DISCUSSION PAPER SERIES

No. 4734

PROPERTIES OF SCORING AUCTIONS

John Asker and Estelle Cantillon

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP4734.asp**

ISSN 0265-8003

PROPERTIES OF SCORING AUCTIONS

John Asker, Stern School of Business, New York University
Estelle Cantillon, ECARES, Université Libre de Bruxelles, Harvard Business School and CEPR

Discussion Paper No. 4734
November 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4734

November 2004

ABSTRACT

Properties of Scoring Auctions*

This Paper studies scoring auctions, a procedure commonly used to buy differentiated products: suppliers submit offers on all dimensions of the good (price, level of non monetary attributes), and these are evaluated using a scoring rule. We provide a systematic analysis of equilibrium behaviour in scoring auctions when suppliers' private information is multidimensional (characterization of equilibrium behaviour and expected utility equivalence) and show that scoring auctions dominate several other commonly used procedures for buying differentiated products.

JEL Classification: D44, L14 and L24
Keywords: auction, multi-attribute, multidimensional private information and procurement

John Asker
Stern School of Business
New York University
Henry Kaufman Management Center
Suite 7-79, 44 West 4th Street
New York NY 10012
USA
Tel: (1 212) 998 0062
Email: jasker@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=161521

Estelle Cantillon
ECARES
Universite Libre de Bruxelles
CP 114
50, Av. FD Roosevelt
1050 Bruxelles
BELGIUM
Tel: (32 2) 650 44 75
Email: ecantillon@hbs.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=151115

*This Paper supersedes an earlier paper distributed under the title 'Equilibrium in Scoring Auctions'. We thank Max Bazerman, Matthew Jackson, Luca Rigotti, Al Roth, Nicolas Sahuguet as well as audiences at Case Western Reserve, Decentralization Conference 2004, ECARES, Gerzensee, Harvard/MIT, Inform2003, LBS, NYU, Ohio State University, Rome Auction Conference and WZB Berlin for useful conversation and suggestions. Ioannis Ioannou provided excellent research assistance. The financial support of the Division of Research at Harvard Business School is gratefully acknowledged.

Submitted 24 October 2004

1 Introduction

In many procurement situations the buyer cares about attributes in addition to price when evaluating the offers submitted by suppliers. Examples of non monetary attributes that buyers care about include lead time, time to completion, and various other measures of quality. Buyers have adopted several practices for dealing with these situations. Some have recourse to fairly detailed request-for-quotes (RFQ) that specify minimum standards that the offers need to satisfy, and then evaluate the submitted bids based on price only.[1] Others decide to 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 RFQ with the flexibility in terms of contract specification offered by negotiation. Several procedures belong to this category. The buyer can let suppliers offer several combinations of price and non monetary attributes, and choose the combination that best suits his needs. We refer to this procedure as a "menu auction."[2] Alternatively, the buyer can request a single offer from each supplier and again choose the one he prefers among the submitted offers. We call this procedure a "single-bid auction with secret scoring rule." Finally, a procedure that is receiving increasing attention from academic and practitioner communities is the scoring auction. In a scoring auction, the buyer 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 paper, 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 US, 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,[3] and auctions for electricity reserve supply (Bushnell and Oren, 1994; Wilson, 2002). In the EU, a legislative package intended to simplify and modernize existing public procurement laws was recently adopted. As before, 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 criteria used

[1] For example, this is a format proposed by FreeMarkets (see www.freemarkets.com).

[2] This is not to be confused with the menu auctions considered by Bernheim and Winston (1986) which are combinatorial auctions.

[3] The road user cost is the (per day) value of time lost due to construction. By 2003, 38 states in the US were using "A+B bidding." "A+B bidding" is used mainly for large projects for which time is a critical factor. Typically, these represent 5-10% of the total highway construction projects in these states. See, for instance, Arizona Department of Transport (2002) and Herbsman et al. (1995).

when best economic value is the basis for the award.[4] In effect, the new law is mandating the use of scoring auctions. This is significant as public procurement in the EU is estimated at about 16% of GDP.[5] 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.[6]

A distinguishing feature of our model is that suppliers' private information about their cost of providing the good can be multidimensional. In particular, 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, such as timeliness. It also 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 close to what is observed in the data. When private information is one-dimensional, equilibrium offers can be parametrized by a single parameter and describe a curve in the price - attributes space. Our model generates equilibrium offers that can be "all over the place" in the price - attributes space.

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 the equilibrium in these auctions when the scoring rule is linear in price and private information is independent across bidders (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, ...). 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). Theorem 2 generalizes the classic revenue equivalence theorems of Myerson (1981) and Riley and Samuelson (1981). In particular, it implies that the buyer is indifferent among many standard auction formats for the scoring auction when bidders are symmetric in their pseudotypes.

Our second set of results compare scoring auctions to other commonly used procedures to buy differentiated products. Specifically, we show that, from the buyer's perspective, scoring auctions dominate price-only auctions with minimum quality standards. They are equivalent to a menu

[4] 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.

[5] http://europa.eu.int/comm/internal_market/publicprocurement/introduction_en.htm

[6] For example, eBreviate, PurchasePro, Clarus, IBM/DigitalUnion, Oracle Sourcing, Verticalnet and Perfect.

auction and a single-bid auction with secret scoring rule 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 dominate the single-bid auction with secret scoring rule. Finally, when a sealed bid "first price" format is used, we show that scoring auctions dominate from an efficiency perspective. Note that our purpose in this paper is not to determine how to optimally 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 second set of results provide further motivation for focussing 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 anodine when private information is multidimensional. Specifically, we show that such reduction in dimensionality requires that (1) the scoring rule be linear in price, and (2) that private information be independently distributed across suppliers, unless the auction format admits a dominant strategy equilibrium. The two papers we are aware of that allow for multidimensional private information, Bushnell and Oren (1994 and 1995), happen to satisfy these conditions (these papers derive the scoring rule that induces productive efficiency in an environment with multidimensional private information). Our contribution is to show exactly under what conditions using pseudotypes to derive an equilibrium is appropriate when private information is multidimensional. Additionally, our paper proves the stronger result that the pseudotype is a sufficient statistic: not only does it allow us to construct an equilibrium, all equilibria in the scoring auction are only a function of a supplier's pseudotype. We illustrate the benefits of having a sufficient statistic, from equilibrium characterization (Corollary 1) to expected utility ranking (Theorem 2).

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 and Parkes and Kalagnanam, 2002).

Some recent papers study other auction environments with multidimensional private information. Multidimensional private information creates complex incentive situations, including the non-existence of equilibria (Jackson, 1999) or the loss of monotonicity of these equilibria (Reny and Zamir, 2002). In our case, we are able to reduce the relevant dimensionality of private information to one, by exploiting the one-dimensionality of the allocation rule *and* the independence of types across bidders. We can then appeal to the analogy between our environment and the standard IPV environment. A similar property (though through a much more subtle analogy to the standard IPV model) is exploited by Che and Gale (2002) to rank revenue in single-object auctions with multidimensional types and non linear payoffs. In both our and Che and Gale's approach, the one-dimensionality of the allocation decision and the independence of private information across bidders are necessary for reducing the dimensionality of the relevant private information. No such reduction is possible for multi-unit auctions, or for single object auctions where private information is not independent (see Fang and Morris, 2003, for an example).

