2003s-31

Implementing the Optimal Auction

Bernard Caillaud, Jacques Robert

Série Scientifique
Scientific Series

Montréal
Juin 2003

CIRANO
Centre interuniversitaire de recherche en analyse des organisations

ISSN 1198-8177

Implementing the Optimal Auction*

Bernard Caillaud†, *Jacques Robert*‡

Résumé / *Abstract*

Dans un cadre du modèle d'enchères avec des valeurs privées indépendantes, nous proposons un jeu, ayant une interprétation économique simple, qui permet de mettre en oeuvre les enchères optimales même quand le vendeur ignore les distributions des volontés à payer des différents soumissionnaires. Dans cette procédure robuste (*detail-free*), une enchère au deuxième prix est organisée et le gagnant de cette enchère propose un paiement au vendeur; ce paiement peut alors être contesté par un autre soumissionnaire qui connaît la distribution de l'évaluation du gagnant.

Mots clés : Enchères, design de mécanismes, doctrine de Wilson.

In a general framework with independent private values of the bidders, we propose a game, with a simple economic interpretation, that allows implementing the optimal auction outcome when the seller ignores the distributions of the different bidders' valuations. In this robust or detail-free implementation procedure, a second-price auction is organized and the winner volunteers a payment to the seller; this payment can then be challenged by another bidder who knows the distribution of the winner's valuation.

Keywords: *Auction, mechanism design, Wilson doctrine.*

---

* We have benefited from comments by seminar participants at the Econometric Society North American Winter Meeting and European Meeting, PURC University of Florida, University of Wisconsin. The authors also wish to acknowledge the hospitality of IDEI in Toulouse where they started this research. The second author acknowledges the financial support of CRSHC, CIRANO and the Bell University Labs.

† CERAS-ENPC (CNRS), 48 Boulevard Jourdan, 75014 Paris, France, et CEPR, London, United Kingdom. E-mail: caillaud@enpc.fr.

‡ Service d'enseignement des technologies de l'information, HEC Montréal, 3000 ch. Côte-Ste-Catherine, Montréal, Québec, Canada, H3T 2A7, et CIRANO. E-mail: jacques.robert@hec.ca.

1 Introduction

Since the late seventies, the mechanism design literature has been successful is determining the form and properties of desirable institutions in situations where informational problems arise.[1] As a leading example, Myerson [1981] characterizes the revenue-maximizing auction when potential buyers have private and independent valuations for the good on sale and are risk-neutral.

The mechanism design approach has however been criticized on the following grounds. First, optimal institutions, as derived by this approach, are often much more complex than real-life institutions. For example, the optimal auction mechanism with ex ante asymmetric participants or in the non-regular case turns out to be much more complicated than a simple first or second-price auction. Second and most importantly, the mechanism design approach is said to be *informationally demanding:* the design of optimal institutions requires an unrealistic degree of knowledge concerning details of the economic environment. For example, even in the symmetric and regular framework, the optimal auction requires the appropriate choice of the reserve price that strongly depends upon the knowledge of the prior distribution of tastes in the population of potential bidders. In repeated environments, simulation-based estimation methods can help figure out the objective distribution of tastes in a stable population,[2] although they usually miss the strategic dimension that is precisely due to this repeated-game setting. But for unusual auction situations (defense procurement, auction for monopoly franchises...), the crucial data are missing. The spirit of the so-called Wilson Doctrine has then been to ask for robustness, that is to try to reach detail-free conclusions within the mechanism design approach.

In this paper, we take these critiques seriously and we look for a detail-free implementation procedure of the optimal auction in Myerson's framework. We propose

[1]See e.g. the recent experiences of spectrum auctions in the US (Cramton [1995], McAfee-McMillan [1996]) and the UK (Binmore-Klemperer [2002], Klemperer [2002]).

[2]See among others Laffont-Ossard-Vuong [1995] and Donald-Paarsch [1996]. McAfee-Quan-Vincent [1996] uses econometric estimations to calculate the optimal reserve prices in housing auctions.

and analyze a game that could be designed by an ignorant seller who has no information not only on the bidders' valuations but also on the objective distributions of these valuations. Assuming that, for each bidder i, there exists one other bidder who knows the objective distribution of bidder i's valuation, the game relies on the possibility of challenges, in the spirit of Moore-Repullo [1988] and Glazer-Ma [1989], and it provides a way to implement the revenue-maximizing auction with an easy economic interpretation.

We start from the well-known implemenation of the optimal revenue-maximizing auction through an ascending-price auction, where the winner is asked to pay a price according to a pre-specified formula conditional on the winning bid. In the spirit of Bulow-Roberts [1989], this price can be viewed as a monopolistic pricing decision against the winner of the auction, where the monopolist's cost is determined by the winning bid in the auction. In order to do so, the seller needs to be able to compute the optimal monopoly price for each participant; but she cannot do so when she is ignorant, that is, when she has no knowledge of the distribution of private information.

So, we look for a mechanism that is independent of the parameters of the problem, but nevertheless achieves the maximal revenue for the seller. We simply assume that it is common knowledge that one bidder $j \neq i$ knows the distribution of bidder i's valuation.[3] Formally, we are looking for *game forms that are independent of the distribution of types* and which induce as an equilibrium the desired optimal allocation: we call this a *universal implementation procedure* of the optimal auction.

