Biased Procurement

[Job Market Paper]

Leonardo Rezende*

November 19, 2002

Abstract

In a large procurement, sellers' products may differ in ways that matter to the buyer. If the buyer is willing to pay different prices to each seller, it may elect to introduce a bias that makes the preferred seller more likely to win. This paper studies the economic incentives that determine the gathering, disclosure and use of information about this bias. It is shown that a buyer able to commit to any mechanism will elect to bias the auction towards the preferred seller, but by an amount smaller than its original preference. A buyer unable to commit always finds in its interest to fully disclose its preferences to the sellers prior to the auction. The buyer's incentives to gather information about its own preferences depend on the ability to commit: while under commitment the value of information for the buyer is always positive, under no commitment the buyer will often elect not to learn about its preferences.

*Dept. of Economics, Stanford University. email: lrezende@stanford.edu. I gratefully acknowledge the support of a John M. Olin Dissertation Fellowship and a Melvin and Joan Lane Stanford Graduate Fellowship. Thanks to Pat Bajari, Tim Bresnahan, Vinicius Carrasco, Jonathan Levin, Paul Milgrom and Ilya Segal for valuable suggestions (and Jon for the title). All errors are mine.

1 Introduction

Comparing sellers is a complex task in a procurement auction. Bid proposals are lengthy documents, that differ in many dimensions besides price, such as technical specifications, time of delivery, amount and quality of service, supplier reputation, payment conditions, and aesthetic properties of the item to be supplied.

Such complexity introduces several important strategic issues that do not exist in a simple auction where the winner is always the lowest bidder. This paper presents a model that focuses on two of those issues: the fact that the preferences of the buyer may not be perfectly known by the sellers, and the fact that the procurement process involves extensive negotiations, including after the bidding and the selection of the winner is done.

In environments where it may not be clear what are the buyer's preferences across sellers, several interesting questions arise: will it be in the buyer's interest to reveal this preferences to the sellers? Will the buyer always follow its preferences in selecting the winner, or can it do better by distorting its choice? Are there circumstances where the buyer will prefer not to know about its own preferences?

On the other hand, the complexity of the products being traded require extensive negotiations and a large degree of flexibility in the auction protocol. There is a significant amount of exchange of information and negotiation both before and after the formal bidding step. In such environment, the specifics of the auctioning process are less important than recognizing the possibility of further negotiation afterwards. When further negotiations are possible, and the buyer cannot commit to a future course of action, sellers anticipate that and bid differently in the auction, trying to strengthen their bargaining position afterwards. Taking into account this possibility in the analysis leads to predictions that differ sharply from the ones obtained in the standard optimal auction framework, where sellers have full commitment power.

Both private and public procurement practice allow for considering other dimensions besides price in the evaluation of bids, and for ex-post negotiations. Gene Richter, former Chief Purchasing Officer of IBM, argues that simple lowest-bid auctions should not be used even for the simplest procurements:

> "There is nothing that a company buys that I can think of where only the price is important. There is a price, quality, delivering

> and technology issue in everything. So the purest use of auctions where the lowest bidder gets the business no matter what is terrible." (ISM, 2002)

Such issues are also relevant in government procurement. The Federal Acquisition Regulation (FAR), the set of rules that govern U.S. Federal Procurement, allow for "procurement by negotiation". In this protocol, bidders are evaluated according to a weighted average of pre-specified criteria. However, the weights do not have to be announced before bidding. This effectively allows a federal agency to procure without revealing its preferences. Furthermore, according to the FAR the buyer can engage in bargaining with a selected group of sellers over all aspects of their proposal, including price.

The importance of incorporating dimensions besides price in complex procurement auctions has been recognized in the literature. These other dimensions have either been modeled as choice variables in the suppliers side (Che, 1993) or as exogenous, private information on the supplier's side (Zheng, 2000). The current paper focuses instead on another possibility: that the selection criteria depend on information the buyer possesses about its preferences.

The rest of the paper is organized as follows. Sections 2 and 3 describe the model and compares it with other related models found in the literature. Sections 4 and 5 solve the model for the commitment and no commitment case, respectively. Section 6 discusses how changing some simplifying assumptions would affect the results. Two final sections conclude: Section 7 discusses the first steps towards testing the theory by qualifying some of the findings and investigating how they relate to some institutional characteristics of U.S. Government procurement practices. Section 8 presents the main policy lessons of the theory.

2 Environment

A buyer needs to procure a single, indivisible item that can be acquired from one of two potential suppliers.[1] Products of different suppliers have different costs and different values for the buyer. The costs are the suppliers' private information; the values are the buyer's private information.

[1]For simplicity, we assume that the goods must be purchased eventually. This simplifies the analysis by avoiding the need to establish reservation prices.

Let c_i denote the cost of producing the item by supplier i and θ_i the value of this good to the buyer. The costs c_1 and c_2 are independent random variables, with the same absolutely continuous distribution F, with positive density f over a compact support $[\underline{c}, \bar{c}]$. Throughout the paper we shall assume that F is *regular*, in the sense that

$$\vartheta(x) = x + \frac{F(x)}{f(x)}$$

is a monotone increasing function. The distribution of (θ_1, θ_2) is independent of c_1 and c_2. Before the auction begins, each seller i observes c_i. The buyer may directly observe (θ_1, θ_2), if it so wishes. It is convenient to reparameterize (θ_1, θ_2) as the valuation difference

$$\Delta = \theta_1 - \theta_2$$

and the average valuation $\bar{\theta} = (\theta_1 + \theta_2)/2$.

The novel aspect of the environment is the presence of θ_i. If $\theta_1 = \theta_2$ were known, then the model would be just like a standard independent private values procurement auction. Here, it will be assumed that θ is exogenous and the private information of the *buyer*. Alternatively, θ could also be the outcome of private information on either side of the market, or of decisions of the suppliers about specifications of their products. Che (1993) studies a procurement model in which the attributes other than price are a choice variable of the sellers, and the auction should be designed in a way that does not distort the incentives to select the most beneficial bundle of characteristics.[2] If the θ_i are exogenous but the private information of the sellers, one would obtain a mechanism design model with multidimensional types similar to the ones studied by Laffont, Maskin, and Rochet (1987), McAfee and McMillan (1988), Armstrong (1996) and Zheng (2000). Assuming that θ comes purely from the buyer's private information allows us to focus on the aspect of the problem that is novel in the literature.

The present framework is also different from the scale auction studied by Athey and Levin (2001). In a scale auction, the criterion for selecting the winner is an average of bids in different dimensions. That is similar to the present hypothesis that the criterion is a combination of price and other aspects. However, two crucial distinctions exist. First, in the current model,

[2] Rezende (2000) provides another example of a model in which product specifications are chosen by the suppliers.

information about perceived quality is on the buyer's side. Second, in a scale auction, bids are on a per-unit basis, and the final price depends also in the quantity traded, which is determined after the auction. This gives rise to the possibility of strategic skewing of the bids. This possibility is not present in the current model, since prices considered here are the total amounts paid upfront.

As the subsequent analysis will reveal, the model has more similarities to the asymmetric auctions literature. Since Griesmer, Levitan, and Shubik (1967), most of this literature focuses on the existence and characterization of first-price auction equilibria (Lebrun, 1996; Bajari, 1999; Maskin and Riley, 2000b). Some work has been done on the comparative statics of revenues across auction rules, and results have been obtained for specific forms of asymmetry (McAfee and McMillan, 1989; Maskin and Riley, 2000a). In the current framework, not only is there a natural form of asymmetry, but this form also allows for sharp conclusions regarding how the buyer's information should be gathered, disclosed and utilized at the auction.

3 The Game Form

In order to provide a comprehensive analysis of the effect of θ, the procurement process is modeled as a game with four stages:

Information Gathering. At this stage the buyer decides if it will learn θ or not.

Information Disclosure. The buyer should then decide how much information about its valuation to disclose to the sellers.

Auction. A procurement auction takes place: sellers compete by making price offers to the buyer.

Allocation decision. Based on information revealed at the earlier stages, the buyer selects the winner.

The following subsections explain in detail how each stage is formally modeled.

3.1 Information Gathering

The information gathering stage is modeled in a simple and natural way. We assume that the buyer initially has a prior belief about θ,[3] and can decide whether to learn the realization of θ or to proceed without knowing it. We assume that learning is costless — it would be straightforward to incorporate a fixed cost of learning in the analysis.

3.2 Information Disclosure

Information disclosure is here represented formally by *event disclosure*, in the spirit of Grossman (1981) and Milgrom (1981): the informed party communicates by disclosing an event that has happened (i.e., a set that contains the true realization of θ). We assume the buyer is able to disclose any event that contains the truth. So it cannot lie, but it can be vague.