A variant of scoring auctions are auction environments that involve the sale or purchase of multiple items but where the auctioneer or the procurement authority cannot commit, at the time of the auction, to the quantity sold or purchased. Examples include the sale of timber rights or the purchase of electricity reserve supply. In these auctions, bidders also submit multidimensional bids (often a fixed and a variable price) which are evaluated using a scoring rule. The weight given to the variable price is based on the auctioneer / procurement agency's estimate. The scoring rule is used for allocating the contract, though the final contract often depends on the realized quantities. This creates interesting incentive problems (see Athey and Levin, 2001, Chao and Wilson, 2002, and Ewerhart and Fieseler, 2003). We ignore these aspects in the current paper.

Finally, a few papers consider some of the 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 (Dasgupta and Spulber, 1989, make a similar argument in a slightly different setting). 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. We consider menu auctions, single-bid auctions with secret scoring rule and price-only auctions with minimum quality standards and systematically compare the outcome of equilibrium in these auctions with that in scoring auctions.

The rest of the paper 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, single-bid auctions with secret scoring rule and auctions with minimum quality standards. Section 6 gathers additional remarks.

2 Model

2.1 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 $M \geq 1$ non monetary attributes, $Q \in \mathbb{R}_+^M$.

Preferences. The buyer values the good (p, Q) at $v(Q) - p$, where $v_Q > 0$ and v_{QQ} is a negative definite matrix. Supplier i's profit from selling good (p, Q) is given by $p - c(Q, \theta_i)$, 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 non monetary attributes they can supply.[7] We assume that the marginal cost of producing each attribute is positive, $c_Q > 0$, and that c_{QQ} is positive semi-definite. In particular, this allows for costs that are independent across attributes and convex in individual attributes. We normalize the type space by assuming that $c_{\theta_i} > 0$. Note that the buyer and the suppliers are risk neutral.

These assumptions imply that social welfare $v(Q) - c(Q, \theta_i)$ is strictly concave in Q. The first best level of non monetary attributes for each supplier, $Q^{FB}(\theta_i) = \arg\max\{v(Q) - c(Q, \theta_i)\}$ is well-defined and unique.

Information. Preferences are common knowledge among suppliers and the buyer, with the exception of suppliers' types, θ_i, $i = 1, ...N$, which are privately observed. Types are independently distributed according to the well-behaved joint density function $f_i(\theta_i)$ with support on a bounded and convex subset of $\mathbb{R}^K$, $\Theta_i \subset \mathbb{R}^K$. These density functions are common knowledge.

2.2 Allocation mechanism

We now introduce the scoring auction. We start with two definitions:

A **scoring rule** is a function $S : \mathbb{R}_+^{M+1} \to \mathbb{R} : (p, Q) \to S(p, Q)$, that associates a score to any potential contract between the buyer and a supplier, and represents a continuous preference relation over the contract characteristics (p, Q). A scoring rule is quasi-linear if it can be expressed as $\phi(Q) - p$ or any monotonic increasing function thereof. We assume that the scoring rule in its linear form is strictly increasing and concave in the non-monetary attributes and strictly decreasing in price.[8] Moreover, for simplicity, we let $\phi_Q(0) - c_Q(0) > 0$ to ensure that all suppliers participate in the auction at equilibrium.

[7] Rezende (2003) studies a procurement model with fixed levels of non monetary attributes. In our model, the level of non monetary attributes is determined during the auction process.

[8] It can be shown that any other scoring rule (i.e. putting no weight or negative weight on some attributes) is dominated from the buyer's point.

A **scoring auction** is an allocation mechanism where suppliers compete by submitting bids of the type $(p, Q) \in \mathbb{R}_+^{M+1}$. 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 make in case he does not, t_i^L. A scoring auction is quasi-linear when it uses a quasi-linear scoring rule.

For example, in a first score scoring auction, the winner must deliver a contract that generates the value of his winning score, i.e. $t_i^W = S(p_i, Q_i)$, $t_i^L = 0$. In an ascending "English" scoring auction, the buyer progressively raises the required score to fulfill until all suppliers but one drop out. t_i^W is the value of that score and $t_i^L = 0$. The first score scoring auction and the ascending scoring auction (with its multiple variants) are common implementations of scoring auctions. A useful benchmark is the second score scoring auction. In a second score scoring auction, the winner must deliver a contract that generates a score equal to the second highest score submitted.

Consider supplier i with type θ_i who has won the contract to supply the good with score to fulfill t_i^W. Supplier i will choose characteristics (p, Q) that maximize his profit, i.e.

$$\max_{(p,Q)} \{p - c(Q, \theta_i)\} \qquad \text{subject to } \phi(Q) - p = t_i^W \tag{1}$$

Substituting for p into the objective function yields

$$\max_{Q} \{\phi(Q) - c(Q, \theta_i) - t_i^W\} \tag{2}$$

An important feature of (2) is that the optimal Q is independent of t_i^W. Define

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

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. They are decreasing in types since costs are increasing in types. 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:

$$x_i \left(k(\theta_i) - t_i^W\right) - (1 - x_i) t_i^L \tag{3}$$

In (3), supplier i's preference over contracts of the type (x_i, t_i^W, t_i^L) is entirely captured by his pseudotype. Only quasi-linear scoring rules have this property when private information is multidimensional. Indeed, consider a more general scoring rule $S(p, Q)$ and let's revisit bidder i's optimization problem in this more general case:

$$\max_{(p,Q)} \{p - c(Q, \theta_i)\} \qquad \text{subject to } S(p, Q) = t_i^W \tag{4}$$

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

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

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

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

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 $\widehat{\theta}_i \neq \theta_i$ belong to such a class. It must be that[9]

$$\begin{aligned} u(x_i, t_i^W, t_i^L; \theta_i) &= u(x_i, t_i^W, t_i^L; \widehat{\theta}_i) \text{ if and only if } u(\widehat{x}_i, \widehat{t}_i^W, \widehat{t}_i^L; \theta_i) = u(\widehat{x}_i, \widehat{t}_i^W, \widehat{t}_i^L; \widehat{\theta}_i) \\ &\quad \text{for all pairs of contracts } (x_i, t_i^W, t_i^L), (\widehat{x}_i, \widehat{t}_i^W, \widehat{t}_i^L). \end{aligned} \tag{5}$$

Let $Q(\theta_i, t_i^W) = \arg\max_Q \{\Psi(Q, t_i^W) - c(Q, \theta_i)\}$. Condition (5) requires that $\frac{\partial}{\partial t_i^W} \Psi(Q(\theta_i, t_i^W), t_i^W) = \frac{\partial}{\partial t_i^W} \Psi(Q(\widehat{\theta}_i, t_i^W), t_i^W)$. This equality will in general not be satisfied for $\widehat{\theta}_i \neq \theta_i$ unless Ψ is separable in Q and t_i^W.[10] In turn, this requires that the scoring rule be quasi-linear ($\Psi(Q, t_i^W) = \phi(Q) - t_i^W$ for a quasi-linear scoring rule).

Note that the requirement of quasi-linearity 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 and pseudotypes (and types) capture preferences over contracts whether or not the scoring rule is linear in price.