The game we propose has the following relatively simple structure: (i) an ascending-price auction is organized and the winner and the winning bid are made public information; (ii) the winner of this auction then volunteers a payment to the seller that is also publicly disclosed; (iii) another participant who knows the distribution of the winner's type is then allowed to challenge the volunteered price, by making a

[3] Our strong results enable us to dispense with an explicit Bayesian setting where bidders' private information not only concerns their own valuation but also their first-degree beliefs on each others' valuations (see Harsanyi [1968] and Mertens-Zamir [1985]).

take-it-or-leave-it offer at a higher price to the ascending-auction winner. This game admits an equilibrium that generates, for all distributions of bidders' valuations, the revenue-maximizing auction outcome. In this equilibrium, instead of charging a monopolistic price against the winner, the seller delegates this right to another bidder who has the relevant piece of information in order to solve the monopolistic pricing problem. The threat of having to face this monopoly price anyway induces the winner of the ascending auction to volunteer this payment, so as to avoid a fee involved in case of a challenge.

The standard characterization of the revenue-maximizing auction does not require any knowledge of the model by the bidders, since dominant strategy implementation is possible (see Mookherjee-Reichelstein [1992]), but it assumes a lot of knowledge by the seller. Our model reverses this assumption: the seller can maximize revenues without any knowledge of the distributions of types among the bidders, by relying on the fact the bidders know the distributions of types of rival bidders. In fact, it is only necessary that one bidder knows the distribution of another bidder's type. This alternative information structure seems relevant for firms within the same sector, whose common experience may have brought some knowledge about their rivals' characteristics.

The paper is organized as follows. Section offers a quick review of Myerson's IPV model and of the implementation of the optimal auction therein. Section 3 deals with the case of an ignorant seller and presents our main result in terms of universal implementation. Section 4 discusses some extensions and limitations of the result. In particular, it shows how to extend the game in the non-regular framework and it explains how to strengthen our weak implementation result in order to get unique implementation of the optimal auction outcome.

2 Revenue-maximizing auction in the benchmark model.

We consider the classical auction setting with Independent Private Values (IPV), as analyzed in Myerson [1981]. A risk-neutral seller wants to maximize her revenue from the sale of an indivisible good for which her valuation is known and normalized to 0. There are n risk-neutral potential buyers, each with private information on his own valuation v_i for the good. Valuations v_i, $i = 1, ...n$ are independently drawn from continuously differentiable distributions $F_i(.)$, with densities $f_i(.)$ and full support $[\underline{v}_i, \bar{v}_i]$. Moreover, we will concentrate on the so-called regular case where each bidder's *virtual valuation function* is monotone increasing: formally, we assume that, for any i,

$$J_i(v_i) \equiv v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$$

is increasing in v_i.[4] These distributions are common knowledge among all agents.

Before addressing the auction design problem, let first consider the corresponding pricing problem of a monopolist with unit cost b facing demand $[1 - F_i(p)]$. Let first $P_i(b)$ denote the optimal monopoly price:

$$P_i(b) \equiv \arg\max_p \ \{(p - b)[1 - F_i(p)]\}. \tag{1}$$

In our regular framework, $P_i(b)$ is single-valued and invertible and we have: $P_i^{-1}(v_i) \equiv \{b | v_i = P_i(b)\} = J_i(v_i)$.

The optimal revenue-maximizing auction has been proved to be such that the good be transferred to bidder i with valuation v_i if:

$$J_i(v_i) \geq \sup\{\sup\{J_h(v_h), h \neq i\}, 0\}, \tag{2}$$

that is to the bidder with the highest non-negative virtual valuation. In case of

[4] As explained in Section 4, our results can easily be extended to the general, i.e. non-regular, case.

equality, any tie-breaking rule can be used.

The outcome of the revenue-maximizing auction can be implemented (in dominant strategies) using a direct revelation mechanism, with the appropriate payment function: the winner pays the lowest valuation that would have made him win. Although in the symmetric regular case, the revenue equivalence theorem shows that a first-price, or a second-price, or else an ascending auction, all with reservation price $J^{-1}(0)$, yield the optimal outcome, it is difficult in more general cases to find a natural indirect mechanism, such as a standard auction, that implements the optimal auction.

Bulow-Roberts [1989], however, shows that the seller's problem is formally equivalent to a third-degree price discriminatory monopoly problem with capacity constraint. Facing perfectly identifiable demand functions $[1 - F_i(p)]$ for $i = 1, 2, ...n$, such a monopolist should optimally sell to the buyers with the highest marginal revenue. In the regular setting, the revenue functions $R_i(Q) \equiv QF_i^{-1}(1-Q)$ are strictly quasi-concave, or equivalently when the functions $J_i(.)$ are all strictly increasing, the optimal monopolistic policy is to compare marginal revenues $R_i'([1 - F_i(v_i)])$, that is $J_i(v_i)$ for $i = 1, ...n$, and to allocate the good to the buyer with highest marginal revenue.

This interpretation of Bulow-Roberts [1998] translates into a simple implementation procedure of the optimal outcome. Bidders participate in an ascending-bid auction with initial bid starting at zero. The winner is the last participant to drop out (or drawn among the last participants to drop out, with equal probabilities). The "winning bid" is the highest bid for which the number of active bidders is larger or equal to 2. If at least one bidder participates and bidder i wins at winning bid $b \geq 0$, he gets the good and pays $P_i(b) = J_i^{-1}(b)$ to the seller; otherwise, the good is not sold.

