consip

QUADERNI CONSIP

Ricerche, analisi, prospettive

III 2004

Equilibrium in Scoring Auctions

MEF

Ministero
dell'Economia
e delle Finanze

QUADERNI CONSIP

Ricerche, analisi, prospettive

III 2004

Equilibrium in Scoring Auctions

John Asker (Economics Department, Harvard University, email: asker@fas.harvard.edu)

Estelle Cantillon (Harvard Business School, email: ecantillon@hbs.edu)

20 Dicembre, 2004

Ministero dell'Economia e delle Finanze

Index

Abstract

This paper studies multi-attribute auctions in which a buyer seeks to procure a complex good and evaluate offers using a quasi-linear scoring rule. Suppliers have private information about their costs, which is summarized by a multi-dimensional type. The scoring rule reduces the multidimensional bids submitted by each supplier to a single dimension, the score, which is used for deciding on the allocation and the resulting contractual obligation. We exploit this idea and obtain two kinds of results. First, we characterize the set of equilibria in quasi-linear scoring auctions with multi-dimensional types. In particular, we show that there exists a mapping between the class of equilibria in these scoring auctions and those in standard single object IPV auctions. Second, we prove a new expected utility equivalence theorem for quasi-linear scoring auctions.

We thank Howard Georgi, Luca Rigotti, Al Roth, as well as seminar audiences at LBS, Ecares, Inform2003, Ohio State University, 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.

1. Introduction

In many procurement situations the buyer cares about attributes other than just price when evaluating the desirability of contracts offered by sellers. Standard non-monetary attributes 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 bidders and negotiate on all the dimensions of the contract with each of them.

A third option is to combine the competition induced by the first option with the flexibility of the second by holding a scoring auction. In a scoring auction, bidders submit offers on all dimensions of the good and the buyer uses a scoring rule to evaluate the offers and select the winner. Scoring auctions can be shown to dominate RFQs with minimum standards and the restricted negotiation, at least from an efficiency standpoint.[2]

In this paper, we study scoring auctions that use quasi-linear scoring rules (i.e. the price enters linearly in the final score). The buyer cares about several (price and non price) attributes of the good and several bidders compete for the contract. 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 a road user cost,[3] and auctions for electricity reserve supply (Bushnell and Oren, 1994; Wilson, 2002).of scoring auctions is also gaining

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

[2] *An argument for why scoring auctions dominate RFQs with minimum standards is provided in Che (1993).*

[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).*

favor in private sector procurement, with several procurement software developers incorporating scoring capability in their auction designs.[4]

A key feature of our environment is that suppliers' private information about their costs of providing the good can be multi-dimensional. 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. As another example, it allows us to consider the likely situation where firms differ in their fixed and variable costs of production.

Our main results are as follows. First, we prove that the multi-dimensionality of suppliers' private information can be reduced to a single dimension (his "pseudotype") that is sufficient to characterize the equilibrium in these auctions when the scoring rule is quasi-linear 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 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, ...).

Second, we prove a new expected utility theorem for the buyer when private information is multidimensional and independently distributed and the scoring rule is quasi-linear (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 a first score, a second score, an ascending (the equivalent of the Dutch format in this procurement setting) or a descending scoring auction when bidders are symmetric in their pseudotypes.

There are several papers studying scoring auctions. Che (1993) derives a series of important results for multi-attribute auctions when private information is one-dimensional.[5] Our paper extends Che's results on equilibrium in scoring auctions, in several ways.

