

Properties of scoring auctions

John Asker*

and

Estelle Cantillon**

This article studies scoring auctions, a procedure commonly used to buy differentiated products: suppliers submit offers on all dimensions of the good (price, level of nonmonetary attributes), and these are evaluated using a scoring rule. We provide a systematic analysis of equilibrium behavior in scoring auctions when suppliers' private information is multidimensional (characterization of equilibrium behavior and expected utility equivalence). In addition, we show that scoring auctions dominate several other commonly used procedures for buying differentiated products, including menu auctions, beauty contests, and price-only auctions with minimum quality thresholds.

1. Introduction

■ In many procurement situations, the buyer cares about attributes other than price when evaluating the offers submitted by suppliers. Examples of nonmonetary attributes that buyers care about include lead time, time to completion, and quality. Buyers have adopted several practices for dealing with these situations. Some use detailed request-for-quotes that specify minimum standards that the offers need to satisfy, and then evaluate the submitted bids based on price only. Others select a small set of potential suppliers and negotiate on all dimensions of the contract with each of them.

A third option is to combine the competition induced by a request-for-quote with the flexibility in terms of contract specification offered by negotiation. Several procedures belong to this category. In a “menu auction,” the buyer lets suppliers submit menus of price and nonmonetary attributes, and choose the combination that best suits his needs. In a “beauty contest,” the buyer tells suppliers he cares about other attributes than price but requests a single offer from them. Again, he chooses the offer he prefers from the received offers. In a scoring auction, the buyer

* Stern School of Business and New York University; jasker@stern.nyu.edu.

** FNRS, Université Libre de Bruxelles (ECARES), and CEPR; estelle.cantillon@ulb.ac.be.

This article supersedes an earlier paper distributed under the title “Equilibrium in Scoring Auctions.” We thank Max Bazerman, Matthew Jackson, Luca Rigotti, Al Roth, Nicolas Sahuguet, Sasha Wolf, as well as audiences at Barcelona Auction Workshop 2005, Case Western Reserve, CORE, Decentralization Conference 2004, ECARES, Essex, Gerzensee, Harvard/MIT, Inform 2003, London Business School, London School of Economics, New York University, Ohio State University, Rome Auction Conference 2004, University College London, and WZB Berlin for useful conversation and suggestions. The editor and two anonymous referees provided useful comments. Ioannis Ioannou provided excellent research assistance. The financial support of the Division of Research at Harvard Business School is gratefully acknowledged.

announces the way he will rank the different offers, that is, the scoring rule; suppliers submit an offer on all dimensions of the product, and the contract is awarded to the supplier who submitted the offer with the highest score according to the scoring rule.

In this article, we study the properties of scoring auctions in which price enters linearly into the scoring rule. Examples of such scoring auctions include “A+B bidding” for highway construction work in the United States, where the highway procurement authorities evaluate offers on the basis of their costs as well as time to completion, weighted by a road user cost,¹ and auctions for electricity reserve supply (Bushnell and Oren, 1994; Wilson, 2002). The European Union has recently adopted a new public procurement directive. The new law allows for two different award criteria: lowest cost and best economic value. The new provisions require that the procurement authority publishes *ex ante* the relative weighting of each criterion used when best economic value is the basis for the award.² In effect, the new law mandates the use of scoring auctions. This is significant, as public procurement in the European Union is estimated at about 16% of GDP.³ The use of scoring auctions is also gaining favor in the private sector, with several procurement software developers incorporating scoring capability in their auction designs.

A distinguishing feature of our model is that suppliers’ private information about their cost is multidimensional. This means that the low-cost supplier for the base option is not necessarily the low-cost supplier when it comes to increasing quality on some other dimension. It allows us to consider the likely situation where firms differ in their fixed and variable costs of production. Our motivation for allowing multidimensional private information is to build a model of scoring auctions that can generate equilibrium predictions that mimic what is observed in the data. When private information is one-dimensional, equilibrium offers can be parameterized by a single parameter and describe a curve in the price-attributes space (Che, 1993). Our model does not suffer from this severe limitation.

We derive two sets of results. First, we characterize equilibrium behavior in scoring auctions when private information is multidimensional and the scoring rule is linear in price. We prove that the multidimensionality of suppliers’ private information can be reduced to a single dimension (their “pseudotype”) that is sufficient to characterize equilibrium outcomes in these auctions (Theorem 1). This allows us to establish a correspondence between the set of scoring auctions and the set of standard single-object one-dimensional independent private value (IPV) auction environments (Corollary 1). The equilibrium in the scoring auction inherits the properties of the corresponding standard IPV auction (existence and uniqueness of equilibrium, efficiency, etc.). We also prove a new expected utility theorem for the buyer when private information is multidimensional and independently distributed, and the scoring rule is linear in price (Theorem 2).

Our second set of results compares scoring auctions to other common procedures used to buy differentiated products. We show that, from the buyer’s perspective, scoring auctions strictly dominate price-only auctions with minimum quality standards. They weakly dominate a menu auction and a beauty contest when an open ascending format is used (the open ascending format is often used for online procurement). When a sealed-bid “second-price” format is used, they weakly dominate a menu auction and strictly dominate a beauty contest. Finally, the ranking between the first-price scoring auction and the first-price menu auction is ambiguous: we find that some buyers prefer the menu auction whereas others prefer the scoring auction. Moreover, we establish that first-price menu auctions are always inefficient. Note that our purpose in this article is not to determine how optimally to buy a differentiated product but, instead, to study the properties of a commonly used and simple procedure for doing so, the scoring auction. Thus, our

¹ The road user cost is the (per-day) value of time lost owing to construction. By 2003, 38 states in the United States were using “A+B bidding” for large projects for which time is a critical factor. See, for instance, Arizona Department of Transport (2002) and Herbsman, Chen, and Epstein (1995).

² See Articles 55 and 56 of Directive 2004/17/EC and Articles 53 and 54 of Directive 2004/18/EC. If the authority does not resort to electronic auctions, it may publish a range of weightings for each criterion instead.

³ http://europa.eu.int/comm/internal_market/publicprocurement/introduction_en.htm.

second set of results provides a motivation for focusing on the scoring auction given its attractive properties.

□ **Related literature.** There are several papers studying scoring auctions. Most papers note, as we do, that, once the scoring rule is given, the maximum level of social welfare a supplier can produce (in our paper, the pseudotype) can be used to construct an equilibrium in these auctions. This involves a benign change of variables when private information is one-dimensional as in Che (1993) and Branco (1997), but the operation is not so anodyne when private information is multidimensional. Specifically, we show that such a reduction in dimensionality requires that (i) the scoring rule be linear in price, and (ii) that private information be independently distributed across suppliers, unless the auction format admits a dominant strategy equilibrium. The papers we are aware of that allow for multidimensional private information, Bushnell and Oren (1994; 1995), happen to satisfy these conditions (these papers derive the scoring rule that induces productive efficiency in an environment with multidimensional private information). There is also a series of papers on scoring auctions published in the computer science and operations research literature. The focus there is on implementability through practical online/iterative processes (see, e.g., Bichler and Kalagnanam, 2003; Parkes and Kalagnanam, 2005).⁴

Several recent papers study other auction environments with multidimensional private information. In some environments, bidder preferences, the structure of information, or the specific allocation mechanism suggest the locus of types likely to use the same bidding strategies at equilibrium. These pseudotypes are used to construct an equilibrium (see, e.g., Che and Gale, 1998; Fang and Parreiras, 2002; de Frutos and Pechlivanos, 2006). Our approach is identical, except for the fact that, in addition, we prove that no other relevant equilibrium exists. This allows us to derive a utility equivalence theorem and to leverage the analogy between our environment and the standard IPV environment.⁵

Che (1993) and Asker and Cantillon (2006) derive the optimal buying mechanism when quality matters. A scoring auction in which price enters linearly into the scoring rule implements the optimal scheme when private information is one-dimensional. Under some conditions on the crosspartial derivative of costs, the optimal scoring rule underweights quality relative to the true preference of the buyer. When private information is multidimensional, Asker and Cantillon show that the buyer is still interested in distorting qualities away from their efficient levels. However, the optimal scheme can no longer be implemented by a scoring auction with a scoring rule that is linear in price. Nevertheless, they provide numerical examples suggesting that such scoring auctions perform almost as well as the optimal scheme.

Finally, a few papers consider alternatives to scoring auctions. Che (1993) provides a qualitative argument for why scoring auctions are better than price-only auctions with minimum quality standards in his one-dimensional framework. Bichler and Kalagnanam (2003) look at the “second-score” menu auction. They focus on the “winner determination problem” for a given set of offers received, not on equilibrium behavior. Menu auctions can be seen as a common agency problem where multiple principals (the suppliers) compete in offering menus of contracts to an agent (the buyer). From suppliers’ perspective, menu auctions are also an example of a screening problem with random participation because a supplier’s offer is accepted only if it is better than the competing offers the buyer received. We draw on these literatures when we study the first-price menu auction. We consider menu auctions, beauty contests, and price-only auctions with minimum quality standards, and systematically compare the outcome of equilibrium in these auctions with that in scoring auctions.