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 (3). Given suppliers' risk neutrality and the linearity of the scoring rule so that the optimal Q in (4) is independent of t_i^W, 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

$$x_i k(\theta_i) - s_i \tag{6}$$

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 $(x_1, ..., x_N)$ and scores to fulfill

[9] In principle, the requirement of equal preferences only entails that $u(x_i, t_i^W, t_i^L; \theta_i) \geq u(\widehat{x}_i, \widehat{t}_i^W, \widehat{t}_i^L; \theta_i)$ if and only if $u(x_i, t_i^W, t_i^L; \widehat{\theta}_i) \geq u(\widehat{x}_i, \widehat{t}_i^W, \widehat{t}_i^L; \widehat{\theta}_i)$ for all pairs of contracts $(x_i, s_i), (\widehat{x}_i, \widehat{s}_i)$. The stronger requirement in (5) follows from the normalization of utilities embodied in the assumption of risk neutrality.

[10] Indeed, $Q(\theta_i, t_i^W) \neq Q(\widehat{\theta}_i, t_i^W)$ usually for $\widehat{\theta}_i \neq \theta_i$.

by each supplier, $(s_1, ..., s_N)$. (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, $\{(p_i, Q_i)\}_{i=1}^N$.[11] Later in the paper, we will switch to a direct revelation mechanism approach where the outcome will be a function of suppliers' pseudotypes, $(k_1, ..., k_N) \in \mathbb{R}^N$. To avoid introducing too much new notation, we shall make these the arguments of the x and s functions. Similarly, 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 possible equilibria of the auction and evaluate the buyer's expected payoff in each case. Suppliers' original multidimensional types become redundant.

In this section, we prove that pseudotypes are sufficient statistics. Specifically, we show that the sets of equilibria in the scoring auction and in a auction where bidders are constrained to submit a bid only as a function of their pseudotypes coincide. Proving this result requires two preliminary 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 $(x_1, ..., x_N)$ and $(s_1, ..., s_N)$. Second, we prove that equilibria in scoring auctions are essentially pure as a function of pseudotypes.

Lemma 1: *All equilibria of a quasi-linear scoring auction are outcome equivalent to an equilibrium where bidders with the same pseudotypes adopt the same strategies.*

Proof: The proof proceeds in two steps.

Step 1: If there exists an equilibrium in this game, one of them is such that bidders with the same pseudotypes adopt the same strategy.

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

$$(p_i^*, Q_i^*) \in \arg\max_{p,Q} E_{\theta_{-i}} \left[x_i \left((p, Q), (p_{-i}^*, Q_{-i}^*) \right) k_i(\theta_i) - s_i((p, Q), (p_{-i}^*, Q_{-i}^*)) \right] \tag{7}$$

[11] Or, more generally, in the case of open formats, the strategies of the bidders.

where the expression for supplier i's expected profit derives from (6). In (7), suppliers' private information enters their objective function only through their pseudotypes. Hence, supplier i of type θ_i is actually indifferent about the strategies played by all the realizations of supplier i's type with the same pseudotype. Therefore, we can construct a new equilibrium $(\widetilde{\mathcal{E}}_1, ..., \widetilde{\mathcal{E}}_N)$, such that:

1. $\widetilde{\mathcal{E}}_i(\theta_i) = \widetilde{\mathcal{E}}_i(\widehat{\theta}_i)$ whenever $k(\theta_i) = k(\widehat{\theta}_i)$.

2. Define $\Theta_i(k) = \{\theta_i \in \Theta_i | k(\theta_i) = k\}$, the set of supplier i's types with pseudotype equal to k. For each k in the support of bidder i's pseudotypes, the distribution of $\widetilde{\mathcal{E}}_i$ for a given $\theta_i \in \Theta_i(k)$ replicates the aggregate distribution of $\mathcal{E}_i$ over all $\theta_i \in \Theta_i(k)$.

By construction, the distribution of bidder i's opponents' strategies under this new equilibrium is the same as before from bidder i's perspective. Moreover, $\widetilde{\mathcal{E}}_i$ is a best response for bidder i. Hence it is an equilibrium. Moreover, in this equilibrium, bidders' strategies are only a function of their pseudotypes.

Step 2: All other equilibria are outcome equivalent to an equilibrium in which bidders bid only according to their pseudotypes.

This follows directly from step 1 since, by construction, $(\widetilde{\mathcal{E}}_1, ..., \widetilde{\mathcal{E}}_N)$ and $(\mathcal{E}_1, ..., \mathcal{E}_N)$ lead to the same distribution of (p, Q) and therefore scores and outcomes. QED.

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 (7), both through their pseudotypes and through their expectations over their opponents' types.

Lemma 1 implies that the set of possible outcomes $(x_1, ..., x_N)$ and $(s_1, ..., s_N)$ 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[12] types, θ_i and $\widehat{\theta}_i$, such that $k(\theta_i) = k(\widehat{\theta}_i)$ and suppose that in equilibrium, they get a different outcome: (x_i, s_i) and $(\widehat{x}_i, \widehat{s}_i)$. By definition, these two types generate expected utility $f_i(\theta_i)s_i + f_i(\widehat{\theta}_i)\widehat{s}_i$ for the buyer, according to the scoring rule. However, this dif-

[12]This simplifying assumption is inessential for the argument.

fers from true expected utility. To know how much expected utility the suppliers generate for the buyer, we need to know how they will satisfy their obligations. Each supplier finds the pair (p, Q) that generates the required score in the most advantageous way for him. Let Q and $\widehat{Q}$ be the resulting levels of non monetary attributes (they are independent of s_i and $\widehat{s}_i$). Since the scoring rule is quasi-linear, the total monetary transfer from the buyer to the suppliers is then given by $x_i\phi(Q) - s_i$ and $\widehat{x}_i\phi(\widehat{Q}) - \widehat{s}_i$, and the buyer's true expected utility is given by:

$$f_i(\theta_i)\left[x_i\left(v(Q) - \phi(Q)\right) + s_i + \widehat{x}_i\left(v(\widehat{Q}) - \phi(\widehat{Q})\right) + \widehat{s}_i\right]$$

This equilibrium is outcome-equivalent to an equilibrium where type θ_i pretends he is $\widehat{\theta}_i$ and vice versa. On the face of it, the buyer gets again utility $f_i(\theta_i)(s_i + \widehat{s}_i)$ from this equilibrium. However, proceeding as above, we find that his true expected utility is given by

$$f_i(\theta_i)\left[\widehat{x}_i\left(v(Q) - \phi(Q)\right) + \widehat{s}_i + x_i\left(v(\widehat{Q}) - \phi(\widehat{Q})\right) + s_i\right]$$

Clearly, the buyer is not indifferent between these two equilibria unless $x_i = \widehat{x}_i$ and $v(Q) = \phi(Q)$.

The next result ensures that suppliers with the same pseudotypes receive the same equilibrium outcome function (x_i, s_i) in any equilibrium, except possibly on a set of measure zero. This rules out the situation described in the previous example. Lemma 2 then implies that outcome equivalent equilibria are also utility equivalent for the buyer, up to a zero measure.

Lemma 2: *All equilibrium strategies in quasi-linear scoring auctions are essentially pure, both when expressed as a function of pseudotypes and (a fortiori) when expressed as a function of types.*

Note that since the only relevant bid information for the purpose of the outcome of the auction is the score generated by suppliers' bids, the statement of Lemma 2 should be understood as all the types of supplier i with the same pseudotypes submit bids generating the same outcome (x_i, s_i) at equilibrium, for all i.