---

[4] *In the US market, the Oracle Sourcing software (via www.oracle.com) is a good example of this. Verticalnet (via www.verticalnet.com) also provides a scoring capability.*

[5] *Branco (1997) extends Che's paper to affiliated costs for the suppliers.*

First, we allow for multi-dimensional private information. Second, while Che already exploited the idea of pseudotypes to derive an equilibrium in the first and second score auction, our characterization result establishes that no other equilibria exists. Third, our characterization result allows us to prove an expected utility equivalence theorem for all quasi-linear scoring rules, and not only the truthful one. Bushnell and Oren (1994 and 1995) derive the scoring rule necessary for productive efficiency in a two dimensional private information setting . They then implicitly exploit the sufficient statistics property of these auctions to derive a symmetric equilibrium in the second score auction. Our Theorem 1 and Corollary 1 establish that the symmetric equilibrium they derive is indeed the only one.

On a more general level, this paper provides a precise exposition of the applicability of the sufficient statistics approach leveraged by these papers. We show why the use of a Multi-dimensional private information creates much more 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 auctions with multi-dimensional 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 one 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 multi-dimensional 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 and Chao and Wilson, 2002). We ignore these aspects in the current paper.

The rest of the paper is organized as follows. Section 2 describes the model and introduces thesufficient statistic is appropriate in these environments. We are also able to outline the limitations of the approach. In particular we show that, when the scoring rule is not quasi-linear, a sufficient statistic approach will not work. Similarly, our characterization result makes clear why independence of signals is critical. Most importantly, we show that the sufficient statistic approach is a powerful and simple tool for the analysis of equilibrium in a far richer class of scoring auction environments than previously investigated, including when private information is multi-dimensional. Some recent papers study other auction environments with multi-dimensional private information. cost,3 and auctions for electricity reserve supply (Bushnell and Oren, 1994; Wilson, 2002). The use notion of pseudotypes. 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 concludes.

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 ≥ 1 non monetary attributes, $Q \in \mathbb{R}$.

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 ≥ 1, is supplier i's type. We allow suppliers to be flexible with respect to the level of non-monetary attributes they can supply.[6] 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 $c_\theta > 0$.
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, $Q^{FB}(\theta_i) = \text{argmax}\{v(Q) - c(Q, \theta_i)\}$ is well-defined and unique.

Information. Preferences are common knowledge among bidders and the buyer, with the exception of suppliers' types, θ_i, i = 1, ...N, which are privately observed by each bidder. 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.

[6] Rezende (2003) studies an example of 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.

2.2 Allocation mechanism

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

A scoring rule is a function $S : \mathbb{R}^{M+1}+1 - \mathbb{R} : (p,Q) - S(p,Q)$, 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 is strictly increasing and concave in non monetary attributes and strictly decreasing in price.[7]

A **scoring auction** is an allocation mechanism where suppliers compete by submitting bids of the type $(p,Q) \in \mathbb{R}^{M+1}+1$. Bids are evaluated according to a scoring rule. The winner is the bidder with the highest score. Bidders are given scores $(s_1, \ldots, s_N)$ to deliver, where these scores are functions of bidders' strategies.[8] A scoring auction is quasi-linear when it uses a quasi-linear scoring rule. For example, in a first score scoring rule, the winner must deliver a contract that generates the value of his winning score. In a second score scoring auction, the winner must deliver a contract that generates the value of 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 si. Supplier i will choose the good characteristics (p,Q) that maximize his profit, i.e.[9]

$$\max_{(p,Q)} \{p - c(Q, \theta_i)\} \qquad \text{subject to } \phi(Q) - p = s_i \quad (1)$$

Substituting for p into the objective function yields

$$\max_{Q} Q \{\phi(Q) - c(Q, \theta_i) - s_i\} \qquad (2)$$

[7] 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 of view by the truthful scoring rule, $S(p,Q) = v(Q) - p$.

[8] The fact that the resulting obligations from the auction are in terms of scores rather than contract characteristics means that only the score of bidders' bid matters. The (p,Q) components of the bids are not binding. While this may appear to contradict standard practice, it does not. Indeed, contracts often recognize that suppliers may not be able to fully commit to some levels of non-monetary attributes and therefore include penalties and rewards for under- and over-supply of such attributes. For example, even though the resulting obligations in the US highway construction projects using "A+B bidding" are in terms of price and a fixed time to completion, delays or faster than planned completion are penalized / rewarded at exactly the same rates as in the scoring rule. This makes the effective resulting obligation in terms of score, instead of specific levels of price and time to completion.