3.3 Auction stage

At this stage the procurement auction takes place. We study an *open biased auction*: a sequential auction in which each seller has a price that starts at $\bar{c}$ and gradually drops. For the sake of concreteness, we assume that the protocol is as follows: sellers alternate turns deciding to reduce its own price by a small decrement or not.[4] As in a button auction, the decrement is exogenous[5] and very small, so that a continuous approximation is valid.

The auction is open: each bidder observes the opponent's current ask price. Unlike the button auction, however, stopping does not mean dropping out of the auction. Each seller is (and should be) allowed to restart dropping its price whenever it wants. The auction only ends when both sellers decide to stop.

While this particular auctioning protocol seems to be quite specific, we claim that the conclusions drawn from it have general validity. First, it will be shown that a buyer with commitment power can use this mechanism to obtain the allocation and price that maximize its revenue in the class of all incentive-compatible mechanisms. Therefore, restricting the analysis to

[3] We also assume that this prior knowledge is shared by the sellers.

[4] We assume that sellers are not allowed to post negative prices, so that the game ends in finite time.

[5] So the possibility of jump bidding or bid delaying is not considered.

this particular protocol is without loss of generality. Second, in the case of a buyer without commitment power, the open auction provides a natural formalization to what commonly occurs in informal negotiations that will take place. The important rules in this auction, such as the ability to "re-enter" and the observability of the opponent's price, reflect the inability of the auctioneer to commit to restricting the bidders actions.[6]

3.4 Allocation decision

In the final stage of the game the buyer decides which seller will supply the good and at which price. This is the point in which the assumption of commitment power is crucial. Under commitment power, the buyer will be able to announce at the outset an allocation rule and a pricing rule based on what happens at the auction and its own private information. As a result, in this step it will simply implement the pre-announced policy. Under no commitment, the buyer is allowed to revise its allocation decision at this point. Not only can it pick the winner, but it is also allowed to modify the transaction price. So the ask prices that arise in the auction are binding only from the point of view of the sellers, not the buyer.

We formally model this last stage as follows: the buyer must select a winner once the auction ends. It may then engage in further price negotiations with the winner. Such bargaining is modeled as a repeated game, in which at each stage the buyer proposes an allocation and a price to the selected seller, and it, in turn, accepts it or not. If not, the stage game is repeated.[7] The repeated bargaining protocol is aimed at capturing the buyer's lack of commitment: If a certain deal is rejected, it cannot refrain from proposing a new one shortly thereafter.

4 Procurement under commitment

Under commitment, a buyer can credibly announce before the auction an allocation rule as a function of the auction prices. We could represent this

[6] The restriction that ask prices only move downwards is a consequence of commitment power by the sellers, not the buyer. An ask price is understood here as a binding promise to deliver the good at that price, and therefore by definition sellers cannot raise their bids.

[7] We assume that there is a small delay between stages, and parties evaluate expected future profits with a discount factor δ. We will only consider the limit case where δ converges to 1.

rule as a function from prices to winning probabilities, as is conventional in the mechanism design literature. Here we chose instead to work with *bias functions*. A bias function ϕ corresponds to the allocation that assigns the good to seller 1 if $p_1 < p_2 + \phi(p_2)$.[8] Intuitively, ϕ represents the bias in the allocation rule in favor of seller 1 — a positive ϕ means that 1 can win the auction at a price higher than the opponent's best offer. We restrict attention to bias functions with the property that $x + \phi(x)$ is strictly increasing in x. That guarantees that the allocation rule is monotonic in both players' prices.

Since the bias function representation does not allow for allocations that are not monotone in prices it is not as general as the usual representation with allocation functions. Using bias functions is however without loss of generality, since as we will see the buyer will be able to implement the revenue-maximizing optimal auction using an appropriate bias function.

Before characterizing the optimal auction, it is useful to discuss how bidding takes place given a bias function. Consider seller 1: its profit is $(p_1 - c_1)$ if $p_1 < p_2 + \phi(p_2)$ and 0 if $p_1 > p_2 + \phi(p_2)$. It is clear that, if p_2 was fixed, it would elect to drop p_1 until it barely wins at $p_1 = p_2 + \phi(p_2)$ or slightly below it, as long as still $p_2 + \phi(p_2) > c_1$, or otherwise stop. Furthermore, since the price target $p_2 + \phi(p_2)$ is monotone in p_2, this alternative is safe, in the sense that an eventual future drop of p_2 would not make 1 regret its decision.

It is easy to see that bidding in this straightforward way constitutes a Nash equilibrium of the bidding game. As in the English auction, other equilibria exist, but this is the only one that is sequentially rational.

Proposition 1 *The only sequentially rational equilibria of the biased open auction involve strategies that specify, for seller 1, to*

$$\begin{cases} \textit{drop } p_1 & \textit{if } p_1 > p_2 + \phi(p_2) > c_1; \\ \textit{stop} & \textit{if } p_1 \leq c_1 \textit{ or } p_1 < p_2 + \phi(p_2); \end{cases}$$

[8] We will not explicitly discuss the problem of how to break a tie. Ties in terms of costs occur with probability zero; so in a game with discrete bidding, the probability of a tie can be made arbitrarily small with small decrements. Alternatively, one can implicitly assume that an endogenous tie-breaking rule is in place: in the event of a tie the sellers are given an opportunity to break the tie voluntarily. Since in all equilibria studied here one of the sellers is indifferent between winning or not, such a tie-breaking protocol would be effective. Endogenous tie-breaking rules are discussed in Milgrom (1989), Simon and Zame (1990, 1999) and Jackson and Swinkels (1999).

and, for seller 2,

$$\begin{cases} \textit{drop } p_2 & \textit{if } p_2 > p_1 - \phi(p_2) > c_2; \\ \textit{stop} & \textit{if } p_2 \leq c_2 \textit{ or } p_2 < p_1 - \phi(p_2). \end{cases}$$

Proof: See Appendix. □

Figure 1 illustrates how bidding takes place in the (p_1, p_2)-space. The bias function divides the space in two regions: above the $p_1 = p_2 + \phi(p_2)$ line, seller 2 wins, while below this line the winner is seller 1.

In any subgame that starts with prices outside the $p_1 = p_2 + \phi(p_2)$ line (and with prices above costs), only the provisional loser at the current price combination will bid. In the example in the figure, at the depicted top-right point 2 is provisionally winning, and initially only seller 1 will bid. The (p_1, p_2) pair falls vertically towards the $p_1 = p_2 + \phi(p_2)$ line. Once this line is reached, seller 2 starts bidding as well, and prices drop along the line as long as they are still above each sellers' costs. Once one of the prices reaches cost —in the example shown, p_1 reaches c_1—, the corresponding seller stops bidding and the auction ends. The final price of the loser is its cost, while the winner's price will be the loser price plus or minus the bias.

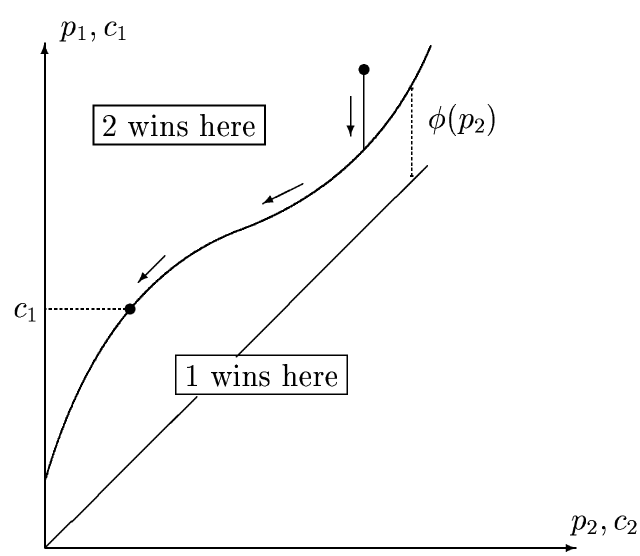

Figure 1: Bidding in a biased auction

4.1 The optimal biased auction

Implicitly define the bias ϕ^* to solve

$$\Delta = \phi^*(x) + \frac{F(x + \phi^*(x))}{f(x + \phi^*(x))} - \frac{F(x)}{f(x)}$$

for each $x \in [\underline{c}, \bar{c}]$, where $\Delta = \theta_1 - \theta_2$, and F and f are the distribution and density of the sellers' costs.

Proposition 2 *The biased auction with bias ϕ^* is the buyer's revenue maximizing auction among all incentive-compatible mechanisms.*

Proof: See Appendix. □

The proof follows Myerson (1981) and McAfee and McMillan (1989) in characterizing the optimal direct revelation mechanism and invoking the Revelation Principle to argue it is optimal across all incentive-compatible mechanisms. It then shows that the proposed biased auction implements the same allocation and payments as the optimal auction, and is therefore also optimal. The optimal allocation is made when the buyer selects the seller that yields the highest marginal revenue $\theta_i - c_i - F(c_i)/f(c_i)$, rather than the true revenue $\theta_i - c_i$. The extra term accounts for the informational rents, and is the reason why $\phi^* \neq \Delta$ in the formula above.