> **Proof:** We first note that if there exists a non trivial mixed strategy equilibrium (where non trivial refers to mixing for a non zero measure of types), then, by Lemma 1, there exists a non trivial mixed strategy equilibrium in the pseudotypes space. Therefore, we shall focus on equilibrium strategies as a function of pseudotypes to rule out non trivial mixed strategy equilibria.
>
> For each pseudotype k, define $\underline{x}_i(k)$ and $\overline{x}_i(k)$ as the lowest and highest expected probabilities of getting the contract among all the bids in the support of bidder i's strategy when he has pseudotype k. (let $\underline{s}_i(k)$ and $\overline{s}_i(k)$ be the resulting expected score to satisfy). By construction, $\underline{x}_i(k) = \overline{x}_i(k)$ when bidder i of pseudotype k uses a pure strategy.

Define $U_i(k)$ as supplier i's equilibrium expected payoff when he has pseudotype k. Incentive compatibility implies that

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

Hence $\underline{x}_i(k)(k - \widehat{k}) \geq \underline{x}_i(\widehat{k})(k - \widehat{k})$ and $\underline{x}_i(k)$ is monotonically increasing in k. The same argument applies to $\overline{x}_i(k)$. Hence $\underline{x}_i(k)$ and $\overline{x}_i(k)$ are almost everywhere continuous.
A similar argument based on the IC constraint establishes that $\underline{x}_i(k) \geq \overline{x}_i(\widehat{k})$ for all $\widehat{k} < k$. Together with the continuity of these functions, this implies that $\underline{x}_i(k) = \overline{x}_i(k)$ (and $\underline{s}_i(k) = \overline{s}_i(k)$) almost everywhere. This rules out mixed strategy equilibria. QED

We are now able to prove the main result of this section:

Theorem 1: *The set of equilibria (mappings from $\Theta_1 \times ... \times \Theta_N$ to $(p_i, Q_i)_{i=1}^N$) in the unconstrained scoring auction is the same as the set of equilibria in the scoring auction where suppliers are constrained to bid only on the basis of their pseudotypes, except possibly on a set of measure zero.*

Proof: Lemma 2 implies that all equilibria in the unconstrained scoring auction are equilibria in the constrained auction (they differ at most on a set of measure zero). To prove that all equilibria in the constrained auction are also equilibria in the unconstrained auction, note that bidders' preferences and beliefs are entirely determined by their pseudotypes. Therefore, if a strategy is a best response when a supplier is constrained to adopt a strategy based on his pseudotype, this strategy remains a best response for all types θ_i consistent with that pseudotype. QED

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 and 1995). Theorem 1 suggests that doing so does not discard any other equilibria of interest. While this may not be totally surprising when types are one-dimensional, this result is not trivial for environments where types are multidimensional. Indeed, it means that the richness introduced by the higher dimensionality of types has no strategic consequences for the set of equilibria. This property is a consequence of the combination of the quasi-linear 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 quasi-linearity 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.[13] (Multi-unit auctions offer an example of multidimensional allocation mechanisms where there is no reduction of dimensionality possible.)

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

Corollary 1: *The equilibrium in quasi-linear scoring auctions with independent types inherits the properties of the equilibrium in the related single object auction where (1) bidders are risk neutral, (2) their (private) valuations for the object correspond to the pseudotype k in the original scoring auction and are distributed accordingly, (3) the highest bidder wins, and (4) the payment rule is determined as in the scoring auction, with bidders' scores being replaced by bidders' bids.*

Corollary 1 relies on the expression for suppliers' expected payoff in the direct revelation mechanism equivalent of the scoring auction: $x_i(k)k - s_i(k)$. This is identical to the direct revelation mechanism expression for bidders' expected payoff in the standard independent private values single object auction with risk neutral bidders. For example, Corollary 1 implies that an equilibrium exists in a wide variety of formats (e.g. first price, second price, third price, ascending, all-pay, ...). It is unique in the first price scoring auction. See Krishna (2002) for a survey.

Corollary 1 has practical implications for the derivation of the equilibrium in scoring auctions. It forms the basis for 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 $G_i(k)$, $b_i(k)$;

(3) The equilibrium bid in the scoring auction is any (p, Q) such that $S(p, Q) = b_i(k)$. (The actual (p, Q) delivered are easy to derive given $b_i(k)$ and the solution to equation (2).)

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 preferences, i.e. $\phi(Q) = v(Q)$. Theorem 2 shows that this result extends to multidimensional private information and scoring rules that do not correspond to the buyer's true

[13] For example, Branco (1997) extends Che's model to correlated private information.

preferences.[14]

Theorem 2 (Expected Utility Equivalence). *Any two scoring auctions that:*

(a) use the same quasi-linear scoring rule,

(b) use the same allocation rule $x_i(k_i, k_{-i})$, $i = 1, ..., N$, *and*

(c) yield the same expected payoff for the lowest pseudotype $\underline{k}_i$, $i = 1, ..., N$.

generate the same expected utility for the buyer.

Proof: Since the buyer's utility is quasi-linear, his expected utility from a given auction is

$$\sum_{i=1}^{N} E_{k_i,k_{-i}}\left[x_i\left(k_i,k_{-i}\right)ESS\left(k_i\right)-U_i\left(k_i\right)\right]=\sum_{i=1}^{N} E_{k_i}\left[x_i\left(k_i\right)ESS\left(k_i\right)-U_i\left(k_i\right)\right] \tag{8}$$

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, (b) and (c) imply that $U_i(k)$ is the same across both auctions.

Next, fix k_i and let $(p^*(\theta_i, s_i), Q^*(\theta_i, s_i))$ 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 quasi-linear, $Q^*(\theta_i, s_i)$ is only a function of the scoring rule and θ_i, and not of s_i (cf. (2)). Hence,

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

is independent of s_i and therefore equal across the two auctions given (a). The claim follows. QED.

Four points are worth noting concerning this result. First, the assumption that the scoring rule is quasi-linear is key. Without it, suppliers' choice of product characteristics (p, 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

[14]The reason for being interested in $\phi(Q) \neq v(Q)$ is that it may be in the interest of the buyer to misrepresent his true preferences. See Che (1993).

Lemma 2). Indeed, in that case, the social surplus associated with a bidder of type θ_i is his pseudotype k_i, so $EES(k_i) = k_i$ 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. But note that 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.

Finally, one could prove an alternative version of Theorem 2 where (b) is replaced by the requirement that the allocation as a function of the original types, $x_i(\theta_i, \theta_{-i})$, is the same, and (c) is replaced by the requirement that the expected payoff of bidders at a point on the boundary of the types set is the same across auctions. The proof for this alternative version adapts an argument made by Krishna and Perry (2000) in proving a payoff equivalence result for allocation mechanisms with multiple goods and multidimensional types. Note however that the conditions for the alternative version are more restrictive than (b) and (c). The result is therefore weaker. In particular, if we used that approach, we could only establish the equivalence between the first score and the second score auction when bidders are symmetric in the original type space.