⁴ A variant of scoring auctions are auctions that involve the purchase of multiple items but where the buyer cannot commit, at the auction, to the quantity purchased. The scoring rule is used for allocating the contract, although the final contract depends on the realized quantities. This creates incentive problems we ignore (see Athey and Levin, 2001; Chao and Wilson, 2002; Ewerhart and Fieseler, 2003).

⁵ A similar property (although through a much more subtle analogy to the standard IPV model) is exploited by Che and Gale (2006) to rank revenue in single-object auctions with multidimensional types and nonlinear payoffs.

The rest of the article is organized as follows. Section 2 describes the model and introduces the notion of pseudotype. Section 3 proves that the pseudotypes are sufficient statistics in our environment, and establishes the correspondence between scoring auctions and regular IPV auctions. Our expected utility equivalence theorem is proved in Section 4. Section 5 compares the outcome of scoring auctions with that of menu auctions, beauty contests, and auctions with minimum quality standards. Section 6 concludes.

2. Model

■ **Environment.** We consider a buyer seeking to procure an indivisible good for which there are N potential suppliers. The good is characterized by its price, p , and $M \geq 1$ nonmonetary attributes, $\mathbf{Q} \in \mathbb{R}_+^M$.

Preferences. The buyer values the good (p, \mathbf{Q}) at $v(\mathbf{Q}, t) - p$, where $t \in [\underline{t}, \bar{t}]$ indexes the buyer's taste for quality.⁶ Supplier i 's profit from selling good (p, \mathbf{Q}) is given by $p - c(\mathbf{Q}, \theta_i)$, where $\theta_i \in \mathbb{R}^K$, $K \geq 1$, is supplier i 's type. We allow suppliers to be flexible with respect to the level of nonmonetary attributes they can supply.⁷ We assume that v and c are twice continuously differentiable with $v_{\mathbf{Q}}, c_{\mathbf{Q}} > 0$, $v - c$ bounded, and $v_{\mathbf{Q}\mathbf{Q}} - c_{\mathbf{Q}\mathbf{Q}}$ negative definite. In particular, this allows for costs that are independent across attributes and convex in individual attributes. We partially order the type space by assuming that $c_{\theta_i} > 0$. When we analyze the first price menu auction, we will also impose $v_{\mathbf{Q}t} > 0$ and $v_{\mathbf{Q}\mathbf{Q}t}$ negative semidefinite.

Because social welfare is bounded and strictly concave in \mathbf{Q} , the first-best level of nonmonetary attributes for each supplier, $\mathbf{Q}^{FB}(\theta_i) = \arg \max\{v(\mathbf{Q}, t) - c(\mathbf{Q}, \theta_i)\}$, is well defined and unique.

Information. Preferences are common knowledge among suppliers and the buyer, with the exception of suppliers' types, θ_i , $i = 1, \dots, N$, and the buyer's taste parameter, t , which are privately observed. Types are independently distributed according to the continuous joint density function $f_i(\cdot)$ with support on a bounded and convex subset of \mathbb{R}^K with a nonempty interior, Θ_i . Taste is distributed according to the continuous density $h(\cdot)$. These density functions are common knowledge.

□ **Allocation mechanism.** We now introduce the scoring auction. We start with two definitions:

A *scoring rule* is a function $S : \mathbb{R}_+^{M+1} \rightarrow \mathbb{R} : (p, \mathbf{Q}) \rightarrow S(p, \mathbf{Q})$ that associates a score to any potential contract and represents a continuous preference relation over contract characteristics (p, \mathbf{Q}) . A scoring rule is quasilinear if it can be expressed as $\phi(\mathbf{Q}) - p$ or any monotonic increasing function thereof. We assume that the scoring rule is twice continuously differentiable and strictly increasing in \mathbf{Q} , and that the resulting "apparent social welfare," $\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta)$, is bounded and strictly concave in \mathbf{Q} for all θ . For simplicity, we let $\max_{\mathbf{Q}}\{\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta)\} \geq 0$ for all θ to ensure that all suppliers participate in the auction at equilibrium.

A *scoring auction* is an allocation mechanism where suppliers submit bids of the type $(p, \mathbf{Q}) \in \mathbb{R}_+^{M+1}$. Bids are evaluated according to a scoring rule. The winner is the bidder with the highest score. The outcome of the scoring auction is a probability of winning the contract, x_i , a score to fulfill when the supplier wins the contract, t_i^w , and a payment to the buyer in case he does not, t_i^l . A scoring auction is quasilinear when it uses a quasilinear scoring rule.

For example, in a first-score scoring auction, the winner must deliver a contract that generates the value of his winning score, that is, $t_i^w = S(p_i, \mathbf{Q}_i)$, $t_i^l = 0$. In an ascending scoring auction, the buyer progressively raises the required score to fulfill by any standing offer until all suppliers

⁶ Until Section 5, in which we consider alternative mechanisms to the scoring auction, nothing is lost if t is assumed to be common knowledge. We introduce the notation here for completeness.

⁷ Rezende (2004) studies a procurement model with fixed levels of nonmonetary attributes. In our model, the level of nonmonetary attributes is determined during the auction process.

but one drop out. t_i^W is the value of that score and $t_i^L = 0$. In a second-score scoring auction, the winner must deliver a contract that generates a score equal to the score of the second-best offer received.

Note that when the scoring rule does not correspond to the buyer's preference—something which might be in his interest (Che, 1993; Asker and Cantillon, 2006)—commitment is essential. In public procurement, this might be easily done. The process must often abide by a strict set of rules and procedures, so that, in effect, the call for tender (and thus the scoring rule) is legally binding for the buyer. In private procurement, this might be harder, although, in principle, the buyer could sign a contract with the bidders before bidding takes place in which he commits to use the scoring rule. Such a contract could be enforced through an independent third-party audit. Repetition is an alternative mechanism.

We now proceed to the analysis of bidding behavior in the scoring auction. Consider supplier i with type θ_i who has won the contract with a score to fulfill t_i^W . Supplier i will choose characteristics (p, \mathbf{Q}) that maximize his profit, that is,

$$\max_{(p, \mathbf{Q})} \{p - c(\mathbf{Q}, \theta_i)\} \quad \text{subject to } \phi(\mathbf{Q}) - p = t_i^W.$$

Substituting for p into the objective function yields

$$\max_{\mathbf{Q}} \{\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta_i) - t_i^W\}. \tag{1}$$

An important feature of (1) is that the optimal \mathbf{Q} is independent of t_i^W . Define

$$k(\theta_i) = \max_{\mathbf{Q}} \{\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta_i)\}.$$

We shall call $k(\theta_i)$ supplier i 's pseudotype. It is the maximum level of apparent social surplus that supplier i can generate. Bidders' pseudotypes are well defined as soon as the scoring rule is given. The set of supplier i 's possible pseudotypes is an interval in \mathbb{R} . The density of pseudotypes inherits the smooth properties of f_i . With this definition, supplier i 's expected profit is given by

$$x_i (k(\theta_i) - t_i^W) - (1 - x_i)t_i^L. \tag{2}$$

In (2), supplier i 's preference over contracts of the type (x_i, t_i^W, t_i^L) is entirely captured by his pseudotype. Only quasilinear scoring rules have this property when private information is multidimensional. Indeed, consider a more general scoring rule $S(p, \mathbf{Q})$. Assume that S is twice continuously differentiable, strictly increasing in \mathbf{Q} and strictly decreasing in p . Bidder i 's optimization problem becomes

$$\max_{(p, \mathbf{Q})} \{p - c(\mathbf{Q}, \theta_i)\} \quad \text{subject to } S(p, \mathbf{Q}) = t_i^W. \tag{3}$$

Let $\Psi(\mathbf{Q}, t_i^W)$ be the price required to generate a score of t_i^W with nonmonetary attributes \mathbf{Q} (Ψ is well defined because S is strictly decreasing in p and strictly increasing in \mathbf{Q} ; it is strictly increasing in \mathbf{Q} and strictly decreasing in t_i^W). The objective function of bidder i becomes

$$\max_{\mathbf{Q}} \{\Psi(\mathbf{Q}, t_i^W) - c(\mathbf{Q}, \theta_i)\},$$

and his expected payoff from contract (x_i, t_i^W, t_i^L) is given by

$$u(x_i, t_i^W, t_i^L; \theta_i) = x_i \max_{\mathbf{Q}} \{\Psi(\mathbf{Q}, t_i^W) - c(\mathbf{Q}, \theta_i)\} - (1 - x_i)t_i^L.$$

Suppose we can organize types in equivalence classes such that all types in a given class share the same preferences over contracts. Concretely, suppose that types θ_i and $\hat{\theta}_i \neq \theta_i$ belong to such a class. It must be that

$$u(x_i, t_i^W, t_i^L; \theta_i) = u(x_i, t_i^W, t_i^L; \hat{\theta}_i) \quad \text{if and only if } u(\hat{x}_i, \hat{t}_i^W, \hat{t}_i^L; \theta_i) = u(\hat{x}_i, \hat{t}_i^W, \hat{t}_i^L; \hat{\theta}_i) \tag{4}$$

for all pairs of contracts $(x_i, t_i^W, t_i^L), (\hat{x}_i, \hat{t}_i^W, \hat{t}_i^L)$.

