g_a'(z_a) > c_a'(z_a)$ for all $\lambda \in [0,1]$. By the above analysis and f_a', g_a', and $-c_a'$ strictly decreasing, for fixed λ the root $x_a(\lambda)$ of $\lambda f_a'(z_a) + (1-\lambda) g_a'(z_a) - c_a'(z_a)$ must lie to the right of $x_a(0)$, which is what we wanted to show.

Next, we consider perturbing λ to $\lambda + \delta$, for $|\delta|$ small (where δ must be positive if $\lambda = 0$, negative if $\lambda = 1$), and show that for any fixed $\eta > 0$, $\exists\, \hat{\delta} > 0$ such that $|x_a(\lambda) - x_a(\lambda+\delta)| < \eta$ for all $|\delta| < \hat{\delta}$.

First, note that for $z_a \in [x_a(0), x_a(1)]$ and δ fixed,

$$
\begin{aligned}
& |((\lambda+\delta) f_a' + (1-\lambda-\delta) g_a' - c_a')(z_a) - (\lambda f_a' + (1-\lambda) g_a' - c_a')(z_a)| \\
=\ & |\delta (f_a' - g_a')(z_a)|, \\
<\ & |\delta| \max_{z_a \in [x_a(0), x_a(1)]} |(f_a' - g_a')(z_a)|, \\
\leq\ & |\delta| B,
\end{aligned}
$$

where some $B < \infty$ exists since $[x_a(0), x_a(1)]$ is compact and $f_a' - g_a'$ is continuous. In other words, over $[x_a(0), x_a(1)]$, the curve $(\lambda+\delta) f_a' + (1-\lambda-\delta) g_a' - c_a'$ lies between the two curves $\lambda f_a' + (1-\lambda) g_a' - c_a' + \delta B$ and $\lambda f_a' + (1-\lambda) g_a' - c_a' - \delta B$. Let $w_a(\delta)$ be the root of the strictly decreasing function $\lambda f_a' + (1-\lambda) g_a' - c_a' + \delta B$. By the bound given above, and the fact that

$x_a(\lambda+\delta) \in [x_a(0), x_a(1)]$, we have that $w_a(-\delta)$ and $w_a(\delta)$ sandwich $x_a(\lambda+\delta)$. Since the same is true for $x_a(\lambda)$, we have our desired result if we can find $\hat{\delta}$ such that $|\delta| < \hat{\delta}$ implies $|w_a(\delta) - w_a(-\delta)| < \eta$.

Pick $\hat{\delta}$ such that $\hat{\delta}B < -(\lambda f'_a + (1-\lambda)g'_a - c'_a)(x_a(\lambda) + \eta/2)$. The right side is positive by the the fact that $x_a(\lambda)$ is a root of the decreasing function $\lambda f'_a + (1-\lambda)g'_a - c'_a$, and $\eta > 0$. Furthermore, choosing $\hat{\delta}$ in this way implies that $(\lambda f'_a + (1-\lambda)g'_a - c'_a + \hat{\delta}B)(x_a(\lambda) + \eta/2) < 0$. Since $(\lambda f'_a + (1-\lambda)g'_a - c'_a + \hat{\delta}B)(x_a(\lambda)) = \hat{\delta}B > 0$, we have that $w_a(|\delta|) \in (x_a(\lambda), x_a(\lambda) + \eta/2)$ for $|\delta| < \hat{\delta}$. By also ensuring that $\hat{\delta}B < (\lambda f'_a + (1-\lambda)g'_a - c'_a)(x_a(\lambda) - \eta/2)$, we have $w_a(-|\delta|) \in (x_a(\lambda) - \eta/2, x_a(\lambda))$ for $|\delta| < \hat{\delta}$, and the continuity of $x_a(\lambda)$ follows.

Since $(\lambda f_a + (1-\lambda)g_a - c_a)(x_a(\lambda))$ is a composition of continuous functions, and the sum of continuous functions is continuous, the righthand side of (45) is continuous in λ – i.e., $S_s(\lambda)$ is continuous in λ for all s.

By our choice of g we have that $S_1(0) - \pi > S_s(0)$ for all $s \neq 1$. Set

$$\lambda^* = \min\{\lambda | S_s(\lambda) = S_1(\lambda) - \pi \quad \text{for some } s \neq 1,\ 0 < \lambda < 1\}. \tag{46}$$

By the continuity of the $S_s(\lambda)$'s and the fact that $S_k(1) > S_1(1) - \pi$, such a λ^* exists. As noted below equation (45), for such a λ^*, $\lambda^* f + (1-\lambda^*)g$ enforces $(c_1(\vec{x}) + \pi, \vec{x})$ in the S-supplier auction. Notice that the scoring rule in this case is a function of 18 parameters: seven come from the function f, ten come from the function g, and the last is the parameter λ^*.