5 Comparison with alternatives

So far, we have examined the properties of scoring auctions. As mentioned in the introduction, other procedures are used in practice. In this section, we consider three such alternatives and show that, where revenue results are availiable, the alternatives generate a lower expected utility for the buyer than a scoring auction that uses the true preferences of the buyer. While this does not prove that scoring auctions are optimal in the class of buying mechanisms, it shows that scoring auctions weakly dominate many other commonly used mechanisms.

The first alternative we consider is an auction procedure where the buyer does not reveal his preferences. Instead, suppliers are asked to submit full (p, Q) schedules. We will call this a menu auction.[15] The buyer then selects the offer that he prefers, that is, the one that generates the highest level of utility. This alternative comes in three versions. In the "English" version (E), 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 offer is made. The winner is the supplier who is offering the best schedule in the last round. The resulting contract is the (p, Q) in his schedule

[15] Bichler and Kalagnanam (2003) use the expression "auctions with configurable offers" to describe such procedures.

that the buyer prefers. In the "First Price" version (FPA), suppliers each submit a schedule of offers. The resulting contract is the (p, Q) contract that generates the highest utility to the buyer; the winner is the supplier that offered this contract in his schedule. Finally, in the "Second Price" version (SPA), suppliers again submit a schedule of offers. The winner is the supplier offering the (p, Q) contract that generates the highest utility to the buyer. The resulting contract is $(\widehat{p}, Q)$ where $\widehat{p}$ is adjusted so that $(\widehat{p}, Q)$ generates the same score as the best offer by the losers.

The second alternative procedure we consider is one where the buyer does not reveal his preferences and the suppliers are asked to submit a single bid (p, Q). We call this auction a single bid auction with a secret scoring rule. It potentially comes in three forms: the English, the First Price and the Second Price auctions. These formats are self-explanatory given their description for the menu auction.

The third alternative we consider is one where the buyer publishes a detailed request-for-quote (RFQ) that sets minimum levels for each attribute, and all offers satisfying these conditions are evaluated on a price basis. Again, it comes in three guises: English, First Price and Second Price.

Using as a point of comparison a scoring auction, of the type described in the previous sections, with the scoring rule corresponding to the true preferences of the buyer, we have the following results:

Theorem 3: *Let buyer have unknown preferences, indexed by t, $v(Q,t) - p$, $v_Q > 0$ and v_{QQ} negative definite as before. The distribution of t is common knowledge. We have the following:*

(a) $U^E_{\text{score}} = U^E_{\text{menu}} = U^E_{\text{single}}$ as the bidding increment goes to zero,

(b) $U^{SPA}_{\text{score}} = U^{SPA}_{\text{menu}} > U^{SPA}_{\text{single}}$,

where U^k_l is the expected utility of the buyer, in format $k \in \{E, SPA\}$, and proceedure $l \in \{$score, menu, single$\}$.

The proof for Theorem 3 is in the Appendix. For the English auction, we build an equilibrium for each procedure and show that the equilibrium outcome of any symmetric equilibrium is unique. The equivalence between all three procedures then stems from the direct comparison of the final allocations. For the SPA, we show that submitting a schedule $\mathbb{S} = \{(p, Q)$ such that $p = c(Q, \theta)$, $Q \in R^M\}$ is the unique dominant strategy equilibrium in the menu auction. The equivalence between the menu auction and the scoring auction with $\phi(Q) = v(Q)$ follows directly. For the single-bid procedure, we argue that the equilibrium bid (p^*, Q^*) must belong to $\mathbb{S}$. Since there is a positive probability that (p^*, Q^*) does not belong to $\max_{(p,Q)\in\mathbb{S}}\{v(Q,t) - c(Q,\theta)\}$ for the actual preference - unknown to the suppliers - of the buyer, $U^{SPA}_{\text{menu}} > U^{SPA}_{\text{single}}$ follows.

There are two things to note about Theorem 3 from a practical perspective. First, menu auctions where suppliers do submit full menus are likely to be rare given the cost of preparing bids. In practice, offers might contain up to ten price attributes combinations. Second, the equilibrium in the single bid English auction involves a very high number of bids since suppliers will first exhaust all price attributes combination for a given profit level before reducing their profit targets. Both factors suggest that, in practice, scoring auctions are likely to dominate both procedures in both formats because they save on bidding costs for suppliers.

Another way in which Theorem 3 understates the superiority of scoring auctions over menu and single bid auctions is that the comparison is with a scoring auction with scoring rule $\phi(Q) = v(Q)$. As suggested by the one-dimensional model, the buyer might be better off announcing $\phi(Q) \neq v(Q)$.

The comparison with the menu auction and the single bid auction more complex when the FPA format is used. What we can prove is that, maybe counter-intuitively, the menu auction and the scoring auction are not equivalent when the first price auction format is used. In particular, this implies that the first price menu auction is inefficient

Theorem 4: *Let U_l^{FPA} be the expected utility of the buyer when the strategy space for suppliers and the buyer is $l \in \{score, menu\}$ (where "score" stands for a scoring auction of the type described in the previous sections with the scoring rule corresponding to the true preferences of the buyer). The buyer has unknown preferences, indexed by $t \in T = [\underline{t}, \overline{t}]$, $v(Q,t) - p$, $v_Q > 0$, v_{QQ} negative definite as before, and $v_{Qt} > 0$. Then $U_{\text{score}}^{FPA} \neq U_{\text{menu}}^{FPA}$ and the menu auction is inefficient.*

The proof of Theorem 4 can be found in the Appendix. The intuition for the proof is as follows. Fix supplier i of type θ. Let $(p(t), Q(t))$ the equilibrium first price bid submitted by this supplier in the scoring auction when the announced scoring rule is $v(Q,t) - p$ (this equilibrium can be derived from Corollary 1). We show that submitting schedule $\{(p(t), Q(t))$ for all $t\}$ is not part of the equilibrium in the first price menu auction. This implies that $U_{\text{score}}^{FPA} \neq U_{\text{menu}}^{FPA}$. Moreover, we show that any equilibrium of the menu auction must involve productive inefficiency in the sense that the chosen Q does not maximize $v(Q,t) - c(Q,\theta)$ for the preference realization t and the winner type realization θ.

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

Theorem 5: *A buyer is always better off using a scoring auction with a scoring rule that corresponds to his true preferences 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 $Q = \underline{Q} \in R^M$. Since costs are increasing, suppliers will set their quality levels at $\underline{Q}$. We are now back to a standard procurement auction with symmetric bidders and costs $c(\underline{Q}, \theta_i) \in R$. Let $[x]^{(n)}$ denote the expected value of the n^{th} lowest order statistics from N draws from x, given the distribution of x induced by the distribution of θ_i. The expected utility of the buyer from this minimum quality standard auction is

$$
\begin{aligned}
& v(\underline{Q}) - [c(\underline{Q}, \theta_i)]^{(2)} \\
= \;& [v(\underline{Q}) - c(\underline{Q}, \theta_i)]^{(N-1)} \\
< \;& [\max_Q \{v(Q) - c(Q, \theta_i)\}]^{(N-1)} \\
= \;& \text{Expected utility from the truthful scoring auction by Theorem 2.} \qquad \text{QED}
\end{aligned}
$$

It is straightforward to see why the scoring auction dominates a price only auction with minimum quality standards. By imposing the minimum standards and then only selecting on price the auctioneer has removed any incentive for the bidder to maximise the gains from trade by adjusting quality. Thus, by design, the size of the pie to be divided between auctioneer and bidder has shrunk, without any offsetting advantage in screening.