By proposition 2 the characterization of the mechanism through bias functions is without loss of generality; by implementing a biased auction with the optimal bias ϕ^*, the buyer can obtain an expected revenue as high as that of any other game the sellers are willing to participate in.

Comparing the optimal bias ϕ^* with the "honest" bias Δ provides some insight on the economic trade-off the buyer faces.

Proposition 3 *If F is regular, ϕ^* has the same sign as Δ at all points. If F is log-concave,*[9] *$|\phi| \leq |\Delta|$ at all points.*

Proof: See Appendix. □

Proposition 3 shows that the buyer wants to bias the auction towards the preferred supplier, but by an amount less than Δ, its true preference. Intuitively, this is because the optimal bias is determined by trading off the two effects of a bias in the buyer's profits:

[9] A distribution F is *log-concave* if $\log F$ is concave or, equivalently, if $\frac{F(x)}{f(x)}$ is monotone increasing in x.

The efficiency effect: Moving the bias towards Δ leads to an improvement in the value of the trade, since it makes the best supplier more likely to win;

The competition effect: Moving the bias away from zero reduces competition between the suppliers, increases their mark-ups, and increases the price to be payed.

In order to maximize efficiency, the bias should be equal to Δ; in order to maximize competition, it should be zero. The propositions above show that in general the optimal bias lays between these two targets.

4.2 Information Gathering and Disclosure

Once the buyer credibly commits to an allocation rule, information about θ is irrelevant for the sellers, since it does not affect their profits. As a result, under commitment, no information disclosure beyond the announcement of the bias rule is necessary (or relevant).

As the formula for ϕ^* reveals, the buyer needs to know Δ in order to determine the appropriate optimal bias. If it elects not to learn about θ, it will be forcing itself to implement a sub-optimal auction, with a bias that does not depend on Δ.

So we conclude that

Proposition 4 *A buyer with commitment power always prefers to learn about Δ. Once the bias has been announced, information disclosure about Δ does not affect the game play.*

5 Procurement under no commitment

The analysis of the procurement auction made in the previous section depends heavily on the hypothesis that the buyer can credibly commit to a pre-announced bias schedule. And yet, when chosen optimally, this bias schedule is inefficient — it requires the buyer to pick the worse deal with positive probability. So in this game there is always a temptation to renege on the established optimal bias in the final stage. If there is scope for renegotiation —and, as discussed in the Introduction, this is often the case in complex procurements— the sellers may anticipate that the buyer will revise the bias afterwards.

In this section we solve the model in the absence of commitment power by the buyer. As we will see, this absence not only affects the behavior at the final stage, but critically affects the incentives to acquire and disclose information. Since predicted behavior at the later stages affects the incentives at the earlier stages, it is convenient to discuss the stages backwards.

5.1 Allocation decision under no commitment

A buyer without commitment can do two things once the auction stage is over. First, it can select a prospective winner arbitrarily — it does not need to follow any pre-specified rule. Second, it can bargain with the selected supplier in an attempt to bring the price further down.

Bargaining is formally modeled as follows. The buyer makes a sequence of price offers, and the selected seller accepts or rejects them. We assume that the buyer has the option of stopping the process and taking the auction price of the selected seller. The process ends when the seller accepts an offer or if the buyer decides to stop the process and take instead the auction price. We assume players face a time discount factor across stages in the bargaining process, and are interested in characterizing the equilibrium when this factor converges to one.

Let T_i be the support of the belief the buyer has about c_i, after observing i's behavior in the auction stage, and let p_i be i's final auction price. Since bidding below cost is not rational, in equilibrium $\sup T_i \leq p_i$. We therefore can use the following characterization of behavior in the bargaining step:

Proposition 5 *Consider a situation where i is selected and $\sup T_i \leq p_i$. Then there is an equilibrium of the bargaining stage where the final price converges to $\sup T_i$ as the discount factor goes to 1.*

Proof: This is a consequence of the Coase conjecture (Coase, 1972; Gul, Sonnenschein, and Wilson, 1986). □

So the buyer can take advantage of information that arises in the previous stages, but it cannot elicit more information at the negotiation stage. In particular, as long as $\sup T_i = p_i$, the final price will be the same as i's auction price.

However, the presence of a potential bargaining stage has a important effect on bidding in the auction, because it prevents sellers from acting in a way that reveals their types. A corollary of proposition 5 is:

Theorem 1 (Secrecy) *If the buyer cannot commit, sellers will only reveal their cost when their expected profits are zero.*

Proof: If they do, later in the negotiation stage the buyer will drive price to their cost. □

We call bid strategies that satisfy this condition *secretive*. The Secrecy theorem states that secretive bidding is a necessary condition for an equilibrium. If strategies are such that, once the auction ends, $\sup T_i = p_i$, we say that bidding is *fully secretive*.

5.2 The auction with perfect information about Δ

We begin the analysis of the auction stage by the case where the sellers know the value of Δ at the outset.

As long as sellers act fully secretively, a buyer cannot do better than obtain a profit of $\theta_i - p_i$ by selecting i, where p_i is the auction price. So it will select the most efficient deal, and with that effectively implement a bias $\phi = \Delta$.

In this case, the analysis of section 4 follows, and we can conclude that bidders follow the strategies described in proposition 1. All that remains to be done is to verify whether the strategies of such equilibrium are indeed fully secretive.

Proposition 6 *Consider a pair of strategies that specify, for seller 1, to*

$$\begin{cases} \textit{drop } p_1 & \textit{if } p_1 > p_2 + \Delta > c_1, \\ \textit{stop} & \textit{if } p_1 \leq c_1 \textit{ or } p_1 < p_2 + \Delta, \\ \textit{act in any way independent of } c_1 & \textit{if } p_1 > c_1 \geq p_2 + \Delta; \end{cases}$$

and, for seller 2,

$$\begin{cases} \textit{drop } p_2 & \textit{if } p_2 > p_1 - \Delta > c_2, \\ \textit{stop} & \textit{if } p_2 \leq c_2 \textit{ or } p_2 < p_1 - \Delta, \\ \textit{act in any way independent of } c_2 & \textit{if } p_2 > c_2 \geq p_1 - \Delta. \end{cases}$$

In an auction where Δ is publicly known, bidders following this strategy are fully secretive. As a consequence, these are sequentially rational equilibria of the bidding game.

Proof: See Appendix. □

5.3 The auction with uncertainty about Δ

The results of the previous section characterize the situation whenever the buyer finds to be in its own interest to reveal Δ. The bidding game when the buyer does not reveal Δ fully is significantly more complex, since bidders should not only evaluate the effect their actions have in each other's incentives, but also be concerned in revealing information in the process, and in anticipating the possibility that the buyer itself may disclose information in the future.

However, when the buyer cannot commit and the Secrecy Theorem is valid, uncertainty about Δ leads to a breakdown of the auction: not only no bidding is an equilibrium, but also there are no equilibria where the winning bid is lower than the highest possible cost $\bar{c}$.[10] Let $\underline{\Delta}$ and $\overline{\Delta}$ be the infimum and supremum of the support of the belief the sellers have about Δ. Using the Secrecy Theorem we show that:

Proposition 7 *In equilibrium, no seller 1 ever bids below $\bar{c} + \overline{\Delta}$, and no seller 2 ever bids below $\bar{c} - \underline{\Delta}$. As a result, the winner's final price is always $\bar{c}$.*

Proof: See Appendix. $\square$

In order to maintain secrecy, each type of the seller has to act as if it were a type with higher cost until its expected profit is zero. But when there is uncertainty about Δ there is always hope of winning with a price above cost. The auction does not unravel, and all types pool by effectively not bidding. The cost revelation mechanism of the auction collapses under the combination of uncertainty and lack of commitment, and equilibrium prices are well above the ones observed otherwise.

5.4 Information Disclosure

Because of the break down of the auction if information about Δ is not disclosed, it is clear that the buyer always finds in its best interest to reveal as much as it can about Δ.

[10]Proposition 9 in the appendix establishes that there are indeed sequential equilibria that involve no bidding, and are supported by out-of-equilibria beliefs that are reasonable according to conventional criteria, such as the Cho and Kreps Intuitive Criterion, Grossman and Perry Perfection and Banks and Sobel Universal Divinity.

Theorem 2 *In any equilibrium, a buyer without commitment fully discloses Δ before the auction.*

Proof: The price obtained when Δ is not disclosed is $\bar{c}$, at least as high as any price that arises in an auction otherwise. □

5.5 Information Gathering

We have shown that a buyer without commitment power will fully disclose the information it has about Δ prior to the auction. The ensuing open auction is a particular instance of a biased procurement auction, where the bias is the efficient one rather than the one that maximizes the buyer's *ex ante* profits.