□

A2. Proof of Proposition 2. In the statement of the Proposition, we tacitly assume that $M_i > \epsilon$; otherwise, the auctioneer's valuation is everywhere below the suppliers' costs plus the minimum bid increment ϵ, meaning that no auction would be pursued (which is not interesting). We break the proof of Proposition 2 into three cases, and begin with an observation that will be useful in all three.

Recall that the definition of optimality was made in relation to the optimum of (7)-(9), (13). Let u^* be the optimum of the general version of (7)-(9), (13), (i.e., (7)-(9), (13), with v, c_1 and c_2 instead of their parameterized counterparts). Towards proving the optimality of supplier i, we first show that $u^* \leq M_i - \epsilon$. In contrast to the formalization of (7)-(9), (13), here we make no assumption on how the suppliers' subscripts are ordered for a given scoring rule. For scoring rules f such that k and l are the top bidders and $S_k > S_l + \epsilon$, the appropriate mathematical program is

$$\begin{aligned} \max_{f} \quad & v(\vec{x}_k^*) - f(\vec{x}_k^*) + S_l \\ \text{subject to} \quad & S_s = f(\vec{x}_s^*) - c_s(\vec{x}_s^*), \quad s = k, l \\ & S_l > \epsilon. \end{aligned} \tag{47}$$

Examining the objective function (47), we see that

$$\begin{aligned} v(\vec{x}_k^*) - f(\vec{x}_k^*) + S_l &= v(\vec{x}_k^*) - c_k(\vec{x}_k^*) - (f(\vec{x}_k^*) - c_k(\vec{x}_k^*)) + S_l, \\ &= v(\vec{x}_k^*) - c_k(\vec{x}_k^*) - S_k + S_l, \\ &\leq v(\vec{x}_k^*) - c_k(\vec{x}_k^*) - \epsilon \qquad \text{since } S_k > S_l + \epsilon, \\ &\leq \max_{\vec{x}} \{v(\vec{x}) - c_k(\vec{x})\} - \epsilon, \\ &= M_k - \epsilon. \end{aligned}$$

Hence, we've shown that $u^* \leq M_i - \epsilon$. This observation motivates the first case, which provides conditions in which the auctioneer is guaranteed to do no worse than u^* by simply announcing the true valuation function v as the scoring rule; in this case, we consider v to be optimal (as discussed below equation (13)).

(c1): $\exists\, k \neq i$ such that $M_k \geq M_i - \epsilon$. Suppose we announce the true valuation function v as the scoring rule. We need to show that, no matter who (i or k) wins, the post-auction utility to the auctioneer is at least $M_i - \epsilon$ (and hence both suppliers are optimal). Suppose supplier i wins the auction. Since we assume the winning supplier wins at the losing supplier's highest possible score, the profit at which supplier i wins the auction is at most $S_i - S_k < \epsilon$. Because the auctioneer announced his true valuation function as the scoring rule, the utility to the auctioneer equals the winner's maximum dropout score minus the winner's profit, i.e., is at least $S_i - \epsilon = M_i - \epsilon$. Next, suppose that supplier k wins the auction. Supplier i drops out at the highest possible losing score; since $S_i < S_k$, supplier k wins the auction with zero profit, and the post-auction utility of the auctioneer is $S_k \geq M_i - \epsilon$. Here we have assumed that supplier k actually bids, i.e., $S_k \geq \epsilon$. If instead $S_k = M_k < \epsilon$, then the utility for the auctioneer is still at least $M_i - \epsilon$: supplier i's winning score must have magnitude at least ϵ (ϵ is the minimum bid increment), and $M_k < \epsilon$ implies that $M_i \leq 2\epsilon$.

In the remaining two cases of the proof, announcing the true valuation function does not generate at least $M_i - \epsilon$ in utility for the auctioneer, and we show that supplier i wins under an optimal solution to (7)-(9), (13). Let $M_i = v(\vec{x}^*) - c_i(\vec{x}^*)$, and let T_i be the hyperplane tangent to c_i at $\vec{x}^*$. Because v is concave and c_i is convex, notice that we have $T_i + M_i$ tangent to v at $\vec{x}^*$.