The key feature of this procedure is that it is a (weakly) dominant strategy for each bidder to bid his virtual valuation. For i of type v_i it is a dominant strategy to stay active whenever b is such that $v_i > J_i(b)$ and drop out whenever $v_i < J_i(b)$, i.e. to stay active until the price reaches the value $J_i(v_i)$. Moreover, one can easily check that the payments are identical to the ones specified in Myerson [1981] (page 69, expression (6.8)). We summarize this in the following proposition.

Proposition 1. *: (Bulow-Roberts) The optimal auction can be implemented in dominant strategies through an ascending-bid auction with payments given by the functions $P_i(.)$.*

Game 1 follows a common index procedure, "a common clock": the index increases until only one participant remains, thereby determining the winning bid and the winner. This bid only serves as a basis to determine the actual payment by the winner, based on the winner's virtual valuation function. The final allocation and payment rules can be seen as resulting from a monopoly pricing decision against the winning bidder where the monopolist's cost is determined by the winning bid. This payment could be extracted in a take-it-or leave-it offer game by any intermediary facing a cost equal to the winning bid and knowing the prior distribution of valuations for the winner. We exploit this idea below.

3 Universal implementation in the regular case

The previous implementation game is not informationally demanding for bidders. The dominant strategy implementation procedure puts no requirement on the knowledge bidders have upon each others and upon the market conditions as a whole. The auctioneer, however, needs to be able to maximize $(p - b)\,[1 - F_i(p)]$ for all i and b. Basically, she needs to know the distributions of valuations $F_i(.)$ for each bidder i. Such a requirement may be unrealistic. In many cases, the auctioneer has little knowledge about the parameters of the market, at least compared to the knowledge actual participants in the market have from long years of practice and competition within the market.

In this section, we propose a "universal", or detail-free mechanism which the auction designer can set up independently of her knowledge of the bidders' tastes.

Doing so, we completely reverse the informational requirement by assuming that participants in the auction are better informed than the auction designer upon the distribution of tastes among bidders. More precisely, we will assume that the distributions $F_i(\cdot)$ are common knowledge among bidders.[5] On the other hand, we do not need specify what are the seller's prior, if any, on these distributions, since the mechanism we present implements the optimal auction whatever the distributions $F_i(.)$. In this mechanism, competition among participants is used as a device to induce participants to implicitly reveal the optimal sale price.

It is important to note there are many mechanisms which can force bidders to reveal to the auctioneer the information necessary to implement the optimal auction. Clearly, selecting among the possible games raises methodological issues. Following the seminal work by Maskin [1977], many studies have been done on Bayesian Nash implementation.[6] The basic conclusion of this literature is that information which is common knowledge among agents can be revealed at no cost to the principal. One way is to ask agents to reveal simultaneously their joint information and, if they fail to send the same information, to impose infinite penalties on them. We propose an easily interpretable procedure where the seller organizes the auction in a rather conventional fashion and let competitive pressure works in her favor to extract the maximum expected rents. The game is meant to be intuitive and practical.[7]

Let us consider the following game Γ:

1. Bidders participate in an *ascending-bid auction,* starting from bid zero. The winner i and the winning bid b are determined as in a standard ascending bid auction and the outcome (i, b) is publicly disclosed.

2. If i wins, he *volunteers a price* p which is also publicly disclosed.

[5]As will appear, we could simply assume that it is common knowledge that for each i, there is another bidder $c(i) \neq i$ who knows $F_i(.)$.

[6]See Moore [1992] for a survey on implementation under complete information and Palfrey [1992] for a survey on Bayesian implementation.

[7]Pursuing a similar quest for simple implementation procedures, Glazer-Ma [1989] have analyzed implementation with the possibility of challenges. The procedure we propose also incorporates the possibility of challenges.

3. If i has won at winning price b and has volunteered a price p, another agent is randomly designated and can *challenge* p by proposing a take-it-or-leave-it offer at a higher price $q > p$, which i can accept or refuse; payments are as follows:

   - if p is unchallenged, then i receives the unit and pays p to the seller.
   - if a challenge $q > p$ is accepted by i, then i receives the unit, pays a fee Δ to the seller and pays q to the challenger. The challenger pays p to the seller.
   - if a challenge $q > p$ is rejected, the seller keeps the good, receives a fee Δ from i and a payment $p - b$ from the challenger.

In contrast with the ascending-bid auction with pre-specified payments described in Proposition 1, the winner is asked to volunteer a payment and the game is designed so that the winner has indeed an incentive to volunteer the optimal auction payment. He is disciplined in doing so by the possibility of challenges from an informed outside agent, who is able to compute the optimal auction payment based on his knowledge of the prior distribution of the winner's valuation and on the public information about the winner and the winning bid.

More precisely, we have the following central result:

Theorem 1. *: Game Γ universally implements the optimal auction outcome. That is, the following strategies and beliefs form a perfect Bayesian equilibrium of Γ in which the seller obtains the same revenues as if she knew the distribution $F_i(.)$:*

- *bidder i of type v_i exits at bid $J_i(v_i)$;*
- *when winning at bid b, bidder i volunteers $p = P_i(b)$;*
- *after he has won at winning bid b and volunteered price p, bidder i of types v_i accepts all challenges such that $q \leq v_i$ and refuses others; challengers update their prior $F_i(.)$ conditional on the event $\{v_i \geq P_i(b)\}$, they propose a price $q = P_i(b)$ if and only if $p < P_i(b)$ and abstain otherwise.*

Proof. After winning at b, bidder i should accept a challenge if and only $v_i - q - \Delta \geq -\Delta$, hence if and only if $q \leq v_i$.