A final question remains: should the buyer invest in learning Δ in the first place, given that it anticipates that future course of action? Intuition tells that this information benefit is likely to be lower in this circumstance, since the outcome is less than optimal. This section establishes that in fact the benefit of learning Δ can be *negative*: the buyer may be *worse-off* by doing so.

The argument hinges on the shape of the buyer's expected profit function. Let $\Pi(\bar{\theta}, \Delta, \phi)$ be the expected profit of a buyer that learns $(\bar{\theta}, \Delta)$ and runs an auction with bias ϕ. We have

$$\begin{aligned}\Pi(\bar{\theta}, \Delta, \phi) &= E_c[(\bar{\theta} + \Delta/2 - c_2 - \phi(c_2))\mathbb{1}\{c_1 < c_2 + \phi(c_2)\} \\ &\quad + (\bar{\theta} - \Delta/2 - c_1 + \phi(c_2))\mathbb{1}\{c_1 > c_2 + \phi(c_2)\}],\end{aligned}$$

where the c subscript means that the expectation is taken over costs and $\mathbb{1}\{\}$ represents the indicator function.

If the buyer decides not to learn about $(\bar{\theta}, \Delta)$, its future choice of ϕ will not depend on the realization of these variables. Its expected profit is therefore

$$\begin{aligned}E_\theta[\Pi] &= E_\theta[E_c[(\bar{\theta} + \Delta/2 - c_2 - \phi(c_2))\mathbb{1}\{c_1 < c_2 + \phi(c_2)\} \\ &\quad + (\bar{\theta} - \Delta/2 - c_1 + \phi(c_2))\mathbb{1}\{c_1 > c_2 + \phi(c_2)\}]] \\ &= E_c[(E[\bar{\theta} + \Delta/2] - c_2 - \phi(c_2))\mathbb{1}\{c_1 < c_2 + \phi(c_2)\} \\ &\quad + (E[\bar{\theta} - \Delta/2] - c_1 + \phi(c_2))\mathbb{1}\{c_1 > c_2 + \phi(c_2)\}]],\end{aligned}$$

since Π is linear on θ in this case, the buyer's ex ante problem is exactly the same as if it faced a realization of θ equal to its expected value. In particular, the bias that it will effectively use is $\phi = E[\Delta]$.

If instead the buyer learns about Δ, then ϕ is a function of it, and Π is no longer linear. In the case of commitment, $\Pi(\bar{\theta}, \Delta, \phi^*(\Delta))$ is convex; however, in the case of no commitment, $\Pi(\bar{\theta}, \Delta, \Delta)$ can be concave:

Proposition 8 *Suppose that, for all values of in the support of* Δ, $2f_{c_1-c_2}(\Delta) \geq f(\bar{c} - |\Delta|)$.[11] *Then* $\Pi(\bar{\theta}, \Delta, \Delta)$ *is concave, and the buyer elects not to acquire information.*

Proof: See Appendix. □

The intuition behind this surprising conclusion may become clearer if we think of the buyer's profits as the difference between expected value and expected price to be paid. A realization of Δ one unit higher than expected (holding $\bar{\theta}$ fixed) means that θ_1 is up by half unit and θ_2 is down by half unit. So the expected effect is $1/2\,\Pr(1 \text{ wins}) - 1/2\,\Pr(2 \text{ wins})$, and expected value is convex in Δ.

However, the effect is almost the opposite for the expected price. A one-unit-higher Δ leads to an increase in p_1 by one unit (as long as $p_1 < \bar{c}$), and similarly a decrease in p_2 by one unit. So ignoring the cases where $p_i = \bar{c}$, we find that expected price is twice as convex as expected value, and as result the expected profit is concave. The $2f_{c_1-c_2}(\Delta) \geq f(\bar{c} - \Delta)$ condition in the proposition guarantees that the effect of the region where $p_i = \bar{c}$ is not strong enough to upset this conclusion.

Intuitively, the condition needed is that the distribution of costs and Δ is such that the event of a tie across bidders ($\{c_1 = c_2 + \Delta\}$) is not much less likely than the event $\{c_2 = \bar{c} - \Delta\}$.[12] If the distribution of costs does not assign a large probability for costs close to $\bar{c}$, this condition is satisfied for Δ near zero; and it is always vacuously satisfied for $|\Delta|$ very large.

6 Extensions

In order to obtain a clear theory of the effect of exogenous bias in a procurement auction, the analysis so far has imposed simplifying assumptions that avoid other issues likely to be present in practice. This section lists some of these issues and indicates how the conclusions would change if they were incorporated into the analysis.

[11] Here $f_{c_1-c_2}$ represents the density of the random variable $c_1 - c_2$.

[12] Or $\{c_1 = \bar{c} + \Delta\}$, if $\Delta < 0$.

6.1 Reserve prices

In practice a buyer often has the alternative of not purchasing the good if it finds out that the ask price of any seller is too high. As in a standard auction, allowing for this possibility leads to the introduction of *reserve prices*, price ceilings above which no trade is done.

Under commitment, optimal reserve prices are determined in a way entirely analogous to Myerson (1981). For concreteness, suppose the buyer's opportunity profit for no trade is zero. In this case, the optimal direct revelation mechanism calls for allocating the contract to 1, 2 or nobody depending on which of $\theta_1 - c_1 - \frac{F(c_1)}{f(c_1)}$, $\theta_2 - c_2 - \frac{F(c_2)}{f(c_2)}$ or 0 is highest.

This optimal allocation can be implemented by a variation in the biased open auction, in which the optimal bias is exactly the same as before and bidding starts at pre-specified levels (p_1^0, p_2^0),[13] where p_i^0 solves $\theta_i = p_i^0 - \frac{F(p_i^0)}{f(p_1^0)}$.

Allowing for no procurement will also change the analysis in the no commitment case. One relevant change is that information about $\bar{\theta}$ is now valuable: if its realization is too low, the buyer has the option of not acquiring the good. Thus the shape of the buyer's expected profit as a function of θ resembles a saddle, being convex in the $\bar{\theta}$ direction and concave in the Δ direction. Which particular effect dominates in order to assess if information should be acquired depends on the region in which the buyer's prior beliefs lie. If realizations of θ are believed to be close to the cost range, making losses a real possibility, then it will be in the buyer's interest to verify that the transaction is indeed profitable by learning θ. If, however, the value of the transaction is very likely to be positive anyway, then the effect identified in Proposition 8 will dominate and information may be harmful.

6.2 More than two sellers

Introducing more than two potential sellers in the model does not affect the strategy to characterize the optimal auction described in the proof of proposition 2 in any fundamental way. Still, the optimal allocation selects the seller with the index that maximizes $\theta_i - c_i - \frac{F(c_i)}{f(c_i)}$, and this allocation rule can still be represented by a bias scheme. However, describing the bias scheme is more cumbersome now, since there must be a bias schedule $\phi_{\{i,j\}}$

[13] The interpretation of these initial prices is the same as a reserve price: i is allowed to abstain from bidding at the level p_i^0, in which case the opponent immediately wins or, in case of mutual abstention, the item is not procured.

for each pair $\{i, j\}$ of sellers. As before, we define each bias schedule as the function that implicitly solves

$$\theta_i - \theta_j = \phi_{\{i,j\}}(x) + \frac{F(x + \phi_{\{i,j\}}(x))}{f(x + \phi_{\{i,j\}}(x))} - \frac{F(x)}{f(x)},$$

for every pair of seller indices with $i < j$ and every x in the support of c_j.

A seller i is the winner if its price beats each opponent's price plus the particular bias: i.e., if $p_i < p_j + \phi_{\{i,j\}}(p_j)$ for every $j > i$ and $p_i < p_j - \phi_{\{j,i\}}(p_i)$ for every $j < i$. So the allocation rule is the outcome of several pairwise comparisons identical to the one studied in Section 4.1, and the same conclusions apply to the pairwise bias schedules: if the cost distribution is regular, they have the same sign as the true bias $\theta_i - \theta_j$ pointwise; if it is log-concave, they are smaller than the true bias in absolute value pointwise.

As for the analysis of the game without commitment, several results do not depend on the number of sellers. Among them, the analysis of the bidding dynamics with $\theta_1, \ldots, \theta_n$ publicly known, the Secrecy Theorem, and the conclusion that full disclosure is optimal.

The part of the analysis that must be changed more substantially to apply to the case with many sellers is the one regarding the incentives to acquire information. The next section discusses the necessary changes.

6.2.1 Information gathering with many sellers

This section discusses whether it is still true that it may be optimal not to learn about θ under no commitment. To establish that, we must investigate the concavity of the buyer's profit function Π as a function of the differences between θ_i's. As we will see, the answer will depend on which particular difference that is taken. Π may be still concave if the difference is between the winner and the runner-up in the auction, but will be convex or flat for other comparisons. So the answer will depend on the particular region of the θ space where the comparison is made. We start with a computation of the derivative of Π:[14]

Lemma 1 $\frac{\partial}{\partial \theta_i}\Pi$ *equals the probability that* $\theta_i - c_i$ *is the second highest among all sellers.*

[14]For simplicity, we ignore the possibility that $p_i = \bar{c}$ in what follows.