Let $\mathbf{Q}(\theta_i, t_i^W) = \arg \max_{\mathbf{Q}} \{\Psi(\mathbf{Q}, t_i^W) - c(\mathbf{Q}, \theta_i)\}$. Condition (4) requires that $\frac{\partial}{\partial t_i^W} \Psi(\mathbf{Q}(\theta_i, t_i^W), t_i^W) = \frac{\partial}{\partial t_i^W} \Psi(\mathbf{Q}(\hat{\theta}_i, t_i^W), t_i^W)$. This equality will in general not be satisfied for $\hat{\theta}_i \neq \theta_i$ unless Ψ is separable in \mathbf{Q} and t_i^W . In turn, this requires that the scoring rule be quasilinear ($\Psi(\mathbf{Q}, t_i^W) = \phi(\mathbf{Q}) - t_i^W$ for a quasilinear scoring rule).⁸

Finally, we carry out one last simplification of the problem. Let $s_i = x_i t_i^W + (1 - x_i) t_i^L$ in (2). Given suppliers' risk neutrality and the linearity of the scoring rule, there is no loss in defining the outcome of a scoring auction as the pair (x_i, s_i) , rather than (x_i, t_i^W, t_i^L) . Suppliers' expected payoff is thus given by

$$x_i k(\theta_i) - s_i. \tag{5}$$

Notation. For the remainder, we adopt the following notation and conventions. The outcome function of a scoring auction is a vector of probabilities of winning (x_1, \dots, x_N) and scores to fulfill by each supplier, (s_1, \dots, s_N) . (If the outcome in a given scoring auction is stochastic, these are *distributions* over vectors of probabilities of winning and scores.) The arguments in these functions are the bids submitted by all suppliers, $\{(p_i, \mathbf{Q}_i)\}_{i=1}^N$.⁹ Later in the article, we will switch to a direct revelation mechanism approach where the outcome will be a function of suppliers' pseudotypes, $(k_1, \dots, k_N) \in \mathbb{R}^N$. To avoid introducing too much new notation, we shall make these the arguments of the x and s functions. We shall also write $x_i(k_i)$ to denote the expectation of x_i over the types of the other suppliers, $E_{k_{-i}} x_i(k_i, k_{-i})$. The arguments will be spelled out whenever confusion is possible.

3. A sufficient statistics result

■ Suppliers' pseudotypes are sufficient statistics in this environment if knowing the distribution of suppliers' pseudotypes is all one needs in order to describe the set of equilibrium outcomes of the auction and evaluate the buyer's expected payoff.

In this section, we prove that pseudotypes are sufficient statistics. Proving this result requires two steps. First, we show that all equilibria of the scoring auction are outcome equivalent to an equilibrium where suppliers are forced to submit bids only as a function of their pseudotypes. We define two equilibria as *outcome equivalent* if they both lead to the same distribution of outcomes (x_1, \dots, x_N) and (s_1, \dots, s_N) in the aggregate. Because outcome equivalence is not enough to guarantee that the buyer is indifferent among these equilibria, we next prove the stronger result that the equilibria in the scoring auction and in the constrained scoring auction have the same distribution of outcomes, *conditional on types*.

Lemma 1. All equilibria of a quasilinear scoring auction are outcome equivalent to an equilibrium where bidders with the same pseudotypes adopt the same strategies.

Proof. Consider any equilibrium $(\mathcal{E}_1, \dots, \mathcal{E}_N)$, where \mathcal{E}_i is a mapping from Θ_i to a distribution over $(p, \mathbf{Q}) \in \mathbb{R}^{M+1}$. Then, for all i , for all θ_i and all (p_i^*, \mathbf{Q}_i^*) in the support of supplier i 's equilibrium strategy,

$$(p_i^*, \mathbf{Q}_i^*) \in \arg \max_{p, \mathbf{Q}} E_{\theta_{-i}} [x_i((p, \mathbf{Q}), (p_{-i}^*, \mathbf{Q}_{-i}^*)) k_i(\theta_i) - s_i((p, \mathbf{Q}), (p_{-i}^*, \mathbf{Q}_{-i}^*))], \tag{6}$$

where the expression for supplier i 's expected profit derives from (5). In (6), suppliers' private information enters their objective function only through their pseudotypes. Thus, supplier i is indifferent among the strategies played by all the realizations of supplier i 's type with the same pseudotype.

⁸ The requirement of quasilinearity of the scoring rule is only needed when private information is multidimensional. When private information is one-dimensional there is a one-to-one mapping between types and pseudotypes. The equivalence classes of types with the same preferences are thus singletons.

⁹ Or, more generally, in the case of open formats, the strategies of the bidders.

We can construct a new equilibrium, $(\tilde{\mathcal{E}}_1, \dots, \tilde{\mathcal{E}}_N)$, such that:

1. $\tilde{\mathcal{E}}_i(\theta_i) = \tilde{\mathcal{E}}_i(\hat{\theta}_i)$ whenever $k(\theta_i) = k(\hat{\theta}_i)$.
2. Define $\Theta_i(k) = \{\theta_i \in \Theta_i \mid k(\theta_i) = k\}$. For each k in the support of bidder i 's pseudotypes, the distribution over (p, \mathbf{Q}) generated under $\tilde{\mathcal{E}}_i$ for a given $\theta_i \in \Theta_i(k)$ replicates the aggregate distribution over (p, \mathbf{Q}) over all $\theta_i \in \Theta_i(k)$ under \mathcal{E}_i .

By construction, the distribution of bidder i 's opponents' strategies is the same as before from bidder i 's perspective. Moreover, $\tilde{\mathcal{E}}_i$ is a best response for bidder i . Hence it is an equilibrium, and bidders' strategies are only a function of their pseudotypes. By construction, $(\tilde{\mathcal{E}}_1, \dots, \tilde{\mathcal{E}}_N)$ and $(\mathcal{E}_1, \dots, \mathcal{E}_N)$ lead to the same aggregate distribution of (p, \mathbf{Q}) and therefore scores and probabilities of winning. *Q.E.D.*

An aspect of Lemma 1 worth stressing is the role played by the assumption that types are independent across bidders. From the expression of suppliers' expected profit, $x_i k(\theta_i) - s_i$, we already know that their payoffs are only a function of their pseudotypes. Independence ensures that their beliefs are also independent of their types beyond their pseudotypes (actually, independence is stronger: it makes suppliers' beliefs independent of their types *and* pseudotypes). Without independence, bidders' private information would enter their expected payoff in (6), both through their pseudotypes and through their expectations over their opponents' types.

Lemma 1 implies that the set of possible outcomes (x_1, \dots, x_N) and (s_1, \dots, s_N) can be generated by equilibria where suppliers bid exclusively on the basis of their pseudotypes. However, it does not imply that nothing is lost by restricting attention to these equilibria. Outcome equivalence does not imply utility equivalence for the buyer. To see this, consider the following example.

Consider two equally likely types, θ_i and $\hat{\theta}_i$ (this assumption is inessential for the argument) such that $k(\theta_i) = k(\hat{\theta}_i)$ and suppose that, in equilibrium, they get a different outcome: (x_i, s_i) and (\hat{x}_i, \hat{s}_i) . By definition, these two types generate expected utility $f_i(\theta_i)s_i + f_i(\hat{\theta}_i)\hat{s}_i$ for the buyer, according to the scoring rule. However, this differs from true expected utility if $\phi(\cdot) \neq v(\cdot, t)$. To know how much expected utility the suppliers generate for the buyer, we need to know how they will satisfy their obligations. Let \mathbf{Q} and $\hat{\mathbf{Q}}$ be the choice of θ_i and $\hat{\theta}_i$, respectively (these are independent of s_i and \hat{s}_i). The total monetary transfer from the buyer to the suppliers is then given by $x_i \phi(\mathbf{Q}) - s_i$ and $\hat{x}_i \phi(\hat{\mathbf{Q}}) - \hat{s}_i$, and the buyer's true expected utility is given by

$$f_i(\theta_i)[x_i(v(\mathbf{Q}, t) - \phi(\mathbf{Q})) + s_i + \hat{x}_i(v(\hat{\mathbf{Q}}, t) - \phi(\hat{\mathbf{Q}})) + \hat{s}_i].$$

This equilibrium is outcome equivalent to an equilibrium where type θ_i adopts $\hat{\theta}_i$'s strategy and vice versa. In that equilibrium, the buyer's true expected utility is given by

$$f_i(\theta_i)[\hat{x}_i(v(\mathbf{Q}, t) - \phi(\mathbf{Q})) + \hat{s}_i + x_i(v(\hat{\mathbf{Q}}, t) - \phi(\hat{\mathbf{Q}})) + s_i].$$

Clearly, the buyer is not indifferent between these two equilibria unless $x_i = \hat{x}_i$ or $v(\mathbf{Q}, t) = \phi(\mathbf{Q})$. The next result ensures that suppliers with the same pseudotypes receive the same equilibrium outcome function (x_i, s_i) . This rules out the situation described in the previous example.