A challenger looks for the challenge price q that maximizes his expected profit, where the expectation is taken with respect to his posteriors on the winner's type v_i after observing that i wins at winning bid b and volunteers price p. Suppose that the challenger's posteriors are given by the Bayesian updating of his priors, conditional on the event $\{v_i \geq P_i(b)\}$.[8] In this case, if indeed $b < \bar{v}_i$,[9] the program for the optimal challenging price is given by:

$$\max_{q \geq 0} \left[-(p-b) + (q-b) \inf \left\{ \frac{1 - F_i(q)}{1 - F_i(P_i(b))}; 1 \right\} \right]. \tag{3}$$

where the infimum corresponds to the challenger's probability assessment that the challenge q will be accepted. The optimum challenge price then corresponds to: $q = P_i(b)$. This is indeed the challenge price if it is strictly larger than p; otherwise, the value of the program above is strictly negative when restricted to $q > p$, and p is not challenged.

Suppose that bidder i wins at winning bid b. Proposing a price $p < P_i(b)$ generates a challenge $q = P_i(b)$. It cannot be an equilibrium strategy for bidder i to encourage a challenge that he will accept, since he could have immediately proposed $p = P_i(b)$, thereby avoiding to pay the fee Δ. Therefore, the maximal profits that bidder i of type v_i can obtain after winning at winning bid b are equal to: $\sup \{v_i - P_i(b); -\Delta\}$. These profits are non-negative if and only if $b \leq J_i(v_i)$ or $v_i \geq P_i(b)$. It follows that by dropping out in the ascending auction precisely at bid $J_i(v_i)$ and by proposing the unchallengeable price $p = P_i(b)$ when he wins, bidder i of type v_i maximizes his expected gains.

Finally, the beliefs that have been posited are actually Bayesian consistent. They

[8]Note that posteriors do not depend upon the price proposal p, although this proposal could serve as a signalling device for the winner.

[9]If $b > \bar{v}_i$, the beliefs proposed in the text are not compatible with the history of the game. One can fix beliefs to be concentrated on $\{v_i = \bar{v}_i\}$, resulting in $q = \bar{v}_i$. The same challenge price should be considered if $b = \bar{v}_i$.

are deduced from Bayes rule after history (i, b). The fact that p does not induce a further updating is consistent with the fact that among all types $v_i \geq P_i(b)$, it is a full pooling equilibrium to volunteer the same price $p = P_i(b)$; therefore, p does not convey any additional information about the winner's type. □

To summarize, *Game* Γ *admits a perfect Bayesian equilibrium that implements the optimal auction outcome whatever the actual distribution of types* $F_i(.)$. In this equilibrium, bidders drop out at their virtual valuation $J_i(v_i)$, the winning bid therefore coincides with the highest second virtual valuation. When they win, bidders submit the lowest unchallengeable price to avoid paying the fee Δ; this price corresponds to the highest valuation that would have still enabled them to win the ascending auction, $P_i(b)$, i.e. the corresponding optimal auction payment.

The challenger's program (3) can be viewed as the program of a monopolist facing demand $Q = 1 - F_i(q)$ and unit cost b (for $q \in [P_i(b), \bar{v}_i]$). So, the game relies on the existence of other agents who have priors $F_i(.)$ on i's valuation and can challenge the volunteered payment when it is lower than the corresponding monopoly price. The seller in fact delegates the monopoly pricing decision discussed in the previous section to one better informed intermediary as an off-equilibrium threat that serves as a disciplining device to induce the correct price proposal by the winner. Note that for the system to work, the challenging agent need not be interested in purchasing the good for himself, he need only be motivated by the possibility of a profitable arbitrage. Note also that the challenge could alternatively be organized as a competition game between potential challengers where only the largest challenge price is considered, or as a (common value) auction between potential challengers so as to win the right to make a take-it-or-leave-it offer to the winner of the first auction.

4 Extensions and discussion

4.1 Implementation in the non-regular case

The analysis is more technical when the functions $J_i(.)$ are not monotone increasing and we only sketch it here. (1) now defines a set of optimal monopoly prices. Let $t_i(b)$ and $r_i(b)$ denote respectively the lowest and the largest elements in $P_i(b)$ and let $H_i(v_i)$ be the (unique) value b such that $v_i \in [t_i(b), r_i(b)]$.[10] The function $H_i(.)$ is obtained by ironing out the function $J_i(.)$ so that it is weakly increasing. The optimal revenue-maximizing auction now requires to transfer the good to bidder i with valuation v_i if:

$$H_i(v_i) \geq \sup \left\{ \sup \left\{ H_h(v_h), h \neq i \right\}, 0 \right\}. \tag{4}$$

The maximal revenue can be obtained via an ascending bid auction, as in Proposition 1, except for the exact form of the payment function: if bidder i wins at winning bid b, he must now pay a weighted average of $t_i(b)$ and $r_i(b)$ whose weights depend upon the number of ties at the winning bid.[11]

The game form Γ, consisting of an ascending-bid auction followed by a challenge stage, provides a universal implementation of the revenue-maximizing auction in this case provided it is modified so as to mimic the more complex structure of the payments in the revenue-maximizing auction. Formally, a winner of the ascending-bid auction is now required to volunteer two prices p_1 and p_2, with.$p_2 \geq p_1$, among which the seller randomizes so that the expected price corresponds to the appropriate weighted average price where p_1 and p_2 are treated as $t_i(b)$ and $r_i(b)$ respectively. Another bidder can then challenge the price randomly drawn with a strictly higher price.