Proof: See Appendix. □

The concavity of Π for the n-sellers case depends on which particular differential is considered.[15] If one considers information about $\Delta_{ij} = \theta_i - \theta_j$, where i and j are the likely winner and runner-up in the auction, then concavity holds and the same conclusions of the model with 2 sellers obtains — this is not surprising, since in the 2 sellers case the only possible differential is between winner and runner-up.

If however the difference considered is between the runner-up and some other loser, then the conclusion is reverted, and buyer profits are convex in that variable. In this case, it would be profitable to gather and disclose information about this difference.

So the particular information gathering policy depends on the region where the buyer's priors lie. If it is believed that there are two stronger sellers in the market, then the analysis made before holds for the value differences among them. If not, the analysis is more delicate and depends on the particular prior about θ and the shape of the cost distribution.

7 Practical Implications and Caveats

Before concluding, it is worthwhile to discuss how do the predictions of the theory relate to procurement practice. This section discusses each of the three main conclusions in turn: the predictions that the bias should be underrepresented under commitment; that information about the final allocation rule should be disclosed early on; and that the value of information about Δ may be negative in the no commitment case.

As a preliminary attempt to test the first two results, we again use the FAR and literature on the U.S. Government acquisition practices to investigate whether they can be explained by the incentives predicted by the model.

The third prediction requires some qualifications before it can be tested appropriately. So instead of investigating its practical validity, subsection 7.3 points out some caveats and argue for the interpretation thought to be most appropriate to the result.

[15] The appendix provides a calculation when the number of bidders is 3.

7.1 Underrepresentation of Bias under Commitment

One of the main results of this paper is that it is optimal to distort the auction rule towards introducing excessive competition among suppliers. In this section we discuss the possible connection between that finding and the layered structure of the administration of government procurement in the United States.

In his study of U.S. Government procurement practices, Kelman (1990) argues that the personnel in charge of the procurement process is divided in two groups with distinct and opposing cultures: "the technical (...) people on one hand and the contracting people on the other" (Kelman, 1990, p. 24).

The "technical people" are usually members of the agency that is going to use the good being purchased. They have expertise about the technical requirements on the good, and tend to be directly affected by its quality.

The "contracting people" have less knowledge about the technical aspects of the procurement and know more about the its legal and procedural aspects. They are not directly affected by the outcome of the contract, but they are the ones to answer if impropriety is suspected.

In the procurement process these two groups deal with different tasks. The technical people are in charge of the initial specification of the product. Once this is done, the procurement process is taken up by the contracting people, which administer the bidding process and determine the winner.

Interestingly, the weighting system that is used to select the winner is usually a compromise between the two groups. The technical people tend to favor a criterion with bigger weights on dimensions besides price, such as quality, supplier reputation and technical assistance, while contracting people prefer to attach a bigger weight on price.

One conventional explanation for this disagreement is that it matches (immediate) preferences. The technical people tend to enjoy higher quality, since they are going to be using the product, and do not care for price, since it is taxpayer, and not them, that will pay the bill.[16]

Under this view the role of the contracting people layer in the procurement process is to reduce this distortion. This group does not benefit from quality and their aim is to minimize the risk of protests and the appearance of

[16]Rogerson (1990) is an example of a paper that studies the distortions that arise when the procuring agent (in that case, the military) has objectives that are distinct from the principal (the Congress). In that paper, the military seeks to maximize value alone, while the Congress cares about value minus cost.

impropriety. In order to do so, this group prefers to choose the cheapest supplier, and, in doing so, may indirectly enhance welfare, by reducing an excessive bias towards quality.

The current theory provides an alternative explanation. The observed behavior would take place even if the technical people fully internalized the taxpayer cost of the procurement and defended the efficient winner selection criterion.

According to this interpretation, the role of the contracting people in reducing the utilized bias is to implement something closer to the optimal mechanism. After the technical people reveal the true socially efficient bias Δ, the contracting people implement an auction with a smaller bias that explores the competition among suppliers and increases the buyer's revenue.

To distinguish between the two theories, one would need to obtain a good assessment of what is the appropriate size of the efficient bias. A crude way to do so is to compare the government procurement practices with those in the private sector. A stylized fact from the literature is that the Government selects much more often the lowest bidder (i.e., employs smaller biases). According to Kelman (1990), "65% of the most recent major contracts discussed by respondents to the Government Computer Managers Survey were awarded to the low-priced bidder, compared with only 41% in the private sector Computer Managers Survey" [p. 60]. Burt (1984) states that "competitive bidding is underutilized in the private sector and overused in the public sector."[17]

The implication of this fact for the theories above depends one what one is willing to assume about how procurement is done in the private sector. One possibility is that it is efficient.[18] In this case, the former theory would predict that the bias proposed by the technical people would be larger than the ones employed by the private sector, and the one eventually used would be about the same. In the second theory, the initially proposed bias would be about the same and the one employed would be smaller, which is what is observed in practice. Therefore according to this interpretation the observed differences provide evidence for the second interpretation.

On the other hand, one may prefer to believe that in the private sector

[17]Interestingly, Burt (1984) uses this fact to argue for less quality bias in the private sector, implicitly using the government practices as the benchmark.

[18]Private sector procurement is typically much less formal than its public counterpart. If one is willing to interpret this informality as lack of commitment, then one would predict that the bias in these procurements would be efficient.

the optimal, rather than the efficient, bias is used.[19] In this case, the second interpretation would predict that the bias in the public and private sectors would be the same, while the first would predict that the bias in the public sector would be *larger*. In this case neither theory would be able to predict the observed pattern, but the second theory would still comes closer to it.

In any case the preliminary comparison between practices in the government and the public sector tends to reject the straightforward incentives explanation in favor of the hypothesis that government procurement institutions indeed impose a form of bias underrepresentation.

7.2 Information Disclosure

An obvious evidence in favor of the second main conclusion, namely that it is a bad idea to run an auction without clearly disclosing the winner selection rule, is that this is a common practice in government procurement.

Even though the FAR allows for a significant degree of vagueness about the winner selection criterion before the auction, "in practice most agencies choose to present exact evaluation weights, often in great detail, in RFPs" (Kelman, 1990, p. 56).[20]

Again this behavior can be explained by the incentives in place for the contracting officers: the more precise the RFP is, the least room there is for discretion, and therefore the weaker is the case for a protest. But we can take one step back and look at the situation from the point of view of the designers of the institutional setting that motivates the officer, and ask why those incentives are there in the first place.

The current theory provides an answer to this question. In any case, and in particular in the no-commitment case, full disclosure is indeed a good thing, and a central ingredient for a successful auction.[21]

[19] Perhaps long-time relationships substitute for fixed rules in providing commitment power to the buyer.

[20] RFPs are *requests for proposals*, public documents that initiate the procurement process.

[21] Of course, this conclusion is qualified by the many simplifying assumption made here. For example, (Kelman, 1990) argues in favor of more discretion because he believes a great deal of learning takes place during the procurement process. It isn't clear how learning would change the conclusions made above, since there are potentially many ways to model such process.

7.3 Information Acquisition

The most counterintuitive result in this paper is the one that predicts that a buyer without commitment power is often better off *not* learning about its preferences. Several qualifications should be made about this conclusion.

First, it should be understood that this is a statement about the value of information about θ along one dimension only, namely, $\Delta = \theta_1 - \theta_2$. It just so happens that under the maintained assumptions this is the only relevant dimension. As discussed in section 6, if the buyer had the alternative of not procuring the good or if there were more than two suppliers this conclusion might not hold, since Δ is no longer the only dimension that matters.

In a more general context the result would still hold in the sense that it still may be true that the value of information on the difference in quality between the winner and the runner-up is negative. But this value will be positive along other dimensions, such as the average quality $\bar{\theta}$ and the difference between the runner-up and other losers, if other alternatives are available to the buyer. The overall value of information about θ will then be a mixture of these effects, and its sign will depend on specific distributional assumptions.

Another qualification is about the particular way information acquisition is modeled. In this paper, the decision to learn about θ is publicly observed; the buyer is not allowed to secretly observe θ and then pretend it did not. This way of modeling is conventional in the contracts literature, but may not be realistic in many contexts.

This observation provides an intuitive base to understand why the result is possible: by publicly declining to learn about θ, the buyer can in a rudimentary way commit not to use it later, and that commitment value offsets the information loss.

From this observation it is reasonable to expect that the result is highly dependent on the specific way information gathering is modeled; if the buyer information gathering was not public, then the logic of the last paragraph would break down.