Consider any equilibrium $(\mathcal{E}_1, \dots, \mathcal{E}_N)$. Define $\underline{x}_i(k) = \inf_{(p, \mathbf{Q}) \in \text{supp}_{\mathcal{E}_i(\theta_i)} \text{ for } \theta_i \in \Theta_i(k)} E_{\theta_i} [x_i((p, \mathbf{Q}), (p^*_{-i}, \mathbf{Q}^*_{-i}))]$. Similarly, define $\bar{x}_i(k) = \sup_{(p, \mathbf{Q}) \in \text{supp}_{\mathcal{E}_i(\theta_i)} \text{ for } \theta_i \in \Theta_i(k)} E_{\theta_i} [x_i((p, \mathbf{Q}), (p^*_{-i}, \mathbf{Q}^*_{-i}))]$. Let $\underline{s}_i(k), \bar{s}_i(k)$ be the resulting scores to fulfill. In words, $\underline{x}_i(k)$ is the lowest expected probability of winning the contract among all the bids in the support of bidder i 's strategy when he has pseudotype k . Similarly, $\bar{x}_i(k)$ is his highest expected probability of winning.

Lemma 2. In any equilibrium, $\underline{x}_i(k) = \bar{x}_i(k)$ and $\underline{s}_i(k) = \bar{s}_i(k)$ for all k except possibly on a set of measure zero.

Proof. Define $U_i(k)$ as supplier i 's equilibrium expected payoff when he has pseudotype k . Incentive compatibility (IC) implies that^{10, 11}

$$\begin{aligned} U_i(k) &= \underline{x}_i(k)k - \underline{s}_i(k) \geq \underline{x}_i(\widehat{k})k - \underline{s}_i(\widehat{k}) = U_i(\widehat{k}) + \underline{x}_i(\widehat{k})(k - \widehat{k}) \\ U_i(\widehat{k}) &= \underline{x}_i(\widehat{k})\widehat{k} - \underline{s}_i(\widehat{k}) \geq \underline{x}_i(k)\widehat{k} - \underline{s}_i(k) = U_i(k) - \underline{x}_i(k)(k - \widehat{k}). \end{aligned}$$

Hence $\underline{x}_i(k)$ is monotonically increasing in k . The same argument applies to $\bar{x}_i(k)$. Hence $\underline{x}_i(k)$ and $\bar{x}_i(k)$ are almost everywhere continuous. A similar argument based on the IC constraint establishes that $\underline{x}_i(k) \geq \bar{x}_i(\widehat{k})$ for all $\widehat{k} < k$. Together with the a.e. continuity of these functions, this implies that $\underline{x}_i(k) = \bar{x}_i(k)$ (and $\underline{s}_i(k) = \bar{s}_i(k)$) almost everywhere. *Q.E.D.*

Define two equilibria as typewise outcome equivalent if they generate the same distribution of outcomes (x_1, \dots, x_N) and (s_1, \dots, s_N) , conditional on types in $\Theta_1 \times \dots \times \Theta_N$. We are now able to prove the main result of this section.

Theorem 1. Every equilibrium in the scoring auction is typewise outcome equivalent to an equilibrium in the scoring auction where suppliers are constrained to bid only on the basis of their pseudotypes, and vice versa.

Proof. All equilibria in the constrained auction are also equilibria in the scoring auction because bidders' preferences and beliefs are entirely determined by their pseudotypes. Lemma 2 implies that all types with the same pseudotype get the same x and s a.e. in all equilibria. *Q.E.D.*

Theorem 1 ensures that there is no loss of generality in concentrating on pseudotypes when deriving the equilibrium in the scoring auction, even if the scoring rule does not correspond to the buyer's true preference. (Note that Theorem 1 does not rule out equilibria where different types submit different (p, \mathbf{Q}) bids—but given that they yield the same score and the same probability of winning at equilibrium, they are payoff irrelevant.)

Most theoretical analyses of scoring auctions have implicitly or explicitly taken advantage of pseudotypes to derive an equilibrium in these auctions (Che, 1993; Bushnell and Oren, 1994, 1995). Theorem 1 suggests that doing so does not discard any other equilibria of interest. Although this might not be totally surprising when types are one-dimensional, this result is not trivial for environments where types are multidimensional. This property is a consequence of the combination of the quasilinear scoring rule, the single dimensionality of the allocation decision, and the independence of types across bidders. We cannot reduce the strategic environment to one-dimensional private information if any of these conditions does not hold. As argued in Section 1, the quasilinearity of the scoring rule is necessary to be able to summarize suppliers' preferences over contracts by a single number. As noted after Lemma 1, independence was needed to make suppliers' beliefs independent of their types. Neither condition is necessary to use pseudotypes to derive an equilibrium in the one-dimensional model (for example, Branco, 1997 extends Che's model to correlated private information).

The next result makes the relationship between scoring auctions and standard one-object auctions even more explicit.

Corollary 1. The equilibrium in quasilinear scoring auctions with independent types inherits the properties of the equilibrium in the related single-object auction where (i) bidders are risk neutral, (ii) their (private) valuations for the object correspond to the pseudotype k in the original scoring auction and are distributed accordingly, (iii) the highest bidder wins, and (iv) the payment rule is determined as in the scoring auction, with bidders' scores being replaced by bidders' bids.

Corollary 1 has practical implications for the derivation of the equilibrium in scoring auctions. It suggests the following simple algorithm for deriving equilibria in scoring auctions: (1) given the scoring rule, derive the distribution of pseudotypes, $G_i(k)$, (2) solve for the equilibrium in the related IPV auction where valuations are distributed according to $G_i(k)$, $b_i(k)$, and (3) the

¹⁰ Note that when $\underline{x}_i(k) < \bar{x}_i(k)$, different types with the same pseudotype use different equilibrium strategies or they use the same mixed strategy.

¹¹ Without loss of generality, we focus on equilibria that involve optimal strategies for all realizations of types.

equilibrium bid in the scoring auction is any (p, \mathbf{Q}) such that $S(p, \mathbf{Q}) = b_i(k)$. (The actual (p, \mathbf{Q}) delivered are easy to derive given $b_i(k)$ and the solution to equation (1).)

4. Expected utility equivalence across auction formats

■ In this section, we extend the revenue equivalence theorem (Myerson, 1981; Riley and Samuelson, 1981) to multi-attribute environments. Che (1993) proved the utility equivalence between the first- and second-score scoring auction when types are one-dimensional and the scoring rule corresponds to the buyer’s true preference. Theorem 2 shows that this result extends to multidimensional private information and scoring rules that do not correspond to the buyer’s true preference.

Theorem 2 (Expected utility equivalence). Any two scoring auctions that:

- (i) use the same quasilinear scoring rule,
- (ii) use the same allocation rule $x_i(k_i, k_{-i}), i = 1, \dots, N$, and
- (iii) yield the same expected payoff for the lowest pseudotype $k_i, i = 1, \dots, N$,

generate the same expected utility for the buyer

Proof. Because the buyer’s utility is quasilinear, his expected utility from a given auction is

$$\sum_{i=1}^N E_{k_i, k_{-i}} [x_i(k_i, k_{-i}) ESS(k_i) - U_i(k_i)] = \sum_{i=1}^N E_{k_i} [x_i(k_i) ESS(k_i) - U_i(k_i)], \quad (7)$$

where $ESS(k_i)$ is the expected social surplus generated by awarding the contract to bidder i with pseudotype k_i .

By Theorem 1, we can focus on equilibria which are only functions of pseudotypes. Incentive compatibility implies that $U_i(k_i)$ is almost everywhere differentiable and that $\frac{d}{dk_i} U_i(k_i) = x_i(k_i)$, where $x_i(k_i)$ is a well-defined function almost everywhere by Lemma 2. Hence, (ii) and (iii) imply that $U_i(k)$ is the same across both auctions.

Next, fix k_i and let $(p^*(\theta_i, s_i), \mathbf{Q}^*(\theta_i, s_i))$ be the realized contract of supplier i with type $\theta_i \in \Theta_i(k_i)$, when the score to satisfy is s_i . Because the scoring rule is quasilinear, $\mathbf{Q}^*(\theta_i, s_i)$ is only a function of the scoring rule and θ_i , and not of s_i (cf (1)). Hence,

$$ESS(k_i) = E_{\theta_i \in \Theta_i(k_i)} [v(\mathbf{Q}^*, t) - c(\mathbf{Q}^*, \theta_i)]$$

is independent of s_i and equal across the two auctions given (i). The claim follows. *Q.E.D.*

Three points are worth noting concerning this result. First, the assumption that the scoring rule is quasilinear is key. Without it, suppliers’ choice of product characteristics (p, \mathbf{Q}) would depend on the form of the resulting obligation, that is, the auction format.

Second, the proof of Theorem 2 relies on the fact that any equilibrium is essentially pure as a function of pseudotypes (i.e., x_i are functions). Without this property, expected utility equivalence between two auctions that yield the same distribution of allocations as a function of pseudotypes would only hold when the scoring rule corresponds to the true valuation (cf the argument before Lemma 2). In that case, $ESS(k_i) = k_i$ and the result holds trivially.