(c2): $M_i - \epsilon > M_s$ for all $s \neq i$, and $\exists\ j \neq i$ such that $T_i + \epsilon$ intersects c_j. Corollary 1 implies that $(c_i(\vec{x}^*) + \epsilon, \vec{x}^*)$ can be enforced, yielding a payoff of $M_i - \epsilon$ for the auctioneer. Hence, i is optimal.

(c3): $M_i - \epsilon > M_s$ for all $s \neq i$, and $\nexists\ j \neq i$ such that $T_i + \epsilon$ intersects c_j. Our argument relies on hyperplanes. Let $\pi = \min_{\vec{z}, s \neq i}\{c_s(\vec{z}) - T_i(\vec{z})\}$. The hyperplane $T_i + \pi$ must lie below or tangent to c_s for all $s \neq i$. We use $T_i + M_i$ and $T_i + \pi$ to bound v from above, and c_s from below, respectively. Consider a scoring rule for which supplier $s \neq i$ wins. Let $\vec{w}$ be supplier

s's winning attribute bid. The auctioneer's utility is at most $v(\vec{w}) - c_s(\vec{w})$, where

$$\begin{aligned} v(\vec{w}) - c_s(\vec{w}) &< (T_i(\vec{w}) + M_1) - (T_i(\vec{w}) + \pi), \\ &= M_i - \pi, \end{aligned}$$

which implies that

$$\max_{\vec{w}}\{v(\vec{w}) - c_s(\vec{w})\} = M_s < M_i - \pi.$$

Applying Corollary 1, $(c_i(\vec{x}^*) + \pi + \delta, \vec{x}^*)$ is enforceable as $\delta \to 0^+$, yielding for fixed δ a payoff of $M_i - \pi - \delta$ for the auctioneer. Choosing $\delta < M_i - \pi - \max_{s \neq i} M_s$ shows that with supplier i winning we can improve on the best case for which supplier $s \neq i$ wins. Hence, the optimum u^* is achieved for some scoring rule f for which supplier i wins (supplier i is optimal). □

A3. Simplification of (15)-(18). We simplify the mathematical program (15)-(18) by finding a closed-form solution to the innermost minimization in (18). Since the c_{as} are convex and increasing in x_a, for fixed s the innermost minimization occurs at z_a^* such that, for $a = 1, \ldots, A$, $\left.\frac{\partial c_{as}(z_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial z_a}\right|_{z_a = z_a^*} = \left.\frac{\partial c_{ai}(z_a, \theta_{ai1}, \ldots, \theta_{aiP})}{\partial z_a}\right|_{z_a = x_a}$ if $\left.\frac{\partial c_{as}(z_a, \theta_{as1}, \ldots, \theta_{asP})}{\partial z_a}\right|_{z_a = 0} < \left.\frac{\partial c_{ai}(z_a, \theta_{ai1}, \ldots, \theta_{aiP})}{\partial z_a}\right|_{z_a = x_a}$, or at $z_a^* = 0$, otherwise. In words, if a point exists at which the tangent to c_{as} is parallel to c_{ai}'s tangent at x_a, we take this point as z_a^*; if such a point does not exist, then the minimization occurs at the left endpoint of the feasible interval in dimension a, and thus we take $z_a^* = 0$.

By incorporating z_a^*, we get

$$\begin{aligned} \pi \geq \min_{s \neq i} \Bigg\{ & \sum_{a=1}^{A} c_{ak}(z_a^*, \theta_{ak1}, \ldots, \theta_{akP}) - \sum_{a=1}^{A} c_{ai}(x_a, \theta_{ak1}, \ldots, \theta_{akP}) \\ & - \sum_{a=1}^{A} (z_a^* - x_a) \left.\frac{\partial c_{ai}(z_a, \theta_{ai1}, \ldots, \theta_{aiP})}{\partial z_a}\right|_{z_a = x_a} \Bigg\} \end{aligned} \qquad (48)$$

in place of (18). The last two factors within the right side of (48) are the expanded version of $-T_i(\vec{z}^*)$. Notice that in solving (15)-(17), (48) we need only search over $\vec{x}$'s for which supplier i's cost is below the auctioneer's true valuation, i.e., $\sum_{a=1}^{A} c_{ai}(x_a, \theta_{ai1}, \ldots, \theta_{aiP}) < \sum_{a=1}^{A} v_a(x_a, \psi_{a1}, \ldots, \psi_{aP})$. Also, the search can be terminated if and when the objective value reaches its upper bound of $M_i - \epsilon$.