So a literal interpretation of the information acquisition result is not necessarily appropriate. Rather, it is best to see it as statement about (1) the aspects of information that are more valuable to the buyer — namely, the average quality across sellers rather how much better the winner is from the runner-up — and (2) how valuable is commitment in this game — only by obtaining inefficiently high revenues the buyer finds in its interest to learn about Δ.

In this sense the result is reminiscent of Grossman and Stiglitz (1980): here too there seems to be trade-off where allocational efficiency distorts the incentives to obtain information — in this case, even when information is free.

8 Concluding Remarks

A buyer that contemplates the possibility of using a procurement auction to meet a particular demand does so because the demand is indeed particular: the good in question is often complex, and requires substantial investment to produce, design or even describe. While this complexity certainly makes costs and therefore prices uncertain, it also makes products offered by different suppliers different along other dimensions.

For a buyer of such products, a matter of great importance is how it should act on perceived differences along these other dimensions. This paper has provided answers to all natural questions about the treatment of information on these differences: what are the incentives to collect, disclose, and act on this information. Furthermore, the answers are valid under a wide range of assumptions about the distributions of costs and the variables that affect these other dimensions.

Instead of re-enumerating all the results of the paper, this conclusion discusses the three findings that are believed to be of greatest interest to practitioners.

The first result is that in the optimal auction the buyer should commit to bias its choice of supplier towards the preferred one, but the bias should be less then the true perceived difference in value. Thus, from an efficiency perspective, the optimal allocation favors the weaker supplier. The reason for that is that the optimal bias is the compromise between two conflicting objectives: to maximize the efficiency of the allocation, that would ask for a full bias, and to maximize the degree of competition between suppliers, that would ask for no bias.

The two other results hold when the buyer is believed to lack commitment power to implement such bias. One of them may be translated in practical terms to the following piece of advice for the buyer: "It is a bad idea to let the suppliers bid without knowing how you will choose the winner". The reason for that is that with this uncertainty, open bidding no longer has the critical virtue of revealing the suppliers' costs in a way that is immune to opportunism by the buyer. When this virtue is lost, the resulting equilibria have very high prices as outcomes, since suppliers rationally defend themselves against this possibility by bidding cautiously.

The last and probably most surprising result is that learning about differences in suppliers' products may *decrease* the buyer's expected profits if it lacks commitment power. The reason for that is that while learning about this difference allows the buyer to pick the best supplier, it also implicitly places this supplier in a stronger bargaining position. So the gain in efficiency with this information is more than fully captured by the suppliers through higher equilibrium prices.

From a social perspective that finding creates an interesting trade-off between allocational and informational efficiency: increasing the buyer's ability to commit allows it to exert monopsonistic power (and therefore introduces allocational inefficiency) but provides the right incentives to collect information, while on the other hand without commitment one would have allocational efficiency, but incorrect incentives to collect information.

A Appendix: Proofs

Proof of Proposition 1: We proceed by induction on the state (p_1, p_2). Without loss of generality we consider 1's point of view. Player 2's situation is analogous.

It is certainly true that 1 should follow the prescribed strategy if $p_1 = 0$, since prices are not allowed to become negative. Now suppose both sellers follow the strategy for any state below (p_1, p_2).

If $p_1 > p_2 + \phi(p_2) > c_1$, 1 knows that if it stops the auction will end, because prices are in the region where 2 stops. In this case, it would lose the auction and its profit would be 0. By following the strategy and dropping its price, it may get a positive profit (in any situation where $c_2 + \phi(c_2) > c_1$).

If $p_1 < c_1$, by dropping its price 1 strictly increases its loss in the case of winning. Furthermore, the probability of winning is $\Pr(c_2 > \bar{p}_2)$, which increases if it drops p_1. So 1 is certainly better off stopping.

We now turn to the case when $p_2 + \phi(p_2) > p_1 > c_1$. We first observe that the only effect of stopping rather than dropping prices now is the risk that the auction may end, since otherwise the bidder can drop its price later. So we may evaluate 1's incentives conditional on 2 stopping in this context.

But conditional on that 1 certainly prefers to stop, since by doing that it gets $p_1 - c_1$ with probability 1, while dropping prices would decrease its margin with no effect in the probability of winning . $\square$

Proof of Proposition 2: We will follow the same strategy of Myerson (1981) and McAfee and McMillan (1989) to characterize the optimal auction. We then compare it to the biased auction and verify that they coincide.

Consider an arbitrary direct revelation mechanism. Let $\rho_i(c_1, c_2)$ be i's probability of winning as a function of the agents' reports, and $\chi_i(c_1, c_2)$ be the expected payment to seller i.

Seller i's expected profit by playing this mechanism with cost c_i and reporting $\hat{c}_i$ against a truthful opponent is

$$U_i(\hat{c}_i, c_i) = E[\chi_i(\hat{c}_i, c_j)] - c_i E[\rho_i(\hat{c}_i, c_j)].$$

In order for truth-telling to be rational, it must be that $c_i \in \operatorname{argmax}_{\hat{c}_i} U_i(\hat{c}_i, c_i)$ for all c_i's. By the Envelope Theorem (Milgrom and Segal, 2002), and as long as $E[\rho_i(\hat{c}_i, c_j)]$ is monotone in $\hat{c}_i$, we obtain the following characterization:

$$\begin{aligned} U_i(c_i, c_i) &= U_i(\bar{c}, \bar{c}) - \int_{c_i}^{\bar{c}} \frac{\partial}{\partial c_i} U_i(c_i, c_i) dc_i \\ &= U_i(\bar{c}, \bar{c}) + \int_{c_i}^{\bar{c}} \int \rho_i(x, c_j) dF(c_j) dx. \end{aligned}$$

We can combine the two formulas for $U_i(c_i, c_i)$ to obtain an expression for $E[\chi_i]$,

$$E[\chi_i(c_i, c_j)] = c_i E[\rho_i(c_i, c_j)] + U_i(\bar{c}, \bar{c}) + \int_{c_i}^{\bar{c}} \int \rho_i(x, c_j) dF(c_j) dx,$$

and then express the buyer's expected profit Π as a function of ρ_i only:

$$\begin{aligned} \Pi &= \sum_i E[\rho_i(c_i, c_j)\theta_i - \chi_i(c_i, c_j)] \\ &= \sum_i \left\{ E[\rho_i(c_i, c_j)(\theta_i - c_i)] - \int \int_{c_i}^{\bar{c}} \int \rho_i(x, c_j) dF(c_j) dx dF(c_i) \right\} - U_i(\bar{c}, \bar{c}) \\ &= \sum_i \int \int \rho_i(c_i, c_j) \left(\theta_i - c_i - \frac{F(c_i)}{f(c_i)} \right) dF(c_i) dF(c_j) - U_i(\bar{c}, \bar{c}). \end{aligned}$$

The optimal auction maximizes this expression subject to feasibility ($\rho_i \geq 0, \sum_i \rho_i = 1$) and the participation constraint ($U_i(\bar{c}, \bar{c}) \geq 0$).

It is clear that the optimal auction involves $U_i(\bar{c}, \bar{c}) = 0$ and $\rho_i = \mathbb{1}\{\theta_i - c_i - F(c_i)/f(c_i) > \theta_j - c_j - F(c_j)/f(c_j)\}$.

We now verify that the proposed biased auction satisfies these conditions. 1 wins in this auction if $c_1 < c_2 + \phi^*(c_2)$; under regularity, this happens if and only if

$$\begin{aligned} c_1 + \frac{F(c_1)}{f(c_1)} &< c_2 + \phi^* + \frac{F(c_2 + \phi^*)}{f(c_2 + \phi^*)} \\ c_1 + \frac{F(c_1)}{f(c_1)} &< c_2 + \left[\Delta + \frac{F(c_2)}{f(c_2)} - \frac{F(c_2 + \phi^*)}{f(c_2 + \phi^*)}\right] + \frac{F(c_2 + \phi^*)}{f(c_2 + \phi^*)} \\ \theta_1 - c_1 - \frac{F(c_1)}{f(c_1)} &> \theta_2 - c_2 - \frac{F(c_2)}{f(c_2)}, \end{aligned}$$

and this is the optimal allocation criterion.

As for $U_1(\bar{c}, \bar{c})$, notice that from the definition of ϕ^* we know that $c_2 + \phi^*(c_2) \leq \bar{c}$, for any $c_2 \in [\underline{c}, \bar{c}]$. So a seller never wins at a price above $\bar{c}$, and consequently $U_i(\bar{c}, \bar{c}) = 0$. □

Proof of Proposition 3: For the first claim, we can write the equation that defines ϕ^* as

$$\Delta = \left(x + \phi^*(x) + \frac{F(x + \phi^*(x))}{f(x + \phi^*(x))}\right) - \left(x + \frac{F(x)}{f(x)}\right).$$

Since by regularity $x + \frac{F(x)}{f(x)}$ is monotone increasing, ϕ^* should have the same sign as Δ.