6 Concluding remarks

When products are differentiated and complex, buyers often see their choice to be limited between a (price only) auction and negotiation.[16] Scoring auctions are an alternative that combines the competition advantage of auctions with the flexibility in terms of contractual terms of the second. As such, they are receiving increased attention in practitioner and policy circles. Yet, the theory of scoring auctions is still incomplete.

Our paper provides a systematic analysis of equilibrium behavior in scoring auctions when private information is multi-dimensional. 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 quasi-linear and types are independently distributed. 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

[16] See Bajari et al. (2002) for evidence from the construction industry in California. In fact, a recent survey of corporate purchasers, published in the August 2003 issue of Purchasing magazine, showed that the use of auctions, and in particular, online auctions, is still very limited. A frequent reason that corporate purchasers give for no using auctions is that "you can't source complicated goods through price-only auctions."

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, include some like the menu auction and the single-bid auction with secret scoring rule, that also combine competition with the flexibility of deciding on all the dimensions of the product, are dominated by scoring auctions. These results suggests 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. One desirable extension of our model is to allow for suppliers' costs to be uncertain at the time of bidding. Suppose that the cost of attribute Q is given by $c(Q, \theta + \tau)$ where only θ is known to the supplier at the time of bidding. The consequence for our model depends on whether the level of Q submitted by the supplier is binding or not (that is, only the score is binding and the supplier can reoptimize on the level of non monetary attribute once the uncertainty has been resolved). If Q is binding, our results continue to hold if we replace $c(Q, \theta)$ by $\overline{c}(Q, \theta) = E_{\tau}[c(Q, \theta + \tau)]$. Indeed, $\overline{c}(Q, \theta)$ preserves all the properties of $c(Q, \theta)$ (increasing and convex in Q). By contrast, if suppliers are allowed to reoptimize once they are given their score to fulfill, the extension is not trivial.

Implications for empirical work. Even in the presence of symmetric suppliers, scoring auctions present interesting and non trivial auction design questions (How can the buyer manipulate the scoring rule to its 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 non identified 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 while the information to be inferred is multi-dimensional. This suggests two possible solutions. When the (p, Q) offers rather than the scores are binding, the observed data is again multi-dimensional. Another possibility is to look at auction data when changes in the scoring rule can be exploited. In any case, our paper provides a theoretical basis from which investigation of identification is feasible.

Optimal multi-attribute auctions. As we have noted at several places, while quasi-linear scoring rules are known to be the optimal way to buy a differentiated product when private information is one-dimensional (Che, 1993), they are unlikely to be optimal when private information is multi-dimensional. An area for future research is to characterize the optimal procedure for buying a differentiated product when private information is multi-dimensional. Answering this question requires that we take a different approach (see Asker and Cantillon, 2004).

References

[1] Arizona Department of Transport (2002), A+B Bidding Guide, http://www.dot.state.az.us/roads/constgrp/A+BGuide.pdf (dated 1/18/2002)

[2] Asker, John and Estelle Cantillon (2004) How to Buy a Differentiated Product, mimeo.

[3] Athey, Susan and Jon Levin (2001), Information and Competition in US Timber Auctions, *Journal of Political Economy*, 109(2), 375-417

[4] Bajari, Patrick, Robert McMillan and Steve Tadelis (2002), Auctions vs. Negotiation in Procurement: An Empirical Analysis, mimeo

[5] Bajari, Patrick, Stephanie Houghton and Steve Tadelis (2004), Bidding for Incomplete Contracts, mimeo.

[6] Bernheim, B. Douglas and Michael Whinston (1986), Menu Auctions, Resource Allocation, and Economic Influence, *Quarterly Journal of Economics,* 101, 1-31

[7] Bichler, Martin and Jayant Kalagnanam (2003), Configurable Offers and Winner Determination in Multi-Attribute Auctions, *European Journal of Operational Research, forthcoming.*

[8] Bikhchandani, Suschil, Philip Haile and John Riley (2002), Symmetric Separating Equilibria in English Auctions, *Games and Economic Behavior,* 38, 19-27.

[9] Branco, Fernando (1997), The Design of Multi-dimensional Auctions, *Rand Journal of Economics, 28(1), 63-81.*

[10] Bushnell, James and Shmuel Oren (1994), Bidder Cost Revelation in Electric Power Auctions, *Journal of Regulatory Economics*, 6(1), 5-26.

[11] Bushnell, James and Shmuel Oren (1995), Internal Auctions for Efficient Sourcing of Intermediate Products, *Journal of Operations Management*, 12(3-4), 311-320.

[12] Chao, Hung-po, and Robert Wilson (2002): Incentive-Compatible Evaluation and Settlement Rules: MultiDimensional Auctions for Procurement of Ancillary Services in Power Markets, *Journal of Regulatory Economics*, to appear.

[13] Che, Yeon-Koo (1993), Design Competition through Multi-dimensional Auctions, *Rand Journal of Economics,* 24(4), 668-680.

[14] Che, Yeon-Koo and Ian Gale (2002), Auctions with Nonlinear Payoffs and Multidimensional Types, University of Wisconsin and Georgetown mimeo.

[15] Ewerhart, Christian and Karsten Fieseler (2003), Procurement auctions and unit-price contracts, *Rand Journal of Economics,* 34(3), 569-581.

[16] Herbsman, Zohar, Wei Tong Chen and William C. Epstein (1995), Time is Money: Innovative Contracting Methods in Highway Construction, *Journal of Construction Engineering and Management*, 121(3), 273-281.

[17] Jackson, Matthew (1999), The Non-Existence of Equilibrium in Auctions with Two Dimensional Types, California Institute of Technology, mimeo.

[18] Krishna, Vijay (2002), *Auction Theory*, Academic Press, San Diego, California.

[19] Krishna, Vijay and Motty Perry (2000), Efficient Mechanism Design, Penn State mimeo.

[20] Fang, Hanming and Stephen Morris (2003), Multidimensional Private Value Auctions, Yale mimeo.

[21] Myerson, Roger (1981), Optimal Auctions, *Mathematics Op. Research, 6, 58-73.*

[22] Parkes, David, and J. Kalagnanam (2002), Iterative Multi-Attribute Vickrey Auctions, IBM Research mimeo.

[23] Reny, Philip and Shmuel Zamir (2002), On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions, University of Chicago mimeo.

[24] Rezende, Leonardo (2003), Biased Procurement, Stanford mimeo.

[25] Riley, John G. and William Samuelson (1981), Optimal Auctions, *American Economic Review*, 71, 381-392.

[26] Wilson, Robert (2002), The Architecture of Power Markets, *Econometrica*, 70(4), 1299-1341.