[9] Note that if supplier i must deliver score si without winning the contract, the only way he can fullfill the score is with money.

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 profit when he wins with probability x_i and needs to deliver score s_i is given by[10]

$$x_i\,(k(\theta_i) - s_i) - (1 - x_i)s_i = x_i k(\theta_i) - s_i \qquad (3)$$

An important feature of (3) is that bidder i's preference over contracts of the type (x_i, s_i) is entirely captured by his pseudotype. Only quasi-linear scoring rules have this property. Indeed, consider a more general scoring rule $S(p,Q)$ and revisit bidder i's optimization problem in this more general case:

$\max_{(p,Q)} \{\text{p} - c(Q, \theta_i)\}$ subject to $S(p,Q) = s_i$

Let $\psi(Q, s_i)$ the price required to generate a score of si 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 s_i).
The objective function of bidder i becomes

$$\max_Q \text{ Q } \{\psi(Q, s_i) - c(Q, \theta_i)\},$$

and his expected payoff from contract (xi, si) is given by:

$$u(x_i, s_i; \theta_i) = x_i \max_Q \{\psi(Q, s_i) - c(Q, \theta_i)\} - (1 - x_i)s_i$$

Suppose we could 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 $\theta_i = \theta_i$ belong to such a class.

[10] *The notation we adopt assumes that si is to be fullfilled whether or not the contract is won. This does not assume that we are restricting ourselves to auction formats where bidders have obligations whether they win or not. Instead, si can be interpreted as an expected level of obligation. The fact that in (2), the optimal choice of Q does not depend on si makes it irrelevant when the obligation must be fullfilled.*

It must be that[11]

$u(x_i, s_i; \theta_i) = u(x_i, s_i; \hat{\theta}_i)$ if and only if $u(\hat{x}_i, \hat{s}_i; \theta_i) = u(\hat{x}_i, \hat{s}_i; \hat{\theta}_i)$ (4)

for all pairs of contracts (x_i, s_i), $(\hat{x}_i, \hat{s}_i)$.

Let $Q(\theta_i, s_i) = \text{argmax}_Q\{\psi(Q, s_i) - c(Q, \theta_i)\}$.

Condition (4) requires in particular that

$$\frac{\partial}{\partial s_i} \psi = Q\left((\theta_i, s_i), s_i\right) = \frac{\partial}{\partial s_i} \psi = Q\left((\theta_i, s_i), s_i\right)$$

This equality will in general not be satisfied for $\hat{\theta}_i \neq \theta_i$ unless ψ is separable in Q and s_i.[12] In turn, this requires that the scoring rule be quasi-linear ($\psi(Q, s_i) = \phi(Q) - s_i$ for a quasi-linear scoring rule).

Notation: For the remainder of this paper, we adopt the following notation and conventions. The outcome function of a scoring auction is a vector of probabilities of winning $(x_1, \ldots, x_N)$ and scores to fulfill by each supplier, $(s_1, \ldots, s_N)$. (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, $\{(p_i, Q_i)\}^N = 1$.[13] Later in the paper, we will switch to a direct revelation mechanism approach where the outcome will be a function of suppliers' pseudotypes, $(k_1, \ldots, k_N) \in \mathbb{R}^N$.
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 xi over the types of the other suppliers, $E_{k-1}\, x_i(k_i, k_{k-1})$. The arguments will be spelled out whenever confusion is possible.