For the second claim, it is enough to rewrite the definition of ϕ^* as

$$\Delta - \phi^*(x) = \frac{F(x + \phi^*(x))}{f(x + \phi^*(x))} - \frac{F(x)}{f(x)}$$

and use the definition of log-concavity. □

Proof of Proposition 6:

At any subgame all types play in the same way, except when $p_i = c_i$. So observing bidding the buyer can only tell either that $c_i < p_i$ or $c_i = p_i$. In either case, $\sup T_i = p_i$. The second statement comes from combining this result with proposition 1. □

Proof of Proposition 7: Let a_i be the infimum among the types of i that may eventually reveal their costs (i.e., if $c_i < a_i$ in equilibrium the buyer cannot figure out exactly what is c_i by observing i's behavior, no matter what happens in the game). Only a seller with cost below a_i would ever bid past this point. Consequently, all types below a_i in order to not reveal themselves should stop bidding before that under all circumstances, and certainly in equilibrium $p_i \geq a_i$.

Consider the situation of seller 1 with cost $c_1 < a_2 + \overline{\Delta}$. Such a seller does not need to obtain zero profits by revealing its cost, because it can still make a positive profit with $p_1 \in (c_1, a_2 + \overline{\Delta})$, no matter what happens with its opponent, since there is always a positive probability that $p_1 - p_2 < \Delta$. This means that we must have $a_1 \geq a_2 + \overline{\Delta}$. By an analogous argument, we need $a_2 \geq a_1 - \underline{\Delta}$.

Combining the two inequalities we find that $a_1 \geq a_2 + \overline{\Delta} \geq a_1 + \overline{\Delta} - \underline{\Delta}$; but since $\overline{\Delta} - \underline{\Delta} > 0$, this is impossible. So at least one of the sellers never reveal their types in equilibrium.

If 1 is such seller, then only types of 2 that satisfy $c_2 \geq \bar{c} - \underline{\Delta}$ may reveal their types. this can only be possible if $\underline{\Delta} > 0$. Similarly, when $\overline{\Delta} < 0$ seller 2 does not reveal its type and seller 1 does so only if $c_1 > \bar{c} + \overline{\Delta}$. Finally, if $\overline{\Delta} > 0 > \underline{\Delta}$ neither seller reveal its type.

□

Proposition 9 *Suppose the sellers are uncertain if Δ is positive or negative. Consider the following strategy/belief profile:*

1. *Sellers never bid.*
2. *If seller i bids, then then other players update their beliefs about its cost to a belief concentrated on $[\underline{c}, x_i]$.*
3. *If the other seller bids in response to a seller's deviation, then the other players believe its cost is $\underline{c}$.*

If $x_i \leq \underline{c} + (\bar{c} - \underline{c}) \Pr(\theta_i - \theta_j \geq 0)$, then this profile is a sequential equilibrium, satisfies the Intuitive Criterion, the iterated version of the Intuitive Criterion, the Equilibrium Domination Test (Cho and Kreps, 1987), and is a perfect sequential equilibrium (Grossman and Perry, 1986). If $x_i = \underline{c}$, then the equilibrium satisfies D1 and is universally divine (Banks and Sobel, 1987).

Proof: By the secrecy theorem, the final price a seller receives is the supremum of the support of the belief the buyer has about its cost. In the proposed profile, this price is $\bar{c}$ along the equilibrium path, and is equal or less than x_i in case i deviates. By deviating, at best i can win with certainty a margin of $(x_i - c_i)$. The condition on x_i guarantees that the deviation is not profitable. Also, after i's deviation, j's profits are zero: if it does not bid, it will lose the auction; if it does, it will obtain a price that is lower than its cost. So not bidding is rational and the profile is a sequential equilibrium.

One can verify the Intuitive Criterion using figure 2. The figure shows the profits of seller 1 as a function of its cost. Following the equilibrium, it will win a margin $\bar{c} - c$ with some probability less than one. If i deviates, then the best possible belief update that can happen is that is cost is anywhere at or below its bid. If it wins for sure after such deviation it gets its bid minus c_i which for sufficiently low c_i is higher than the equilibrium profits. The Intuitive Criterion requires that we restrict possible beliefs to the set of types that might contemplate this deviation and to verify that no type is willing to deviate assuming the worst belief update that is compatible with this restriction. In this game, the worst belief update is a belief that $c_i = \underline{c}$. Such belief is always compatible with the restriction and leads to zero expected profits, so the Intuitive Criterion is satisfied. Iterating the criterion would lead to a lower and lower set of possible types, but as long as this set is non-empty, $\underline{c}$ will be contained in it. So the iterative Intuitive Criterion and the Equilibrium Dominance Test are satisfied as well.

To verify if it is a perfect sequential equilibrium, we must check if the belief updating is credible. For this Grossman and Perry (1986) require that new beliefs have support contained in the old support, that they follow Bayes' rule when possible, and if there is a set K of types such that, if new beliefs are the same as the old conditional on K, then all types in K prefer to deviate and types not in K do not, then the belief update must be the old ones conditional on K. In the current game the first two conditions are clearly satisfied, and the third one is vacuous, since no (non-empty) set can be a valid set K. This is because a type in K close to the supremum of that set has a profit close to zero upon a deviation, which is worse than its equilibrium payoff.

The extreme belief update when $x_i = \underline{c}$ additionally satisfies D1 and universal divinity. D1 is satisfied when beliefs attribute probability zero to types c_i that meet the following condition: there is some other type c_i' that given any conjecture about future play that would make c_i weakly prefer to deviate, c_i' would strictly prefer to deviate. In this game, if $c_i > \underline{c}$ is willing to bid, it conjectures that bidding will increase its winning probability in a way that compensates for the lower margin. But then that change will also be even better to any type below c_i. So all types except for $\underline{c}$ are eliminated by D1. Iterating the procedure yields the same answer, so Universal Divinity is also satisfied. □

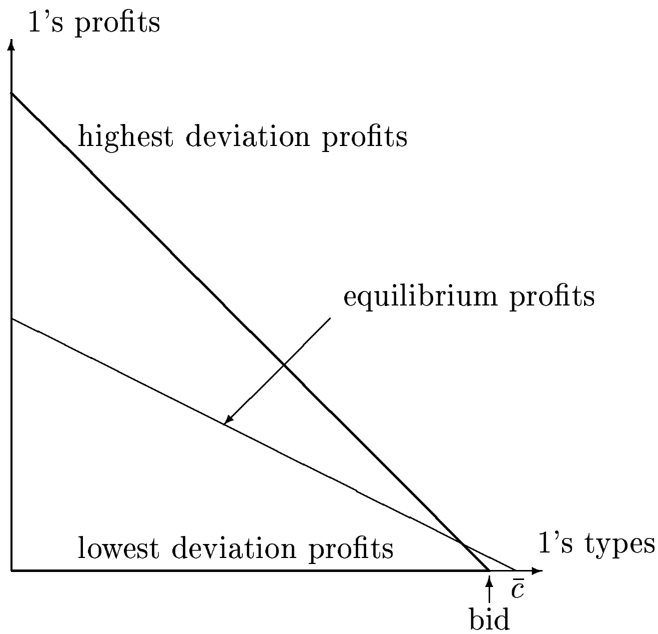

Figure 2: Out-of-equilibrium-path beliefs when auction breaks down

Proof of Proposition 8: The function of interest is

$$
\begin{aligned}
\Pi(\bar{\theta}, \Delta, \Delta) &= \int_{c_1 < c_2 + \Delta} (\bar{\theta} + \Delta/2 - p_1) dF + \int_{c_1 > c_2 + \Delta} (\bar{\theta} - \Delta/2 - p_2) dF \\
&= \int \int_{c_1 - \Delta}^{\bar{c} - \Delta} (\bar{\theta} + \Delta/2 - c_2 - \Delta) dF(c_2) dF(c_1) \\
&\quad + \int \int_{\bar{c} - \Delta}^{\bar{c}} (\bar{\theta} + \Delta/2 - \bar{c}) dF(c_2) dF(c_1) \\
&\quad + \int \int_{c_2 + \Delta}^{\bar{c} + \Delta} (\bar{\theta} - \Delta/2 - c_1 + \Delta) dF(c_1) dF(c_2)
\end{aligned}
$$

$$+\int\int_{\bar{c}+\Delta}^{\bar{c}}(\bar{\theta}-\Delta/2-\bar{c})dF(c_1)dF(c_2),$$

since the final auction price is $p_i = c_j \pm \Delta$ or $\bar{c}$, whichever is lower.