Third, Theorem 2 implies the standard equivalence between the first-score auction, the second-score auction, and the Dutch and English auctions when bidders are symmetric. However, the symmetry requirement is with respect to the distribution of pseudotypes and not the distribution of types. In particular, some bidders can (stochastically) be stronger for one attribute and others for another attribute, yet, when it comes to pseudotypes, they can be symmetric.

5. Comparison with alternatives

■ In this section, we consider three alternatives to scoring auctions under three different auction formats, and for simplicity we focus on the case where suppliers are *ex ante* symmetric. We show

that, except for the first-price menu auction, these alternatives generate equal or lower expected utility for the buyer than a scoring auction that uses the true preference of the buyer. Thus, *a fortiori*, a scoring auction with an optimally chosen scoring rule dominates these alternatives. We next describe these procedures in detail and discuss some of their properties.

*Menu auctions.*¹² In the menu auction, the buyer does not reveal his type. Instead, suppliers are asked to submit (p, \mathbf{Q}) schedules. The buyer selects the offer that generates the highest level of utility. This alternative comes in three versions. In the “ascending” version (A), the auction takes place over several rounds. In each round, the buyer selects the supplier whose schedule generates the greatest utility. In the next round, the other suppliers are invited to submit new schedules and the process continues until no further offers are made. The winner is the supplier who offers the best schedule in the last round. The resulting contract is the (p, \mathbf{Q}) in his schedule that the buyer prefers. In the “first-price” version (FP), the winner is the supplier offering the (p, \mathbf{Q}) contract that generates the highest utility to the buyer and this is the resulting contract. Finally, in the “second-price” version (SP), the winner is the supplier offering the (p, \mathbf{Q}) contract that generates the highest utility to the buyer and the resulting contract is (\hat{p}, \mathbf{Q}) , where \hat{p} is adjusted so that (\hat{p}, \mathbf{Q}) generates the same score as the best offer by the losers.¹³

Menu auctions introduce an interesting new twist: suppliers must now account for the fact that the buyer selects the offer he prefers in the submitted menus. Let $(\mathbf{Q}(t, \theta), p(t, \theta))_{t \in T}$ denote the menu submitted by a supplier of type θ , with the indexing such that $(\mathbf{Q}(t, \theta), p(t, \theta))$ is the offer preferred by the buyer with taste t . Incentive compatibility for the buyer requires that

$$v(\mathbf{Q}(t, \theta), t) - p(t, \theta) \geq v(\mathbf{Q}(\hat{t}, \theta), t) - p(\hat{t}, \theta) \quad \forall t, \hat{t} \in T, \quad (8)$$

that is, using standard arguments and the fact that $v_{Q_t} > 0$ (v_{Q_t} ensures that \mathbf{Q} is monotonic and thus a.e. differentiable),

$$v_{Q_t}(\mathbf{Q}(t, \theta), t) Q_t(t, \theta) = p_t(t, \theta) \text{ for all } \theta \text{ and all } t \text{ at which } \mathbf{Q} \text{ is differentiable.} \quad (9)$$

Lemma 3. Consider any incentive-compatible menu $(\mathbf{Q}(t, \theta), p(t, \theta))_{t \in T}$. This menu induces efficient production for all t, θ if and only if (i) it corresponds to an *ex post* iso-profit curve of supplier with type θ and (ii) $\mathbf{Q}(t, \theta)$ is a.e. differentiable with $Q_t(t, \theta) > 0$ a.e.

Proof. Efficient production requires

$$v_{Q_t}(\mathbf{Q}(t, \theta), t) = c_{Q_t}(\mathbf{Q}(t, \theta), \theta) \text{ for all } t, \theta. \quad (10)$$

Condition (ii) follows directly from the requirement of efficiency together with the assumption that $v_{Q_t} > 0$. Suppliers’ *ex post* iso-profit curves are described by the locus of offers such that $p(t, \theta) - c(\mathbf{Q}(t, \theta), \theta)$ is constant, that is,

$$c_{Q_t}(\mathbf{Q}(t, \theta), \theta) Q_t(t, \theta) = p_t(t, \theta) \text{ for all } \theta \text{ and all } t \text{ at which } \mathbf{Q} \text{ is differentiable.} \quad (11)$$

Because $Q_t(t, \theta) \neq 0$, conditions (10) and (11) are equivalent given (9) (continuity of the efficient production level in t ensures that (10) is satisfied at nondifferentiability points of \mathbf{Q}). *Q.E.D.*

Beauty contests. In a beauty contest, the buyer does not reveal his type and the suppliers are asked to submit a single bid (p, \mathbf{Q}) . It potentially comes in three forms: the ascending format, and the first-price and the second-price auctions. These formats are self-explanatory given their description for the menu auction.

¹² Bichler and Kalagnanam (2003) use the expression “auctions with configurable offers” to describe such procedures.

¹³ Note that commitment will again be essential here. The buyer must be able to convince suppliers that he will not manipulate the (unannounced) scoring rule in order to increase the value of the second-best offer.

Price-only auctions. In price-only auctions, the buyer publishes a detailed request-for-quote that sets minimum levels for each attribute. All offers satisfying these conditions are evaluated on a price basis. Again, it comes in three guises: ascending, first-price and second-price.

We now compare the performance of these alternative procedures with the performance of a scoring auction that uses the true preference of the buyer as its scoring rule. Let $U_l^k(t)$ be the expected utility of the buyer with taste t , in format $k \in \{A, FP, SP\}$ and procedure $l \in \{\text{score}, \text{menu}, \text{beauty}\}$ (where “score” stands for a scoring auction of the type described in the previous sections with the scoring rule corresponding to the true preference of the buyer).

Theorem 3. For all t ,

- (i) $U_{\text{score}}^A(t) = U_{\text{menu}}^A(t) = U_{\text{beauty}}^A(t)$ as the bidding increment goes to zero,
- (ii) $U_{\text{score}}^{SP}(t) = U_{\text{menu}}^{SP}(t) > U_{\text{beauty}}^{SP}(t)$.

Proof. See the Appendix.

For the ascending auction, we build the unique symmetric equilibrium for each procedure. The equivalence between all three procedures then stems from the direct comparison of the final allocations. For the second-price format, we show that submitting schedule $\mathbb{S} = \{(p, \mathbf{Q}) \text{ such that } p = c(\mathbf{Q}, \theta), \mathbf{Q} \in \mathbb{R}^M\}$ is the unique dominant strategy equilibrium in the menu auction. The equivalence between the menu auction and the scoring auction with $\phi(\mathbf{Q}) = v(\mathbf{Q}, t)$ follows directly. For the beauty contest, we argue that the equilibrium bid (p^*, \mathbf{Q}^*) must belong to \mathbb{S} . Because there is a positive probability that (p^*, \mathbf{Q}^*) does not belong to $\text{argmax}_{(p, \mathbf{Q}) \in \mathbb{S}} \{v(\mathbf{Q}, t) - c(\mathbf{Q}, \theta)\}$ for the actual type—unknown to the suppliers—of the buyer, $U_{\text{menu}}^{SP} > U_{\text{beauty}}^{SP}$ follows.

Theorem 3 understates the superiority of scoring auctions in two ways. First, scoring auctions are likely to dominate both procedures because they save on bidding costs for suppliers. In practice, the existence of bidding costs will limit the number of offers made in a menu auction. This favors the scoring auction. Likewise, the equivalence result for the beauty contest in the ascending format requires that suppliers submit a very high number of bids. Second, the comparison in Theorem 3 is with a scoring auction with scoring rule $\phi(\mathbf{Q}) = v(\mathbf{Q}, t)$. As suggested by Che (1993) and Asker and Cantillon (2006), the buyer will in general be better off announcing $\phi(\mathbf{Q}) \neq v(\mathbf{Q}, t)$.

We next consider the first-price menu auction. We first show the following general result.

Theorem 4. Any equilibrium of the first-price menu auction is inefficient.

Proof. See the Appendix.

Theorem 4 follows from the following observations. If the equilibrium in the menu auction involves pooling (suppliers make the same offer to different buyer types), inefficiency is immediate. If, instead, full separation occurs at equilibrium, inefficiency arises from the tension between the requirements of incentive compatibility and those of profit maximization. For the purpose of profit maximization alone, suppliers are tempted to target different profit levels according to the buyer’s type. The buyer’s incentive compatibility constraint limits the ability of suppliers to do this. We argue that the buyer’s incentive compatibility constraint binds generically in any separating equilibrium of the first-price menu auction and show that qualities are distorted as a result.¹⁴

The inefficiency of the first-price menu auction is not necessarily bad news for the buyer if it induces fiercer competition. To investigate this question further, we focus on the more restricted environment where private information is one-dimensional and the buyer can only have two

¹⁴ We find that there is no quality distortion at the top and at the bottom in the separating equilibrium of the first-price menu auction, a result that mirrors Rochet and Stole (2002). Rochet and Stole develop the intuition for this finding. They also reconcile it with the discrete type case in which there is distortion at the bottom (see Theorem 5).

types.^{15, 16} The following theorem suggests that buyers with different types are likely to rank the two procedures differently.