[11] *In principle, the requirement of equal preferences only entails that* $u(x_i, s_i; \theta_i) \geq u(\hat{x}_i, \hat{s}_i; \theta_i)$ *if and only if* $u(x_i, s_i; \hat{\theta}_i) \geq u(\hat{x}_i, \hat{s}_i; \hat{\theta}_i)$ *for all pairs of contracts* (x_i, s_i), $(\hat{x}_i, \hat{s}_i)$. *The stronger requirement in (4) follows from the normalization of utilities embodied in the assumption of risk neutrality.*

[12] *Indeed,* $Q(\theta_i, s_i) \neq Q(\hat{\theta}_i, s_i)$ *usually for* $\hat{\theta}_i \neq \theta_i$.

[13] *Or, more generally, in the case of dynamic formats, the strategies of the bidders.*

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 multi-dimensional 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 $(\varepsilon_1, ..., \varepsilon_N)$ where ε_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,

$$(p_i, Q_i) \in \underset{pQ}{\operatorname{argmax}}\ E_{\theta-i}\left[x_i((p,Q), (p^*_{-i}, Q^*_{-i}))k_i(\theta_i) - s_i((p,Q), (p^*_{-i}, Q^*_{-i}))\right) \qquad (5)$$

where the expression for bidder i's expected profit derives from (3).
In (5), bidders' private information enters their objective function only through their pseudotypes. Hence, bidder i of type θ_i is actually indifferent among the strategies played by the other bidders with the same pseudotype. Therefore, we can construct a new equilibrium $(\tilde{\varepsilon}_1, ..., \tilde{\varepsilon}_N)$, such that:

1.$\tilde{\varepsilon}_1(\theta_i) = \tilde{\varepsilon}_N(\theta_i)$ whenever $k((\theta_i)) = k((\hat{\theta}_i))$.

2. Define $\Theta_i(k) = \{\theta_i \in \Theta_i \mid k(\theta_i) = k\}$, 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 ε_1 for a given $\theta_i \in \Theta_i(k)$ replicates the aggregate distribution of $\tilde{\varepsilon}_1$ 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, $\tilde{\varepsilon}_1$ 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, $(\tilde{\varepsilon}_1, \ldots, \tilde{\varepsilon}_N)$ and $(\varepsilon_1, \ldots, \varepsilon_N)$ 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. Without it, bidders' private information would enter their expected payoff in (5), both through their pseudotypes and through their expectations over their opponents' types. In that case, the argument for the outcome equivalence between all equilibria in the scoring auction and those where suppliers are constrained to bid only according to their pseudotypes breaks down.

Lemma 1 implies that the set of possible outcomes $(x_1, \ldots, x_N)$ and $(s_1, \ldots, s_N)$ 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[14] types, θ_i and $\hat{\theta}_i$, 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.

[14] *This simplifying assumption is inessential for the argument.*

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

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

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

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

Clearly, the buyer is not indifferent between these two equilibria as soon as $x_i \neq \hat{x}_i$. The next result ensures that suppliers with the same pseudotypes receive the same equilibrium outcome function (x, s) 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 on 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 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

$U_i(k) = \underline{x}_i\,(k)\;k - \underline{s}_i\,(k) \geq \underline{x}_i(\hat{k})k - \underline{s}_i(\hat{k}) = U_i(\hat{k}) + \underline{x}_i(\hat{k})(k - \hat{k}\,)$

$U_i(\hat{k}) = \underline{x}_i\,(\hat{k})\,\hat{k} - \;\underline{s}_i\,(\hat{k}) \geq x_i\,k(\hat{k}) - \underline{s}_i(k) = U_i(k) + \underline{x}_i(k)(k - \hat{k})$

Hence $\underline{x}_i(k)(k - \hat{k}) \geq \underline{x}_i(k)(k - \hat{k})$ 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 $\underline{x}_i(k) \geq \overline{x_i(\hat{k})}$ for all $k < \hat{k}$.