7 Appendix

Proof of Theorem 3:

(a) To simplify the argument, we discretize the price grid and therefore the utility and profit grids. Let Δ be the minimum price (and therefore profit and utility) increment. (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, $\{(p, Q);\ Q \in R^M,\ p - c(Q, \theta) = \text{constant and } \max_t \max_{(p,Q)}\{v(Q, t) - p\} = 0\}$. Let π_t be the profit level corresponding to the schedule at period t for a given supplier. At round t, this supplier submits schedule $\{(p, Q);\ Q \in R^M,\ p - c(Q, \theta) = \pi_{t-1} - \Delta\}$ as long as $\pi_{t-1} - \Delta \geq 0$ if his offer was not selected in round $t - 1$. Clearly, this is an equilibrium (moreover, we can adapt the arguments in Bikchandani et al. 2002 to argue that the outcome of this strategy is the unique equilibrium outcome in symmetric strategies). Each supplier stays in the race as long as a positive profit can be made, otherwise they drop out. The selected level of attributes, Q^*, satisfies $Q^* = \arg\max\{v(Q, t) - c(Q, \theta)\}$ for the realization of t. Moreover, the final price satisfies $p = v(Q^*, t) - \max_Q\{v(Q, t) - \widetilde{c}(Q, \widetilde{\theta})\}$ (modulo the increment) where $\widetilde{c}(Q, \widetilde{\theta})$ refers to the cost function of the second best supplier. This is exactly the outcome of the scoring auction. (2) Single-bid auction: 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, $\{(p, Q);\ Q \in R^M,\ p - c(Q, \theta) = \text{constant and } \max_t \max_{(p,Q)}\{v(Q, t) - p\} = 0\}$. 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 $\{(p, Q);\ Q \in R^M,\ p - c(Q, \theta) = \pi_{t-1}\}$ that he has not submitted in the past. If no unsubmitted bid remains in this schedule, then the supplier submits a bid in $\{(p, Q);\ Q \in R^M,\ p - c(Q, \theta) = \pi_{t-1} - \Delta\}$ 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. It guarantees that, conditional on winning, suppliers' profits are maximal. Also the winner in the single bid auction is the same as in the menu auction. However, the buyer may not be equally well off as in the menu auction since here, he cannot choose the (p, Q) pair that maximizes his utility. However, as Δ goes to zero, the winning (p, Q) must maximize $v(Q, t) - c(Q, \theta)$ since otherwise, the winning bidder could have won with a higher level of profit and this is ruled out by his bidding behavior. Thus, the buyer is equally well off.

(b) (1) Menu auction: We first show that $\{(p, Q);\ p = c(Q, \theta),\ Q \in R^M\}$ is a dominant strategy equilibrium. The argument is identical to the argument for the standard second price auction. The portions of the schedule where $\pi = p - c(Q, \theta) > 0$ are weakly dominated by the alternative $(\widehat{p},\ c(Q, \theta))$ where $\widehat{p} = c(Q, \theta)$. The deviation increases the probability of winning a positive profit and does not affect the profit for the cases where the original schedule already won. By the same argument, the portions of the schedule where $\pi = p - c(Q, \theta) < 0$ are weakly dominated by $(\widehat{p},\ c(Q, \theta))$ where $\widehat{p} = c(Q, \theta)$. The argument is represented graphically in Figure 1(a) for the case of $Q \in R$. The iso-profit curves of a given supplier are represented together with a realization of the buyer's preferences. The original schedule is given by the dotted line. Area A is an area where profits could be made and are currently missed. Area B is an area where losses could be made and avoided. The equivalence with the scoring auction follows from the fact that suppliers submit bids such that $p = c(Q, \theta)$ in the dominant strategy equilibrium of the scoring auction, with $Q = \arg\max\{v(Q) - c(Q, \theta)\}$. This is also the bid selected by the buyer in the winning schedule. The best second offers in both auctions are also identical by the same argument.

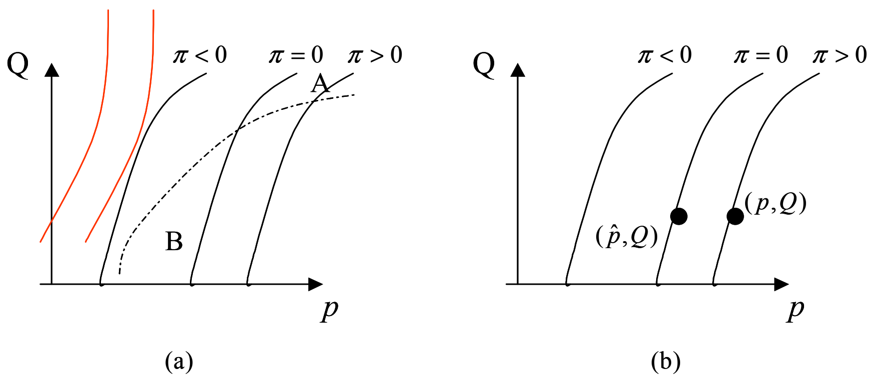

Figure 1: Equilibria in the Second Price format

(2) Single-bid auction: We first show that, at equilibrium, suppliers submit an equilibrium bid (p^*, Q^*) in the schedule $\{(p, c(Q, \theta));\ p = c(Q, \theta),\ Q \in R^M\}$. Consider any alternative bid (p, Q) such that $p - c(Q, \theta) > 0$. The expected profit generated by this bid is equal to

$$\text{prob}\,((p, Q) \text{ generates the highest score})\,(Ep - c(Q, \theta))$$

where Ep is the expected resulting price determined by the second best offer. Consider the deviation $(\widehat{p}, c(Q, \theta))$ where $\widehat{p} = c(Q, \theta)$. This offer will win more often. The expected profit it generates is equal to

$$\begin{aligned}&\text{prob}\,((p, Q) \text{ generates the highest score})\,(Ep - c(Q, \theta)) + \\ &\text{prob}\,((\widehat{p}, Q) \text{ generates the highest score but } (p, Q) \text{ does not})\,(E\widetilde{p} - c(Q, \theta))\end{aligned}$$

where $E\widetilde{p}$ is the expected price given that $(\widehat{p}, Q)$ generates the highest score but (p, Q) does not. Clearly, $E\widetilde{p} - c(Q, \theta) > 0$. A similar argument establishes that (p, Q) such that $p - c(Q, \theta) < 0$ is dominated. The argument is represented graphically in Figure 1(b). From the buyer's point of view, $(p^*, Q^*) \in \{(p, c(Q, \theta));\ Q \in R^M,\ p = c(Q, \theta)\}$ is worse than $\max\{v(Q, t) - p;\ p = c(Q, \theta),\ Q \in R^M\}$ with strictly positive probability. QED

Proof of Theorem 4: Menu auction: Fix supplier i of type θ. Let $(p^s(t), Q^s(t))$ be the equilibrium first price bid submitted by this supplier in the scoring auction when the announced scoring rule is

$v(Q,t)-p$ (this equilibrium can be derived from Corollary 1). This equilibrium entails productive efficiency since each bidder maximizes the total surplus $v(Q,t)-c(Q,\theta)$. We show that submitting schedule $((p^s(t),Q^s(t))_{t\in T}$ is not part of the equilibrium in the first price menu auction.
By definition, $(p^s(t),Q^s(t))$ maximizes supplier i's expected profit:

$$\max_{p,Q}(p-c(Q,\theta))G(v(Q,t)-p;t)^{N-1} \tag{9}$$

where G is the cumulative distribution of scores. Let $\pi^s(\theta,t)$ be the supplier's expected profit associated with $(p^s(t),Q^s(t))$ when the preferences of the buyer are t.