[10] Under the assumption that $F_i(.)$ is continuously differentiable, if $b < b'$, then $r_i(b) < t_i(b')$ and $\lim_{b' \downarrow b} t_i(b') = r_i(b)$. Hence the expression $H_i(v_i)$ is well-defined.

[11] The precise form of the payments in case of simultaneous exits by several bidders at the winning bid can be found in Myerson [1981]. If i is the sole winner and m bidders dropped out at b, bidder i must pay $p = \frac{1}{m+1} t_i(b) + \frac{m}{m+1} r_i(b)$. If i wins as the result of a random draw among several bidders, he must pay $t_i(b)$.

In this modified game Γ, bidder i of type v_i participates up to the point where the ascending price reaches $H_i(v_i)$; when he wins at bid b, bidder i proposes $p_1 = t_i(b)$ and $p_2 = r_i(b)$; these price levels are also the appropriate challenge prices. This equilibrium induces the same expected profit for the seller as if she knew the distributions $F_i(.)$ and had organized the optimal auction in the more usual way.

Again, this mechanism provides a universal, or detail-free implementation procedure of the revenue maximizing auction outcome in a detail-free manner.[12] The appendix sketches the formal and technical analysis.

4.2 Strong implementation

Let us come back to the analysis of the regular case. Take any strictly increasing function $b_i(.)$ such that for all v_i, $b_i(v_i) \leq J_i(v_i)$. Consider the following strategies: bidder i of type v_i drops out in the ascending auction at bid $b_i(v_i)$; if he wins at b, he volunteers a price $p = b_i^{-1}(b)$ (if $b > b_i(\underline{v}_i)$, and $\underline{v}_i$ otherwise) and beliefs on i's type following history (i, b, p) correspond to the updating of priors conditional on the event $\left\{v_i \geq b_i^{-1}(b)\right\}$. It is easy to see that Program (3) is still valid after replacing $P_i(b)$ by $b_i^{-1}(b)$. Under the regularity assumption and since $P_i(b) \leq b_i^{-1}(b)$, the optimal challenge price is then obtained as a corner solution at $q = b_i^{-1}(b)$ (or $\underline{v}_i$). The argument is then similar to the one developed above.

The proposed strategies and beliefs therefore constitute another perfect Bayesian equilibrium of Game Γ. Game Γ has indeed a continuum of equilibria. All the equilibria rely on a less aggressive behavior from bidders in the ascending auction compared to the equilibrium that implements the optimal auction; the challenging stage then induces higher prices, for a given winning bid b. Multiplicity of perfect Bayesian equilibria should not be a surprise. One can construct "unreasonable" equilibria by allowing potential challengers to have unreasonable out-of-equilibrium beliefs leading to high challenges q.

[12] The mechanism is less appealing because it is less simple than in the regular case. But the payment structure in the optimal auction is itself quite complex, so that one should not expect to find a simpler implementation procedure.

Theorem 1 is therefore a weak implementation result, but we can exploit the specific properties of the various equilibria to strengthen our point. A first point that can be made is that the equilibrium that implements the revenue-maximizing auction is not any equilibrium. As a limit equilibrium, it has a focal point property; moreover it is the unique equilibrium where beliefs on and off equilibrium are concentrated on $\{v_i$, such that $v_i \geq p\}$.

To go one step further, we will restrict beliefs out-of-the-equilibrium path and impose a refinement of the perfect Bayesian equilibrium concept that generates a unique equilibrium; then, we can argue that the equilibrium in Theorem 1 can be interpreted as the only reasonable equilibrium in Game Γ. For this purpose, we apply here the concept of explicable equilibrium due to Reny [1992].

The idea underlying the notion of "explicable" equilibrium is that when a deviation is detected, the other participants must interpret this deviation not necessarily as an irrational move on the part of the deviator but, whenever possible, as the result of some confusion over which equilibrium is being played. Whenever possible, a deviation should be interpreted as a best-response to some other equilibrium. Formally, let B be some common standard of behavior and let π be some strategy profile in B. Now suppose that for $i \neq j$, (a) an information set, h, for j is inconsistent with i's strategy profile; (b) h is both consistent with i using the distinct pure strategies s and s'; (c) s is a best response relative to B while s' is not, where s is a best response to B if there exists an element $\gamma \in \mathrm{co}B$ such that s is a best response against γ. Then, according to the notion of explicable equilibrium, if h is reached, j's reference about i's strategy should put zero probability on s' being played.