*Theorem 5 (Adapted from Theorem 2 of Biglaiser and Mezzetti, 2000).*¹⁷ Suppose $Q, \theta \in \mathbb{R}, t \in \{t_L, t_H\}$ with $t_L < t_H$. Suppose $v_Q(Q, t_H) > v_Q(Q, t_L)$ and $v_{QQ}(Q, t_H) \leq v_{QQ}(Q, t_L)$ for all Q and that $c_{Q\theta} > 0$. Then, in the symmetric equilibrium of the menu auction, $U_{\text{score}}^{FP}(t_L) \geq U_{\text{menu}}^{FP}(t_L)$ and $U_{\text{score}}^{FP}(t_H) \leq U_{\text{menu}}^{FP}(t_H)$.

The proof of Theorem 5 closely follows that in Biglaiser and Mezzetti. A sketch is provided in the Appendix pointing out how to adapt their arguments.

At equilibrium, suppliers offer two contracts, one targeted at the low-type buyer, (Q_L, p_L) , and the other targeted at the high-type buyer, (Q_H, p_H) . The inequalities in Theorem 5 are strict whenever one of the incentive compatibility constraints (8), evaluated at the equilibrium offers of the scoring auction, binds. Biglaiser and Mezzetti (2000) argue that this will be the case when t_L and t_H are sufficiently close. Intuitively, when t_L and t_H are sufficiently distinct, the two contracts offered by each supplier are sufficiently different that the low-type buyer is not tempted by the high-type contract and vice versa. In that case, the bidding equilibrium in the truthful scoring auction describes the equilibrium contracts for each type in the menu auction.¹⁸ When the two buyer types are sufficiently close, the incentive compatibility constraints bind. Following the intuition from the single-principal single-agent case, the price and quality of the low-type buyer is distorted downward to satisfy the incentive compatibility constraint of the high-type buyer. However, competition means that the participation constraint of the buyer is now endogenous from the point of view of an individual supplier: it depends on the bids of the other suppliers. This increases the costs of distorting the low-type contract relative to the single-principal single-agent benchmark. As a result, suppliers also use the price offered to the high-type buyer to help ensure his incentive compatibility constraint is satisfied. Thus, the price offered to the high-type buyer decreases relative to the scoring auction, and the high-type buyer is better off.¹⁹

Biglaiser and Mezzetti point out that if the buyer knows his type prior to choosing a procedure, unraveling of the buyer's private information is likely: the low-type buyer chooses a scoring auction, leaving the high-type as the only type to choose the menu auction. Because he no longer has any private information, the menu and the scoring auction become equivalent again.

We now turn to the procedure where the buyer sets minimum quality standards and awards the contract on the basis of price only.

Theorem 6. Consider any standard auction format where the high bidder wins and its equivalent in the scoring auction. A buyer is always better off using a scoring auction with a scoring rule that corresponds to his true taste than imposing minimum quality standards/attribute levels and selecting the winner on the basis of price only.

Proof. Suppose the buyer sets minimum quality standards $\mathbf{Q} = \underline{\mathbf{Q}} \in \mathbb{R}^M$. Because costs are increasing, suppliers will set their quality levels at $\underline{\mathbf{Q}}$. We are now back to a standard procurement auction with symmetric bidders and costs $c(\underline{\mathbf{Q}}, \theta_i) \in \mathbb{R}$. Let $x^{(n:N)}$ denote the n th highest-order statistics from N draws of random variable x . From the revenue equivalence theorem, the expected

¹⁵ When the equilibrium in the menu auction is known to be a separating equilibrium, then it is easy to show, using the techniques developed in the proof of Theorem 4, that some buyer types prefer the menu auction whereas others prefer the scoring auction.

¹⁶ Rochet and Stole (2002) note that pooling is a common feature of equilibrium in this class of models. How pooling affects the welfare of different buyers in the menu auction is unclear without a full characterization of the equilibrium.

¹⁷ Biglaiser and Mezzetti consider a model in which multiple principals bid for the exclusive service of an agent. Each principal has private information about their valuation of the service, while the agent has private information about the disutility of providing the service. The first-price menu auction considered here is the procurement version of this model (with the agent being the buyer and the principals being the suppliers).

¹⁸ Thus the outcome is efficient. This does not contradict Theorem 4 because Theorem 4 applies to the case where the buyer has a continuum of types.

¹⁹ The quality offered to the high-type buyer remains at the first best.

utility of the buyer from this minimum quality standard auction is

$$\begin{aligned} & v(\underline{\mathbf{Q}}) - E_{\theta_i} [c(\underline{\mathbf{Q}}, \theta_i)^{(N-1:N)}] \\ &= E_{\theta_i} [v(\underline{\mathbf{Q}}) - c(\underline{\mathbf{Q}}, \theta_i)]^{(2:N)} \\ &< E_{\theta_i} [\max_{\mathbf{Q}} \{v(\mathbf{Q}) - c(\mathbf{Q}, \theta_i)\}]^{(2:N)} \end{aligned}$$

= expected utility from the truthful scoring auction by Theorem 2. *Q.E.D.*

6. Concluding remarks

■ Our article provides a systematic analysis of equilibrium behavior in scoring auctions when private information is multidimensional. We have characterized the set of equilibria in scoring auctions and have argued that a single number, the supplier's pseudotype, is sufficient to describe the equilibrium outcome in these auctions, when the scoring rule is quasilinear and types are independently distributed. In doing so, we have drawn on the equivalence between the reduced form of a scoring auction and that of a standard single-object IPV auction. We have also derived a new expected utility equivalence theorem for scoring auctions. Both results extend existing theories of scoring auctions.

In addition, we have shown that several other candidate procedures for buying differentiated products, including some, such as the menu auction and the beauty contest, that also combine competition with the flexibility of deciding on all the dimensions of the product, are dominated by scoring auctions. These results suggest that scoring auctions provide a useful mechanism (they are simple straightforward procedures) for buying differentiated products.

We conclude with a few remarks on potential venues for further research.

Suppliers' uncertainty about their costs at the time of bidding. In our model, it was immaterial whether bidders were committed to their offer or to the scores that their offer generated. Suppose now that the cost of attribute \mathbf{Q} is given by $c(\mathbf{Q}, \theta, \tau)$, where only θ is known to the supplier at the time of bidding and τ is known before the contract is executed. Define $k(\theta) = E_{\tau} [\max_{\mathbf{Q}} \{\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta, \tau)\}]$. Our equilibrium characterization results go through with this redefined pseudotype. The only difference is that the delivered quality level now generically differs from the offered quality level because the delivered quality will solve $\max_{\mathbf{Q}} \{\phi(\mathbf{Q}) - c(\mathbf{Q}, \theta, \tau)\}$ for the realization of τ . This provides a rationale for making the scores, rather than the actual offers, binding. Thus, low-cost realizations generate higher levels of quality and higher prices for the supplier, whereas negative cost shocks generate lower qualities and lower prices.²⁰ This added flexibility is another advantage of the scoring auction relative to the other procedures that set the quality to be delivered at the contracting stage.

Noncontractible quality dimensions. An essential assumption for all our results is that quality is contractible. When some dimensions of the good are contractible and others not, contracting can generate perverse incentives, as Holmstrom and Milgrom (1991) have shown. At the same time, it seems desirable to generalize the analysis of procurement mechanisms to such environments (see Che, forthcoming for a discussion of possible mechanisms).

Implications for empirical work. Even in the presence of symmetric suppliers, scoring auctions present interesting auction design questions (e.g., how can the buyer manipulate the scoring rule to his advantage?). However, scoring auctions present two difficulties from the point of view of identification: the identification of the functional form for the costs and the identification of the distribution of private information. One consequence of our sufficient statistics result is that the distribution of types will generally be nonidentified on the basis of auction data, even when the functional form for the costs is known. Indeed, the observed information (the scores) is one-dimensional, whereas the information to be inferred is multidimensional. This suggests two

²⁰ The U.S. highway procurement authorities use such a reward/penalty scheme.

possible solutions. When the (p, \mathbf{Q}) offers rather than the scores are binding, the observed data are again multidimensional.²¹ Another possibility is to look at auction data where changes in the scoring rule can be exploited. In any case, our article provides a theoretical basis from which investigation of identification is feasible.

Appendix

■ This Appendix contains proofs of Theorems 3 and 4 and a sketch of the proof of Theorem 5.

Proof of Theorem 3, part (i). We discretize the price grid. Let Δ be the minimum price (and therefore profit and utility) increment. We proceed in two steps, comparing first the menu auction and then the beauty contest to the scoring auction.