Step 1: A necessary and sufficient condition for $(p^s(t),Q^s(t))_{t\in T}$ to be an equilibrium in the menu auction is that $(p^s(t),Q^s(t))$ corresponds to an iso-profit curve for the supplier.

In the menu auction, the buyer chooses the pair (p,Q) that maximizes his utility among the offers in the schedule, $(p(t),Q(t))_{t\in T}$. Suppose supplier i submits schedule $(p^s(t),Q^s(t))_{t\in T}$. If

$$v(Q^s(t),t)-p^s(t)\geq v(Q^s(\widehat{t}),t)-p^s(\widehat{t}) \qquad \forall t,\widehat{t}\in T, \tag{10}$$

then the buyer with preference t will indeed select the offer intended for him $(p^s(t),Q^s(t))$. As a result, the supplier's maximization problem in the menu auction is decomposable into a continuum of scoring auction problems of the kind in (9). The problem is separable in t.
A necessary and sufficient condition for (10) to hold when t describes a continuum of preferences with $v_{Qt}>0$ is that the price - attribute trade-off involved in schedule $(p^s(t),Q^s(t))_{t\in T}$ corresponds to the price - attribute trade-off present in each individual scoring auction problem, that is,

$$\max_{(p,Q)}\{v(Q,t)-p(t)\} \text{ subject to } p(t)-c(Q,\theta)= \text{ constant}$$

This will be the case only if the supplier's profit associated with $(p^s(t),Q^s(t))_{t\in T}$, $\pi^s(\theta,t)$, is independent of t so that $(p^s(t),Q^s(t))_{t\in T}$ corresponds to an iso-profit curve for the supplier. Figure 2 illustrates this argument graphically for the case where $Q\in R$. In the figure, the buyer has preferences $\widehat{t}$ (represented by two iso-utility curves). The schedule $(p^s(t),Q^s(t))_{t\in T}$ is represented in dotted lines, while the supplier's iso-profit curve going through $(p^s(\widehat{t}),Q^s(\widehat{t}))$ is in full line. When $(p^s(t),Q^s(t))_{t\in T}$ does not correspond to an iso-profit curve, the curvature of $(p^s(t),Q^s(t))_{t\in T}$ is not the same as the curvature of the supplier's iso-profit curve. As a result, $(p^s(\widehat{t}),Q^s(\widehat{t}))$ does not maximize the buyer's utility among the schedule of offers (instead, B does).

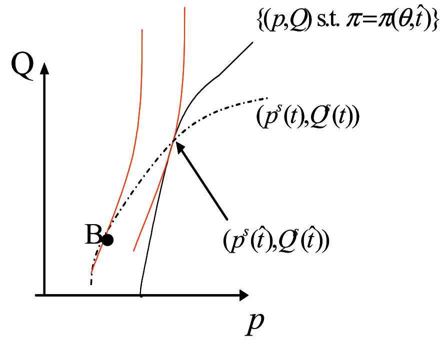

Figure 2: Graphical representation of the argument

Step 2: $(p^s(t), Q^s(t))$ does not correspond to an iso-profit curve.

The reason why $\pi^s(\theta, t)$ depends on t is that t affects the relative advantage of types. To simplify, consider a case where $c(Q, \theta) = \theta_1 + c_2(Q, \theta_2)$. It is sufficient to prove that $\pi^s(\theta, t)$ is not constant in t for at least one θ (continuity will ensure that the statement will be true for a non zero measure of types). Consider $\theta = (\underline{\theta}_1, \underline{\theta}_2) = \underline{\theta}$, the "best" type a supplier can have. Then, $G(v(Q^s(t), t) - p^s(t), t)^{N-1} = 1$ for any t, and from Corollary 1, $(p^s(t) - c(Q^s(t), \underline{\theta})) = (k(\underline{\theta}, t) - k(t)^{(1;N-1)})$, where $k(t)^{(1;N-1)}$ is the expected value of the first order statistics of the pseudotypes of supplier $\underline{\theta}$'s opponents. Consider t and $\widehat{t}$ with $\widehat{t} > t$. Consider the (θ_1, θ_2) space (Figure 3). The linearity of pseudotypes in θ_1, the fact that for the purpose of computing $(k(\underline{\theta}, t) - k(t)^{(1;N-1)})$, we can normalize $k(\underline{\theta}, t) = k(\underline{\theta}, \widehat{t})$ and the fact that the iso-pseudotype curves for $\widehat{t}$ are less steep than those for t implies that $k(\widehat{t})^{(1;N-1)} < k(t)^{(1;N-1)}$ by first order dominance. (In Figure 3, the dotted lines correspond to the iso-pseudotype curves for t. The full lines correspond for those for $\widehat{t}$)

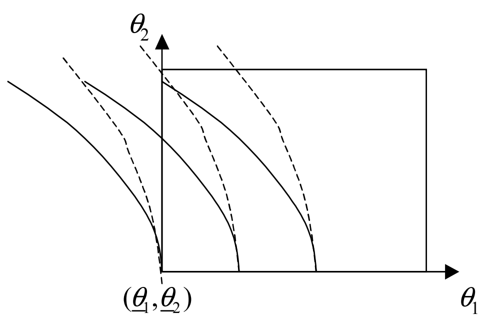

Figure 3: $\pi^s(\theta, t)$ is not constant in t

Steps 1 and 2 imply that $(p^s(t), Q^s(t))$ is not an equilibrium in the menu auction. Our last step proves that there is no other efficient equilibrium.

Step 3: There is no efficient equilibrium $(p(t), Q(t))$ in the menu auction.

Towards a contradiction suppose there is such an efficient equilibrium, $(p(t), Q(t))$. Given t, the buyer prefers to select $(p(t), Q(t))$ in the schedule of offers submitted by the supplier, and the equilibrium involves efficient production and allocation.

Fix t. Efficiency requires that, for all θ, $Q(t) = \arg\max\{v(Q,t) - c(Q,\theta)\}$, and that the winning supplier is the one whose type maximizes $v(Q(t),t) - c(Q(t),\theta)$ among all of the suppliers' type realizations. Incentive compatibility (for a given t realization) then pins down $\pi(\theta, t)$, the expected profit of a type θ supplier when buyer preference is t.

$$\pi(\theta,t) = \int_{\underline{k}}^{k(\theta,t)} x(k,t)dk + \text{constant}(t)$$

where $\underline{k}$ is the lowest pseudotype given t and $k(\theta, t)$ is the pseudotype associated with θ when the buyer's preference is t.

We now argue that in any equilibrium of the menu auction, the constant, which corresponds to the expected payoff of the supplier with the lowest pseudotype (and type, given that $c_\theta > 0$), is equal to zero, for all t. Clearly it cannot be negative otherwise the supplier with the highest type $\overline{\theta}$ (lowest pseudotype) would have a profitable deviation for some t in T. It cannot be positive either because supplier with type $\overline{\theta}$ wins with zero probability in all realizations of t. A positive profit for some t would yield a deviation opportunity as well.

Hence, by the payoff equivalence theorem (a subset of Theorem 2), $\pi(\theta,t) = \pi^s(\theta,t)$, which we have shown, does not correspond to an iso-profit curve. As a result, $(p(t), Q(t))$ is not the offer selected when the buyer's preference is t, contradicting efficiency (the argument is analogous to the argument made in Step 1). QED.