In the context of our game, the notion of explicable equilibrium has sufficient bite if we set B to be the set of Perfect Bayesian Equilibria. It can eliminate all equilibria characterized by a $b_i(.)$ function such that $b_i^{-1}(b) > P_i(b)$ with strictly positive measure. Consider such an equilibrium and suppose an information h is reached where i wins the auction at some price b and, as in the previous subsection, offer $p = P_i(b) < b_i^{-1}(b)$. We know that the strategy s used by bidder i in the implementation equilibrium consists of bidding according to $J_i(.)$ and offering $p = P_i(b)$, and it is a best-response to the set of perfect Bayesian equilibria. However, any strategy s' for i that consists of bidding according to a $b_i(.)$ function such that $b_i^{-1}(b) > P_i(b)$ with positive probability, and of offering $p(b) = P_i(b)$ is not a best-response relative to the set of perfect Bayesian equilibria: $p(b) = P_i(b)$ lies off the equilibrium path and either this price is challenged, in which case i could have profitably proposed a higher p, or it is not challenged, in which case i did not followed an optimal bidding strategy since only types in $\left\{v_i \geq b_i^{-1}(b)\right\}$ bid up to b and may win at price b while it would have been profitable for $v_i \in \left(P_i(b), b_i^{-1}(b)\right)$ to win at price p. Hence, the challengers' inference about i's play must put zero probability on all strategies where i bids up to b only if $v_i \geq b_i^{-1}(b) > P_i(b)$, for any candidate $b_i(.)$-function. Since offering $p = P_i(b)$ is a strictly dominated strategy if $v_i < P_i(b)$, the challengers' inference must then be that $v_i \geq p = P_i(b)$. This again nails down beliefs off-the-equilibrium path. Following the discussion in the previous subsection, it selects a unique perfect Bayesian equilibrium outcome, the equilibrium that implements the optimal auction outcome.

The conclusion of this analysis is that the equilibrium that leads to the optimal auction outcome has very special features that suggest that it constitutes a reasonable outcome of our simple implementation procedure.

5 Conclusion

We have proposed a game form that implements the optimal auction in a relatively simple way without requiring extensive knowledge on the part of the auctioneer. The key features of Game Γ are that the winner of the auction does not pay the winning price of the auction; she pays some price that is determined afterwards through some well-defined bargaining process. This is not very different from some current practices. Often the competitive process is meant only to identify a winner, the actual price and contract conditions are bargained afterwards between the interested parties.

The result of this paper has many limitations. We have restricted our attention to the case of private and independent values and the case of risk-neutral bidders. We also assume that one and only one unit is on sale, we do not consider how the logic here applies to multi-auctions with multi-unit demands. Our last concern relates to the repetition of these auctions. The presumption that participants are well-informed about the distributions of valuations of other participants reflects the notion that they all share a common experience and that these auctions are often repeated. If this is true, then collusion may arise: as a rule buyers may agree never to challenge each other. So, we view this paper as a first step in the pursuit of finding practical implementation procedures for optimal auctions.

References

[1] Binmore, K. and P. Klemperer [2002]: "The biggest auction ever: the sale of the British 3G Telecom Licences," *Economic Journal*, **112 (478)**, March.

[2] Bulow, J. and J. Roberts [1989]: "The simple Economics of Optimal Auction," *Journal of Political Economy*; **97**; 1060-1090.

[3] Cramton, P. [1995]: "Money out of Thin Air: The Nationwide Narrowband PCS Auction", *Journal of Economics and Management Strategy,* **4**, 267-343.

[4] Donald, S.G. and H.J.Paarsch [1996]: "Piecewise pseudo-maximum likelihood estimation in empirical models of auction", *International Economic Review*; **34**; pp.121-148.

[5] Fudenberg, D. and J. Tirole [1991]: "Perfect Bayesian Equilibria and Sequential Equilibria", *Journal of Economic Theory*, **53**, 236-260.

[6] Glazer, J. and C.-A. Ma [1989]: "Efficient Allocation of a "Prize"- King Solomon's Dilemma", *Games and Economic Behavior,* **1**, 222-233.

[7] Harsanyi, J. [1968]: "Games with Incomplete Information Played by Bayesian Players", I-II-III, *Management Sciences,* **14**, 159-182, 302-334, 486-502.

[8] Klemperer, P. [2002]: "What really matters in auction design," *Journal of Economic Perspectives*, **16(1)**, 169-190.

[9] Laffont, J.-J., H. Ossard and Q. Vuong [1995]: "Econometrics of First-Price Auctions", *Econometrica,* **63(4)**, 953-980.

[10] McAfee, P. and J. McMillan [1996]: "Analyzing the Airwaves Auction", *Journal of Economic Perspective,* **10(1)**, 159-176.

[11] McAfee, R.P., D.Quan and D. Vincent [1996]: "How to set Minimum Acceptable Bids", *mimeo.*

[12] Maskin, E. [1977]: "Nash Equilibrium and Welfare Optimality", mimeo.

[13] Mertens, J.-F. and S. Zamir [1985]: "Formulations of Bayesian Analysis for Games of Incomplete Information", *International Journal of Game Theory,* **14(1)**, 1-29.

[14] Mookherjee, D. and S. Reichelstein [1992]: "Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules", *Journal of Economic Theory*, **56**, 378-399.

[15] Moore, J. [1992]: "Implementation, Contracts and Renegotiation in Environments with Symmetric Information", in *Advances in Economic Theory: the 6th World Congress*, ed. J.J.Laffont, Cambridge University Press.

[16] Moore, J. and R. Repullo [1988]: "Subgame Perfect Implementation", *Econometrica*, **56**, 1191-1220.

[17] Myerson, R. [1981]: "Optimal Auction Design," *Mathematics of Operational Research*; **6**; 58-63.

[18] Palfrey, T. [1992]: "Implementation in Bayesian Equilibrium: the Multiple Equilibrium Problem in Mechanism Design", in *Advances in Economic Theory: the 6th World Congress*, ed. J.J.Laffont, Cambridge University Press.