A4. Optimal Scoring Rule in Step 3. Let $\vec{x}^*, \pi^*$ be the optimal solution to (15)-(17), (48), and let T_i be the hyperplane tangent to c_i at $\vec{x}^*$.

(c1): T_i does not intersect any c_s, $s \neq i$. Let $\pi_s = \min_{\vec{z}}\{c_s(\vec{z}) - T_i(\vec{z})\}$ and let $j = \arg\min_{s \neq i} \pi_s$, and note that $\pi^* = \pi_j$ (otherwise π^* is not optimal). For $\hat{\delta} > 0$, small, let f be the scoring rule constructed in Claim 5 to enforce $(c_i(\vec{x}^*) + \pi^* + \hat{\delta}, \vec{x}^*)$ in the two-supplier auction between i and j. By our choice of f, $S_j = S_i - \pi^* - \hat{\delta}$ when f is announced. Since f induces supplier i to bid attribute level $\vec{x}^*$, we must have that $T_i + S_i$ is tangent to f at $\vec{x}^*$, and also that $T_i + \pi^*$ bounds c_s from below for all $s \neq i$. Together these observations imply that, if f is announced to all S suppliers, $S_s < S_i - \pi^*$ for all $s \neq i$, and supplier i wins the auction with profit at most $\pi^* + \hat{\delta}$ (by $S_j = S_i - \pi^* - \hat{\delta}$). That is, f enforces $(c_i(\vec{x}^*) + \pi^* + \delta, \vec{x}^*)$ $(\delta \leq \hat{\delta})$ in the S-supplier auction (since $\hat{\delta}$ is arbitrarily small, f is considered optimal). That is, the optimal scoring rule is f in (19).

(c2): T_i intersects c_j for some $j \neq i$, and v satisfies (8). In this case, we show that Claim 7 can be applied in a special way. In particular, we show that the true valuation function v satisfies the conditions of Claim 6; this implies that Claim 7 can be applied to construct an optimal scoring rule of the form $\lambda^* f + (1 - \lambda^*)v$, where f is the scoring rule constructed in Claim 5 to enforce $(c_i(\vec{x}^*) + \pi^*, \vec{x}^*)$ in the two-supplier auction between i and j. To rephrase the conditions of Claim 6, we need to show that v is a feasible scoring rule satisfying (8)-(9) that causes supplier i to win the auction with attribute bid $\vec{x}^*$ and profit at least π^*.

To see that v induces supplier i to bid $\vec{x}^*$, we suppose not (suppose that v induces bid $\vec{z} \neq \vec{x}^*$) and derive a contradiction by showing that the payoff can be improved (and enforceability maintained) by moving slightly away from $\vec{x}^*$ toward $\vec{z}$. Let a be such that $x_a^* \neq z_a$, and let $d = z_a - x_a^*$. Then, for $0 < \delta \leq 1$,

$$v_a(x_a^* + \delta d) - c_{ai}(x_a^* + \delta d) > v_a(x_a^*) - c_{ai}(x_a^*), \tag{49}$$

since $v_a - c_{ai}$ is strictly concave, and is maximized at z_a. Clearly, for δ small, we can ensure that at the new point $(x_1^*, \ldots, x_a^* + \delta d, \ldots, x_A^*)$ that we move to, the hyperplane tangent to c_i will still intersect the surface c_j. This implies that $(c_i(x_1^*, \ldots, x_a^* + \delta d, \ldots, x_A^*) + \pi^*, (x_1^*, \ldots, x_a^* + \delta d, \ldots, x_A^*))$ is enforceable, which together with (49) contradicts the fact that $\vec{x}^*$, π^* is optimal. Hence, $\vec{x}^* = \vec{z}$, so announcing v induces i to bid $\vec{x}^*$.

Notice that constraint $\pi \geq \epsilon$ is tight for our optimal solution in this case; if not, since $T_i + \pi^*$ intersects c_j, we could decrease π^* and still maintain enforceability, contradicting the optimality of π^*. Furthermore, because we did not exit the three-step method of §2.5 at step 3, we have $M_i - \epsilon > M_s$ for all $s \neq i$. That is, announcing v yields a profit of at least ϵ for supplier i, and v satisfies equation (9). Since equation (8) holds by assumption, we are done.

(c3): T_i intersects c_j for some $j \neq i$, and v does not satisfy (8). In this case, the general construction of Proposition 1's "$\Rightarrow$" direction proof (§A1.2) can be applied to enforce $(c_i(\vec{x}^*) + \pi^*, \vec{x}^*)$.