Taking the derivative with respect to Δ we obtain:

$$\begin{aligned}\frac{d}{d\Delta}\Pi(\bar{\theta},\Delta,\Delta) &= \int\int_{c_1-\Delta}^{\bar{c}-\Delta}-\frac{1}{2}\,dF(c_2)dF(c_1)+\int\int_{\bar{c}-\Delta}^{\bar{c}}\frac{1}{2}\,dF(c_2)dF(c_1)\\ &\quad+\int\int_{c_2+\Delta}^{\bar{c}+\Delta}\frac{1}{2}\,dF(c_1)dF(c_2)+\int\int_{\bar{c}+\Delta}^{\bar{c}}-\frac{1}{2}\,dF(c_1)dF(c_2).\end{aligned}$$

When $\Delta \geq 0$, this simplifies to

$$\begin{aligned}\frac{d}{d\Delta}\Pi(\bar{\theta},\Delta,\Delta) &= -1/2\,\Pr(1\text{ wins and }c_2<\bar{c}-\Delta)+1/2\,\Pr(2\text{ wins})\\ &\quad+1/2\,\Pr(1\text{ wins and }c_2>\bar{c}-\Delta).\end{aligned}$$

The derivative of that expression with respect to Δ is

$$\begin{aligned}\frac{d^2}{d\Delta^2}\Pi(\bar{\theta},\Delta,\Delta) &= -1/2f_{c_1-c_2}(\Delta)+1/2f_{c_1-c_2}(\Delta)+1/2f(\bar{c}-\Delta)\\ &= -f_{c_1-c_2}(\Delta)+1/2f(\bar{c}-\Delta)\end{aligned}$$

Similarly, when $\Delta < 0$ we obtain $\frac{d^2}{d\Delta^2}\Pi(\bar{\theta},\Delta,\Delta) = -f_{c_1-c_2}(\Delta)+1/2f(\bar{c}+\Delta)$.

By assumption, these expressions are negative. So Π is concave. By the Jensen inequality, expected profits are higher when the buyer does not learn Δ. $\square$

Proof of Lemma 1: If i wins the auction and j is the second strongest bidder, i receives a price $p_i = c_j + \theta_i - \theta_j$ and delivers a product worth θ_i to the buyer. This happens if $\theta_i - c_i > \theta_j - c_j > \theta_k - c_k$, for all $k \neq i, j$. So we can write the buyer's profits as follows:

$$\Pi = \int \sum_{i,j\neq i}(\theta_j - c_j)\Pr(\theta_i - c_i > \theta_j - c_j > \theta_k - c_k, \forall k \neq i, j)dF(c)$$

The derivative of this expression with respect to θ_1 has three terms: the first two come from when 1 plays the role of j: $\int\sum_i \Pr(\theta_i - c_i > \theta_1 - c_1 > \theta_k - c_k$, for all $k \neq i, 1)$ and $\int\sum_i(\theta_1 - c_1)\frac{\partial}{\partial\theta_1}\Pr(\theta_i - c_i > \theta_1 - c_1 > \theta_k - c_k$, for all $k \neq i, 1)$. The second corresponds to terms where 1 is first: $\int\sum_j(\theta_j -$

$c_j)\frac{\partial}{\partial\theta_1}\Pr(\theta_1 - c_1) > \theta_j - c_j > \theta_k - c_k$, for all $k \neq 1, j)$. Now, the ordering of $\theta_i - c_i$'s only changes in regions of the c-space where there is a tie; so derivatives of the probability terms vanish everywhere but when $\theta_1 - c_1 = \theta_i - c_i$. So the second and third terms are equal except for an opposing sign, and cancel each other. □

With lemma 1 we can investigate the curvature of Π as a function of $\Delta_{ij} = \theta_i - \theta_j$. It is enough to specialize to the case with 3 sellers; the n-seller case is similar, but with a more cumbersome notation. We can write

$$\begin{aligned}
2\frac{\partial}{\partial\Delta_{12}}\Pi &= \Pr(1 \text{ is runner-up}) - \Pr(2 \text{ is runner-up}) \\
&= \int \{F(x+\theta_2-\theta_1)[1-F(x+\theta_3-\theta_1)] \\
&\quad + F(x+\theta_3-\theta_1)[1-F(x+\theta_2-\theta_1)]\}dF(x) \\
&\quad - \int \{F(x+\theta_1-\theta_2)[1-F(x+\theta_3-\theta_2)] \\
&\quad + F(x+\theta_3-\theta_2)[1-F(x+\theta_1-\theta_2)]\}dF(x).
\end{aligned}$$

When $\theta_3 << \theta_1, \theta_2$,

$$\begin{aligned}
2\frac{\partial}{\partial\Delta_{12}}\Pi &\simeq \int \{F(x+\theta_2-\theta_1)[1-0]+0\}dF(x) \\
&\quad - \int \{F(x+\theta_1-\theta_2)[1-0]+0\}dF(x) \\
&= \int \{F(x-\Delta_{12}) - F(x+\Delta_{12})\}dF(x),
\end{aligned}$$

a decreasing function of Δ_{12}. On the other hand, when $\theta_3 >> \theta_1, \theta_2$,

$$\begin{aligned}
2\frac{\partial}{\partial\Delta_{12}}\Pi &\simeq \int 0 + [1-F(x+\theta_2-\theta_1)]dF(x) \\
&\quad - \int 0 + [1-F(x+\theta_1-\theta_2)]dF(x) \\
&= \int \{F(x+\Delta_{12}) - F(x-\Delta_{12})\}dF(x),
\end{aligned}$$

a function that is now increasing in Δ_{12}.

References

Armstrong, M. (1996): "Multiproduct Nonlinear Pricing," *Econometrica*, 64(1), 51–75.

Athey, S., and J. Levin (2001): "Information and Competition in U.S. Forest Service Timber Auctions," *Journal of Political Economy*, 109(2), 375–417.

Bajari, P. (1999): "Sealed-Bid Auction with Asymmetric Bidders: Theory and Computation," Stanford University.

Banks, J. S., and J. Sobel (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647–661.

Burt, D. N. (1984): *Proactive procurement : the key to increased profits, productivity, and quality*. Prentice-Hall, Englewood Cliffs, N.J.

Che, Y.-K. (1993): "Design Competition through Multidimensional Auctions," *Rand Journal of Economics*, 24(4), 668–680.

Cho, I.-K., and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179–220.

Coase, R. H. (1972): "Durability and Monopoly," *Journal of Law and Economics*, 15(1), 143–149.

Griesmer, J. H., R. E. Levitan, and M. Shubik (1967): "Towards a Study of the Bidding Processes Part IV — Games with Unknown Costs," *Naval Research Logistics Quarterly*, 14(4), 415–434.

Grossman, S. J. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24(3), 461–483.

Grossman, S. J., and M. Perry (1986): "Perfect Sequential Equilibrium," *Journal of Economic Theory*, 39(1), 97–119.

Grossman, S. J., and J. E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), 393–408.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39, 155–190.

ISM (2002): *Enterprise Supply Management: Delivering Pure Profit.* Perfect Commerce and Dun & Bradstreet, Palo Alto, CA, CD-ROM.

Jackson, M. O., and J. M. Swinkels (1999): "Existence of Equilibrium in Auctions and Discontinuous Bayesian Games: Endogenous and Incentive Compatible Sharing Rules," Caltech and Washington University in St. Louis.

Kelman, S. (1990): *Procurement and Public Management: The Fear of Discretion and the Quality of Government Performance.* The AEI Press, Washington, D.C.

Laffont, J.-J., E. Maskin, and J.-C. Rochet (1987): "Optimal Nonlinear Pricing with Two-Dimensional Characteristics," in *Information, Incentives and Economic Mechanisms*, ed. by T. Groves, R. Radner, and S. Reiter, chap. 8, pp. 256–267. University of Minnesota Press, Minneapolis.

Lebrun, B. (1996): "Existence of Equilibrium in First Price Auctions," *Economic Theory*, 7, 421–443.

Maskin, E., and J. Riley (2000a): "Asymmetric Auctions," *Review of Economic Studies*, 67(3), 413–438.

——— (2000b): "Equilibrium in Sealed High Bid Auctions," *Review of Economic Studies*, 67(3), 439–454.

McAfee, R. P., and J. McMillan (1988): "Multidimensional Incentive Compatibility and Mechanism Design," *Journal of Economic Theory*, 46(2), 335–354.

——— (1989): "Government Procurement and International Trade," *Journal of International Economics*, 26(3/4), 921–943.

Milgrom, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

——— (1989): "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3(3), 3–22.

Milgrom, P. R., and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583–601.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73.

Rezende, L. B. (2000): "Modern Manufacturing and Competition," Stanford University.

Rogerson, W. P. (1990): "Quality vs. Quantity in Military Procurement," *American Economic Review*, 80(1), 83–92.

Simon, L., and W. Zame (1990): "Discontinuous Games and Endogenous Sharing Rules," *Econometrica*, 58, 861–872.

——— (1999): "Discontinuous Games and Incomplete Information," UCLA and Berkeley.

Zheng, C. Z. (2000): "Optimal Auction in a Multidimensional World," CMS-EMS Discussion Paper 1282, Northwestern University.