[19] Reny, P. [1992]: "Backward Induction, Normal Form Perfection and Explicable Equilibria," *Econometrica*; **60**(3); 627-650.

A Appendix: the general case

In the non-regular case, we consider the following game:

1. Bidders participate in an ascending-price auction as in Game 2, with public disclosure of the winner and the winning bid (i, b).

2. If i wins, he volunteers two prices p_1 and p_2, with $p_2 \geq p_1$. If i dropped out at the winning bid (tie between potential winners), the seller sets $p = p_1$. Otherwise, the seller sets $p = p_1$ with probability $\frac{1}{m+1}$ and $p = p_2$ with probability $\frac{m}{m+1}$, if m bidders simultaneously dropped out at b.

3. The price p is publicly disclosed, as well as whether p is equal to p_1 or p_2. Another agent can then challenge p as in Game Γ.

We provide a sketch of the formal proof of the following result:

Theorem 2. *: Each Perfect Bayesian equilibrium outcome of the game in this Appendix is characterized by a profile of left-continuous non-decreasing functions* $\mu_i(.)$ *such that for all* $b \geq 0$,

$$\mu_i(b) \in \arg \max_{q \geq \mu_i(b)} \left[(q - b)[1 - F_i(q)] \right],$$

and of non-decreasing functions $\eta_i(.)$ *defined by:* $\eta_i(b) = \lim_{b' \downarrow b} \mu_i(b')$*, such that:*

(i) in the ascending auction, bidder i *bids up to* $b_i(v_i)$ *defined as the value of* b *such that* $v_i \in [\mu_i(b), \eta_i(b)]$*, if it is not smaller than* 0*, or drops out immediately at* 0 *otherwise;*

(ii) when he wins at winning bid b*, bidder* i *submits prices* $p_1 = \mu_i(b)$ *and* $p_2 = \eta_i(b)$*;*

(iii) there is no challenge.

For every such profile, there exists an equilibrium that yields the corresponding outcome.

Note that $t_i(.)$ and $r_i(.)$ satisfy the conditions in the theorem; moreover, the expected price is equal to the price in the optimal auction for $p_1 = t_i(b)$ and $p_2 = r_i(b)$. Hence the general implementation result.

A.1 Sufficiency

We first prove that for any profile of left-continuous increasing functions $\mu_i(.)$ and associated profile of functions $\eta_i(.)$, the following strategies sustain the corresponding outcome as an equilibrium. Bidder i bids up to b such that $v_i \in [\mu_i(b), \eta_i(b)]$. If i wins the initial auction at value b, he volunteers payments $p_1 = \mu_i(b)$ and $p_2 = \eta_i(b)$. These prices are unchallenged. However, if $p = p_1 < \mu_i(b)$, it is challenged and the challenger offers $q = \mu_i(b)$, and if $p = p_2 < \eta_i(b)$, the challenger offers $q = \eta_i(b)$.

Consider the challenger's beliefs when i wins at winning bid b and m bidders apart from i dropped out at b. If $p = p_1$, it must be that either i also dropped out at b and was randomly selected with probability $\frac{1}{m+1}$ or that i dropped out above b but $p = p_1$ was announced, which had probability $\frac{1}{m+1}$. The posterior beliefs therefore coincide with the Bayesian updating of prior beliefs conditional on $\{v_i \geq \mu_i(b)\}$. If $p = p_2$ is announced, posterior beliefs must be the Bayesian updating of prior beliefs conditional on $\{v_i > \eta_i(b)\}$.

Given these posterior beliefs, whenever $p = p_1$, it is sequentially rational to challenge p if $p < \mu_i(b)$ and offer $q = \mu_i(b)$ since $\mu_i(b) \in \arg\max_{q \geq \mu_i(b)}(q - b)[1 - F_i(q)]$. The case $p = p_2$ is similar..

Given the response of potential challengers, if i wins the initial auction at price b, his best strategy is to offer $p_1 = \mu_i(b)$ and $p_2 = \eta_i(b)$ unless $v_i < \mu_i(b) - \Delta$ in which case he offers a challengeable price, pays Δ and ends up not receiving the good. So let $\hat{b}$ be the value at which the last of all other participants drops out. Bidder i gets negative payoffs if he wins and $v_i < \mu_i(\hat{b})$ and he gets positive payoffs if he wins and $v_i > \eta_i(\hat{b})$ or if he drops out exactly at $\hat{b}$ and $\mu_i(\hat{b}) < v < \eta_i(\hat{b})$. His best-response is indeed to bid up to the value b such that $v_i \in [\mu_i(b), \eta_i(b)]$.

A.2 Necessity

Consider an equilibrium with $b_i(.)$ the bid (or exit price) functions, $p_{i1}(b, v_i)$ and $p_{i2}(b, v_i)$ the price proposals, and a decision rule which specifies whether to challenge i and, if so, at which price $q^i(b, p, \{p = p_k\})$. Let $Y_i = \{x | \exists v \text{ such that } b_i(v) = x\}$ denote the support of i's bids and $Y_{-i} = \cup_{j \neq i} Y_j$. Let $dG_i(b)$ denotes the equilibrium measure corresponding to the probability that i wins at the winning bid b.