Step 1: Menu auction. The following is an equilibrium. In round 1, each bidder submits a schedule that generates at most zero utility for the buyer, that is, $\{(p, \mathbf{Q}); \mathbf{Q} \in \mathbb{R}^M, p - c(\mathbf{Q}, \theta) = \text{constant and } \max_t \max_{(p, \mathbf{Q})} \{v(\mathbf{Q}, t) - p\} = 0\}$. Let π_t be the profit level corresponding to the period t schedule for a given supplier. At round t , this supplier submits schedule $\{(p, \mathbf{Q}); \mathbf{Q} \in \mathbb{R}^M, p - c(\mathbf{Q}, \theta) = \pi_{t-1} - \Delta\}$ as long as $\pi_{t-1} - \Delta \geq 0$ if his offer was not selected in round $t - 1$.²² This is an equilibrium (we can adapt the arguments in Bikhchandani, Haile, and Riley, 2002 to argue that the outcome of this strategy is the unique equilibrium outcome in symmetric strategies). Each supplier participates as long as a positive profit can be made, otherwise they exit. The selected level of attributes, \mathbf{Q}^* , satisfies $\mathbf{Q}^* = \arg \max_{\mathbf{Q}} \{v(\mathbf{Q}, t) - c(\mathbf{Q}, \theta)\}$ for the realization of t . The final price satisfies $p = v(\mathbf{Q}^*, t) - \max_{\mathbf{Q}} \{v(\mathbf{Q}, t) - c(\mathbf{Q}, \tilde{\theta})\}$ (modulo the increment), where $\tilde{\theta}$ refers to the cost function of the second-best supplier. This is the outcome of the scoring auction.

Step 2: Beauty contest. The following is an equilibrium. In round 1, each bidder submits a bid in the schedule that generates at most zero utility for the buyer, that is, $\{(p, \mathbf{Q}); \mathbf{Q} \in \mathbb{R}^M, p - c(\mathbf{Q}, \theta) = \text{constant and } \max_t \max_{(p, \mathbf{Q})} \{v(\mathbf{Q}, t) - p\} = 0\}$. Let π_t be the profit level corresponding to the bid in period t for a given supplier. At round t , if this supplier was not the winner in round $t - 1$, he submits any bid in schedule $\{(p, \mathbf{Q}); \mathbf{Q} \in \mathbb{R}^M, p - c(\mathbf{Q}, \theta) = \pi_{t-1}\}$ that he has not submitted in the past. If no unsubmitted bid remains in this schedule, the supplier submits a bid in $\{(p, \mathbf{Q}); \mathbf{Q} \in \mathbb{R}^M, p - c(\mathbf{Q}, \theta) = \pi_{t-1} - \Delta\}$ as long as $\pi_{t-1} - \Delta \geq 0$. The process continues until no further bid is received. As before, the equilibrium strategies yield the unique equilibrium outcome. The winner in the beauty contest is the same as in the menu auction. However, the buyer might not be equally well off as in the menu auction, because here he cannot choose the (p, \mathbf{Q}) pair that maximizes his utility. However, as Δ goes to zero, the winning (p, \mathbf{Q}) must maximize $v(\mathbf{Q}, t) - c(\mathbf{Q}, \theta)$. Otherwise, the winning bidder could have won with a higher level of profit. This is ruled out by his bidding behavior. Thus, the buyer is equally well off.

Proof of Theorem 3, part (ii). We proceed in two steps, comparing first the menu auction and then the beauty contest to the scoring auction.

Step 1: Menu auction. An argument mirroring the argument for the standard second-price auction establishes that $\{(p, \mathbf{Q}); p = c(\mathbf{Q}, \theta), \mathbf{Q} \in \mathbb{R}^M\}$ is a dominant strategy equilibrium. Equivalence with the scoring auction follows from the fact that suppliers submit bids such that $p = c(\mathbf{Q}, \theta)$ in the dominant strategy equilibrium of the scoring auction, with $\mathbf{Q} = \arg \max \{v(\mathbf{Q}, t) - c(\mathbf{Q}, \theta)\}$. This is also the bid selected by the buyer in the winning schedule. The best second offers in both auctions are also identical.

Step 2: Beauty contest. In equilibrium, suppliers submit an equilibrium bid (p^*, \mathbf{Q}^*) in the schedule $\{(p, \mathbf{Q}); p = c(\mathbf{Q}, \theta), \mathbf{Q} \in \mathbb{R}^M\}$. Consider any alternative bid (p, \mathbf{Q}) such that $p - c(\mathbf{Q}, \theta) > 0$. The expected profit generated by this bid is equal to

$$\Pr((p, \mathbf{Q}) \text{ generates the highest score})(E p - c(\mathbf{Q}, \theta)),$$

where $E p$ is the expected resulting price determined by the second-best offer. The deviation $(\hat{p}, c(\mathbf{Q}, \theta))$, where $\hat{p} = c(\mathbf{Q}, \theta)$ dominates. The expected profit it generates is equal to

$$\Pr((p, \mathbf{Q}) \text{ generates the highest score})(E p - c(\mathbf{Q}, \theta)) + \Pr((\hat{p}, \mathbf{Q}) \text{ generates the highest score but } (p, \mathbf{Q}) \text{ does not})(E \hat{p} - c(\mathbf{Q}, \theta)),$$

where $E \hat{p}$ is the expected price given that (\hat{p}, \mathbf{Q}) generates the highest score but (p, \mathbf{Q}) does not. Clearly, $E \hat{p} - c(\mathbf{Q}, \theta) > 0$. (p, \mathbf{Q}) such that $p - c(\mathbf{Q}, \theta) < 0$ is similarly dominated. From the buyer's point of view, the utility from $(p^*, \mathbf{Q}^*) \in \{(p, c(\mathbf{Q}, \theta)); \mathbf{Q} \in \mathbb{R}^M, p = c(\mathbf{Q}, \theta)\}$ is lower than $\max\{v(\mathbf{Q}, t) - p; p = c(\mathbf{Q}, \theta), \mathbf{Q} \in \mathbb{R}^M\}$ with strictly positive probability. *Q.E.D.*

²¹ Bajari, Houghton, and Tadelis (2006) analyze data from highway paving procurement that have this feature. The approach they use involves a theoretical structure that employs aspects of the approach developed here.

²² Note that it is a best response for the buyer to select the offer he truthfully prefers at each round. Note also that Lemma 3 ensures that the menu induces productive efficiency.

Proof of Theorem 4. If the menu auction equilibrium involves pooling of offers, then inefficiency is immediate. Thus, we direct attention to equilibria where full separation occurs (a full menu is offered by each supplier).

Consider the optimization problem faced by a supplier of type θ . Let $U(t, \theta)$ denote the utility received by a buyer of type t from a supplier of type θ in the equilibrium of the menu auction. The buyer's incentive compatibility constraint can be rewritten as

$$\frac{d}{dt}U(t, \theta) = v_t(\mathbf{Q}(t, \theta), t) \tag{A1}$$

(with second-order condition $\nabla_t \mathbf{Q}(t, \theta) \geq 0$). Let $\Pr(U, t)$ denote the equilibrium probability that an offer generating a level of utility U for a buyer of type t wins. Ignoring the second-order condition, the supplier's problem can be written as²³

$$\max_{\mathbf{Q}(\cdot, \theta), U(\cdot, \theta)} \int_{\underline{t}}^{\bar{t}} (v(\mathbf{Q}(t, \theta), t) - c(\mathbf{Q}(t, \theta), \theta) - U(t, \theta)) \Pr(U(t, \theta), t) h(t) dt.$$

subject to (A1)

This is a standard optimal control problem, with U being the state variable and \mathbf{Q} being the control variables. The Hamiltonian is given by

$$H(\mathbf{Q}, \lambda, U, t) = (v(\mathbf{Q}(t, \theta), t) - c(\mathbf{Q}(t, \theta), \theta) - U(t, \theta)) \Pr(U(t, \theta), t) h(t) + \lambda(t, \theta) v_t(\mathbf{Q}(t, \theta), t).$$

If a solution exists, it must satisfy the following first-order conditions (where we have dropped the arguments for simplicity):

$$(v_{\mathbf{Q}}(\mathbf{Q}, t) - c_{\mathbf{Q}}(\mathbf{Q}, \theta)) \Pr(U, t) h(t) + \lambda(t, \theta) v_{\mathbf{Q}_t}(\mathbf{Q}, t) = 0 \tag{A2}$$

$$-(v - c - U) \frac{d}{dU} \Pr(U, t) h(t) + \Pr(U, t) h(t) = \lambda' \tag{A3}$$

with $\lambda(\bar{t}, \theta) = \lambda(\underline{t}, \theta) = 0$.