Together with the continuity of these functions, this implies that $\underline{x}_i(k) = \overline{x_i(k)}$ (and $\underline{s}_i(k) = \overline{(s_i)}$) 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* $\Theta_i \times \ldots \times \Theta_N$ to $(p_i, Q_i)N_{i=1}$) *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 measure zero.*

Proof: Lemma 2 implies that all equilibria in the unconstrained scoring auction are equilibria in the constrained auction (they differ at most by a measure zero). To prove that all equilibria in the constrained auction are also equilibria in the unconstrained auction, note that bidders' preferences over strategies (and therefore outcomes x and s) are entirely determined by their pseudotypes (refer to (3) if needed). Therefore, if a strategy is a best response when a bidder is constrained to adopt a strategy based on his pseudotype, this strategy is again a best response for all types θ consistent with that pseudotype. QED

Most theoretical analyses of scoring auctions have implicitly or explicitly taken advantage of the sufficient statistics property of scoring auction 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,[15] this

[15] *In the one dimensional case as in Che (1993), this is not so much a question of sufficient statistics (there is no reduction of the dimensionality of private information per se) as a simple change of variables.*

result is not trivial for environments where types are multi-dimensional. 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 supplier's beliefs independent of their types. (Multi-unit auctions offer an example of multi-dimensional 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 scoring auctions 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. Indeed, 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. Theorem 2 extends a result obtained by Che (1993) on the utility equivalence between the first and second scoring auction when types are one-dimensional and the scoring rule corresponds to the buyer's true preferences, i.e. $\theta(Q) = v(Q)$.

Theorem 2 (Expected Utility Equivalence). *Any two scoring auctions that:*
(a) use the same quasi-linear scoring rule,
(b) use the same allocation rule $x_i(k_i, k_{-i})$, $i = 1, \ldots, N$, *and*
(c) yield the same expected payoff for the lowest pseudotype k_i, $i = 1, \ldots, N$.
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

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

where ESS(ki) 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 byLemma 2. Hence, (b) and (c) implies that $U_i(k_i)$ 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 si. 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 si 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. Indeed, in that case, the social surplus associated with a bidder of type θ_i is his pseudotype k_i, so $EES(k_i) = k_i$.

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 multi-dimensional 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. Concluding remarks

Auctions with multi-dimensional private information are notoriously tricky to analyze. In this paper, we exploit the simple property that multi-attribute auctions with scoring rules reduce the multi-dimensional decision problem into a one-dimensional variable, the score. This score is used both for deciding whom to award the contract and the resulting contractual obligations of the bidders.

We have exploited this idea in two ways. First, 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. Second, we have derived a new expected utility equivalence theorem for scoring auctions. Any two scoring auctions that use the same quasi-linear scoring rule and have the same allocation rule generate the same expected utility for the buyer, modulo one additive constant. Both results extend existing theories of scoring auctions.

The "sufficient statistics approach" presented here greatly facilitates the practical analysis of existing scoring auctions since standard techniques and results of auction theory can be applied to scoring auctions (Theorem 1 and Corollary 1). Nevertheless, it is likely to be less helpful to answer questions about the optimal choice of a scoring rule. The reason is that the distribution of pseudotypes is endogenous to the choice of a scoring rule. Therefore choosing the best scoring rule comes down to maximizing the expression for the buyer's expected utility (6), over the set of distributions of pseudotypes compatible with a scoring rule. We are skeptical that any useful progress on this question can be achieved using this approach. In work in progress, we adapt and extend techniques used in the multi-dimensional screening literature to study the question of buyer optimal scheme in multi-attributes auctions, including the question of the optimal scoring rule.

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

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

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

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

[6] 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.

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

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

[9] 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.

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

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

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

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

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

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

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

[17] Riley, John G. andWilliam Samuelson (1981), Optimal Auctions, American Economic Review, 71, 381-392.

[18] Wilson, Robert (2002), The Architechure of Power Markets, Econometrica, 70(4), 1299-1341.

CReA Catizone Randi e Associati Srl
www.crea-design.it