Claim (i): *For dG_i-almost all $b \in Y_{-i}$, define the following: $\mu_i(b) \equiv \sup\{v_i, b_i(v_i) < b\}$ and $\eta_i(b) \equiv \sup\{v_i, b_i(v_i) \leq b\}$; then, in equilibrium $b_i(v_i) < b$ if $v_i < \mu_i(b)$, $b_i(v_i) = b$ if $\mu_i(b) < v_i < \eta_i(b)$, and $b_i(v_i) > b$ if $v_i > \eta_i(b)$.*

The claim asserts that in equilibrium, bid functions must necessarily be increasing when they are relevant.

Proof of Claim (i). Let $U_i(v_i, b_i)$ denote the expected payoffs of bidder i of type v_i who bids up to b_i. Consider two possible bids b_i and b_i', with $b_i > b_i'$ and $b_i \in \text{int}Y_{-i}$. If $U_i(v_i, b_i) \geq U_i(v_i, b_i')$, we can write:

$$0 \leq \int_{b_i'}^{b_i} u_i\left(v_i, p_1(b, v_i), p_2(b, v_i), b\right) dG_i(b),$$

where the index $u_i(v_i, p_1, p_2, b)$ stands for bidder i's expected payoff (in equilibrium) when he is of type v_i, wins at winning bid b and proposes prices (p_1, p_2). $u_i(v_i, p_1, p_2, b)$ may coincide with one of i's two equilibrium proposals, $v_i - p_1$ or $v_i - p_2$, or with an accepted challenge, $v_i - q^i(b, p, \{p = p_k\}) - \Delta$, or else with a rejected challenge, $-\Delta$; it is strictly increasing in v_i when it is non-negative. The inequality above implies that within (b_i', b_i), the integrand is non-negative on a set of positive dG_i-measure on which it is then strictly increasing in v_i. Therefore, for all $v_i' > v_i$, if $U_i(v_i, b_i) \geq U_i(v_i, b_i')$,

$$0 < \int_{b_i'}^{b_i} u_i\left(v_i', p_1(b, v_i), p_2(b, v_i), b\right) dG_i(b)$$

and then $U_i(v_i', b_i) > U_i(v_i', b_i')$. The result follows. $\square$

Claim (ii): *Suppose that i wins at winning bid $b \in Y_{-i}$, then along the equilibrium path, i offers $p_1(b, v_i) = \mu_i(b)$ and $p_2(b, v_i) = \eta_i(b)$ which are almost never challenged, for dG_i-almost all $b \in Y_{-i}$.*

Proof of Claim (ii). For $b \in \text{int}Y_{-i}$, let $k_1(b, v_i, i)$ (a similar argument exists for $k_2(b, v_i, i)$ corresponding to p_2) denote the effective price paid by i along the equilibrium path whenever he wins at bid b and $p = p_1$.

Let $k_1^*(b, i) \equiv \inf_{\{v_i \geq \mu_i(b)\}} k_1(b, v_i, i)$. We show that $k_1^*(b, i) = \mu_i(b)$ almost always, in the sense of dG_i. Suppose that $k_1^*(b, i) < \mu_i(b)$, i.e. there exists one $v_i \geq \mu_i(b)$ for which the effective price along the equilibrium path is less than $\mu_i(b)$. It cannot correspond to a rejected challenge since $v_i + \Delta > \mu_i(b)$; it cannot correspond to an accepted challenge or to an unchallenged price either, since the challenger would strictly benefit from challenging at a price q such that $k_1(b, v_i, i) < q < \mu_i(b)$, which would surely be accepted given posterior beliefs concentrated on $\{v_i \geq \mu_i(b)\}$.

Suppose now that there exists an interval $(b_0, b_0 + \epsilon) \subset Y_{-i}$ and $\delta > 0$ such that:

$$\forall v_i \in (\mu_i(b_0), \mu_i(b_0) + \delta), \forall b \in [b_0, b_0 + \epsilon), \quad k_1^*(b, i) > v_i.$$

A bidder i of type v_i in this right-neighborhood of $\mu_i(b_0)$ would have been strictly better off by submitting a bid strictly lower than b_0, which contradicts the definition of $\mu_i(.)$. Hence $k_1^*(b, i) = \mu_i(b)$.

From this, it follows that if there is a challenge in equilibrium, it must be accepted and must occur at price $q = \mu_i(b) - \Delta$. But this price cannot correspond to a rational challenge since a challenger knows that the winner i must have valuation $v_i \geq \mu_i(b)$ and would therefore accept a slightly higher challenge price with probability 1. So it is necessary that i winning at b volunteers $p_1 = \mu_i(b)$ and that there is no challenge whatever her type $v_i \geq \mu_i(b)$. $\square$

Claim (iii): *Along the equilibrium path, when i wins the auction at b for dG_i-*

almost every $b \in Y_{-i}$, then $\mu_i(b)$ must be such that:

$$\mu_i(b) \in \arg\max_{q \geq \mu_i(b)} (q-b)[1-F_i(q)] \tag{5}$$

Proof of Claim (iii). Given Claim (i), when a challenger faces the winner i at winning bid b and with $p = p_1$, he should have beliefs corresponding to the updating of prior beliefs conditional on $\{v_i \geq \mu_i(b)\}$ and should not find any profitable challenge. If there were a $q > \mu_i(b)$ such that $(q-b)[1-F_i(q)] > (\mu_i(b)-b)[1-F_i(\mu_i(b))]$, there would exist such a strictly profitable challenge against p_1. The same holds for $\eta_i(b)$. Hence the Claim. $\square$

This completes the proof of the theorem.