To complete the proof, we now argue that $\lambda(t, \theta) \neq 0$ for a positive measure of t , that is, the buyer's incentive constraint binds. Toward a contradiction, suppose it does not. Then the equilibrium in the menu auction corresponds to the equilibrium in the scoring auction as t varies. Denote by $(\mathbf{Q}^s(t, \theta), p^s(t, \theta))_{t \in T}$ the menu generated from the equilibrium offers submitted by a supplier of type θ in the scoring auction as t varies. From Corollary 1 and the known characterization of the equilibrium in the single-object IPV auction, the supplier's *ex post* profit in the scoring auction is given by

$$k(t, \theta) - E_{\tilde{\theta}}[k(t, \tilde{\theta})^{(1:N-1)} \mid k(t, \tilde{\theta})^{(1:N-1)} < k(t, \theta)],$$

where $k(t, \theta)$ is the pseudotype that corresponds to type θ when the scoring rule is equal to $v(\mathbf{Q}, t) - p$, and $k(t, \tilde{\theta})^{(1:N-1)}$ denotes the first-order statistics of $N - 1$ independent draws of pseudotypes. By the envelope theorem, $k_t(t, \theta) = v_t(\mathbf{Q}(t, \theta), t)$. It is independent of the distribution of θ . By contrast, the second term and its derivative with respect to t is a function of the distribution of θ . Thus the two terms do not cancel out and *ex post* profits depend on t for a given θ (except if supplier θ wins with zero probability for all t). This implies—by Lemma 3—that the menu that corresponds to the equilibrium in the scoring auction is not incentive compatible, a contradiction.

When $\lambda(t, \theta) \neq 0$, equation (A2) implies that qualities are distorted away from their efficient levels. *Q.E.D.*²⁴

Sketch of proof of Theorem 5. Here we show how to adapt the arguments in Biglaiser and Mezzetti (2000) to fit our procurement framework.²⁵

Let $U_j(\theta)$ be the utility of a type j buyer from the contract supplier θ offers to him. Let $\theta^{-1}(U_j)$ be the inverse of $U_j(\theta)$. Let $\Pr(t_L) = p$. The bidding problem faced by each supplier is given by

$$\max_{Q_L, Q_H, U_L, U_H} p [v(Q_L, t_L) - c(Q_L, \theta) - U_L] (1 - G(\theta^{-1}(U_L)))^{N-1} + (1 - p) [v(Q_H, t_H) - c(Q_H, \theta) - U_H] (1 - G(\theta^{-1}(U_H)))^{N-1}$$

subject to the following incentive compatibility constraints:

$$IC_H : v(Q_H, t_H) - p_H \geq v(Q_L, t_H) - p_L$$

$$IC_L : v(Q_L, t_L) - p_L \geq v(Q_H, t_L) - p_H.$$

²³ If the solution to this problem violates the second-order condition, pooling must occur at equilibrium.

²⁴ $\lambda(t, \theta)$ must follow the law of motion (A3) together with the boundary conditions $\lambda(\bar{t}, \theta) = \lambda(\underline{t}, \theta) = 0$. In particular, this implies no distortion at the top and at the bottom.

²⁵ For details, the reader is referred to Biglaiser and Mezzetti (2000).

Assuming that IC_L is slack (this is verified *ex post*) and rewriting IC_H as follows,

$$U_H \geq v(Q_L, t_H) - v(Q_L, t_L) + U_L$$

gives rise to the following Lagrangian for the bidders' optimization problem:

$$\begin{aligned} \mathbb{L}(\theta) = & p[v(Q_L, t_L) - c(Q_L, \theta) - U_L](1 - G(\theta^{-1}(U_L)))^{N-1} \\ & + (1 - p)[v(Q_H, t_H) - c(Q_H, \theta) - U_H](1 - G(\theta^{-1}(U_H)))^{N-1} \\ & + \lambda(\theta)(1 - G(\theta))^{N-1}[U_H - v(Q_L, t_H) + v(Q_L, t_L) - U_L]. \end{aligned}$$

Of interest to us are the first-order conditions with respect to U_L and U_H , which can be written as

$$\frac{\partial U_L}{\partial \theta} = \left[\frac{p(N-1)g(\theta)}{(\lambda(\theta) + p)(1 - G(\theta))} (v(Q_H, t_H) - c(Q_H, \theta) - U_H) \right] \tag{A4}$$

$$\frac{\partial U_H}{\partial \theta} = \left[\frac{(1-p)(N-1)g(\theta)}{(1 - (\lambda(\theta) + p))(1 - G(\theta))} (v(Q_L, t_L) - c(Q_L, \theta) - U_L) \right]. \tag{A5}$$

When $\lambda(\theta) = 0$, IC_H does not bind and the menu auction is equivalent to two independent scoring auctions. When $\lambda(\theta) > 0$, IC_H does bind. In this case, U_H increases and U_L decreases. To see this, note that (A4) implies U_H is increasing in $\lambda(\theta)$ and (A5) implies U_L is decreasing in $\lambda(\theta)$. It follows that $U_{\text{menu}}^{FP}(t_L) \leq U_{\text{scoring}}^{FP}(t_L)$ and $U_{\text{menu}}^{FP}(t_H) \geq U_{\text{scoring}}^{FP}(t_H)$. *Q.E.D.*

References

ARIZONA DEPARTMENT OF TRANSPORT. "A+B Bidding Guide." <http://www.azdot.gov/Highways/ConstGrp/PDF/ATBGuide.pdf>, 2002.

ASKER, J. AND CANTILLON, E. "Procurement When Both Price and Quality Matter." CEPR Discussion Paper no. 6082, 2006.

ATHEY, S. AND LEVIN, J. "Information and Competition in US Timber Auctions." *Journal of Political Economy*, Vol. 109 (2001), pp. 375–417.

BAJARI, P., HOUGHTON, S., AND TADELIS, S. "*Bidding for Incomplete Contracts*." Mimeo, University of Michigan, 2006.

BICHLER, M. AND KALAGNANAM, J. "Configurable Offers and Winner Determination in Multi-Attribute Auctions." *European Journal of Operational Research*, Vol. 160 (2003), pp. 380–394.

BIGLAISER, G. AND MEZZETTI, C. "Incentive Auctions and Information Revelation." *RAND Journal of Economics*, Vol. 31 (2000), pp. 145–164.

BIKHCHANDANI, S., HAILE, P., AND RILEY, J. "Symmetric Separating Equilibria in English Auctions." *Games and Economic Behavior*, Vol. 38 (2002), pp. 19–27.

BRANCO, F. "The Design of Multidimensional Auctions." *RAND Journal of Economics*, Vol. 28 (1997), pp. 63–81.

BUSHNELL, J. AND OREN, S. "Bidder Cost Revelation in Electric Power Auctions." *Journal of Regulatory Economics*, Vol. 6 (1994), pp. 5–26.

_____, _____, AND _____. "Internal Auctions for Efficient Sourcing of Intermediate Products." *Journal of Operations Management*, Vol. 12 (1995), pp. 311–320.

CHAO, H. P. AND WILSON, R. "Multi-Dimensional Procurement Auctions for Power Reserves: Robust Incentive-Compatible Scoring and Settlement Rules." *Journal of Regulatory Economics*, Vol. 22 (2002), pp. 161–183.

CHE, Y.K. "Design Competition through Multi-Dimensional Auctions." *RAND Journal of Economics*, Vol. 24 (1993), pp. 668–680.

_____, _____, _____. "Procurement." In *The New Palgrave Dictionary of Economics*, 2d ed., forthcoming.

_____, _____, AND GALE, I. "Standard Auctions with Financially Constrained Bidders." *Review of Economic Studies*, Vol. 65 (1998), pp. 1–21.

_____, _____, AND _____. "Revenue Comparison for Auctions When Bidders Have Arbitrary Types." *Theoretical Economics*, Vol. 1 (2006), pp. 95–118.

DE FRUTOS, M.A. AND PECHLIVANOS, L. "Second-Price Common-Value Auctions under Multidimensional Uncertainty." *Games and Economic Behavior*, Vol. 55 (2006), pp. 43–71.

EWERHART, C. AND FIESELER, K. "Procurement Auctions and Unit-Price Contracts." *RAND Journal of Economics*, Vol. 34 (2003), pp. 569–581.

FANG, G. AND PARREIRAS, S. "Equilibrium of Affiliated Value Second Price Auctions with Financially Constrained Bidders: The Two-Bidder Case." *Games and Economic Behavior*, Vol. 39 (2002), pp. 215–236.

HERBSMAN, Z., CHEN, W.T., AND EPSTEIN, W.C. "Time Is Money: Innovative Contracting Methods in Highway Construction." *Journal of Construction Engineering and Management*, Vol. 121 (1995), pp. 273–281.

HOLMSTROM, B. AND MILGROM, P. "Multi-Task Principal-Agent Problems: Incentive Contracts, Asset Ownership and Job Design." *Journal of Law Economic and Organization*, Vol. 7 (1991), pp. 24–52.

- MYERSON, R. "Optimal Auctions." *Mathematics of Operations Research*, Vol. 6 (1981), 58–73.
- PARKES, D. AND KALAGNANAM, J. "Models of Iterative Multi-Attribute Vickrey Auctions." *Management Science*, Vol. 51 (2005), pp. 435–451.
- REZENDE, L. "Biased Procurement." Mimeo, University of Illinois, 2004.
- RILEY, J.G. AND SAMUELSON, W. "Optimal Auctions." *American Economic Review*, Vol. 71 (1981), pp. 381–392.
- ROCHET, J.C. AND STOLE, L. "Nonlinear Pricing with Random Participation." *Review of Economic Studies*, Vol. 69 (2002), 277–311.
- WILSON, R. "The Architecture of Power Markets." *Econometrica*, Vol. 70 (2002), pp. 1299–1341.