MANAGEMENT SCIENCE
Vol. 55, No. 5, May 2009, pp. 752–762
ISSN 0025-1909 | EISSN 1526-5501 | 09 | 5505 | 0752

informs®

DOI 10.1287/mnsc.1080.0967

Quasi-Robust Multiagent Contracts

Anil Arya
Fisher College of Business, Ohio State University, Columbus, Ohio 43210,
arya@cob.osu.edu

Joel Demski
Fisher School of Accounting, University of Florida, Gainesville, Florida 32611,
joel.demski@cba.ufl.edu

Jonathan Glover, Pierre Liang
Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
{jglover@andrew.cmu.edu, liangj@andrew.cmu.edu}

A criticism of mechanism design theory is that the optimal mechanism designed for one environment can produce drastically different actions, outcomes, and payoffs in a second, even slightly different, environment. In this sense, the theoretically optimal mechanisms usually studied are not "robust." To study robust mechanisms while maintaining an expected utility maximization approach, we study a multiagent model in which the mechanism must be designed before the environment is as well understood as is usually assumed. The particular model is of an auction setting with binary private values. Our main result is that if the prior belief about the correlation in the agents' values is diffuse enough, the optimal Bayesian-Nash auction must also satisfy dominant strategy incentive constraints. Furthermore, when the optimal auction does provide dominant strategy incentives, it takes one of two forms: (i) if perfect correlation and negative correlation are excluded as possibilities, the auction incorporates all information about the prior belief over the possible correlations, and (ii) if either perfect correlation or negative correlation is a possibility, the auction does not incorporate any correlation information and can be described as a modified Vickrey auction.

Key words: auctions; robust mechanisms; multiagent contracts
History: Received July 18, 2005; accepted November 11, 2008, by Stefan Reichelstein, accounting. Published online in *Articles in Advance* February 12, 2009.

1. Introduction

A criticism of mechanism design theory is that the optimal theoretical mechanisms often seem overly fine-tuned to the given environment (e.g., Wilson 1987). A mechanism tailored for one environment can produce drastically different actions, outcomes, and payoffs in a second, even slightly different, environment. In this paper, we seek to construct more "robust" mechanisms.

To identify robust mechanisms, we maintain a traditional expected utility framework but assume the designer knows less at the time she designs the mechanism than is typically assumed. The mechanism is designed taking expectations over the variety of possible environments (agent characteristics) that might subsequently emerge. We assume the mechanism cannot later be fine-tuned. It is either impossible or prohibitively costly to design a different mechanism for each environment.

The particular setting we study is an auction model that, other than the early design assumption, is a special case of the model of Cremer and McLean (1988). The model also applies to multiagent procurement, as in Laffont and Tirole (1993, chap. 7), because a procurement model can be viewed as changing the signs on the transfers and values of an auction model.

In the Cremer-McLean model, the auctioneer can extract the full surplus from the risk-neutral bidders (if the beliefs satisfy a spanning condition), which leads to the following observation:

> In "nearly all" auctions, the seller should be able to extract the full surplus, which implies that asymmetry of information between buyers and sellers should be of no practical importance. Economic intuition and informal evidence (we know of no way to test such a proposition) suggest this result is counterfactual.... Costly information gathering, not explicitly modeled in auction problems, may result in less profitable but vastly simpler auctions being used in practice. (Cremer and McLean 1988, p. 1254)

Correlation in the bidders' values is the key variable in Cremer-McLean and is assumed to be common knowledge in their model at the time of the auction design. To extract the full surplus from the bidders, the transfers go to negative or positive infinity as the correlation goes to zero. This makes the Cremer-McLean model particularly well suited to study robustness by assuming that the auction has to be designed before the correlation is known, because the solution for the known correlation case is so seemingly lacking in robustness.

Our early auction design assumption can be viewed as a seller who will conduct many sales, each with a different set of bidders, using the same auction mechanism. The auction is designed to maximize the seller's expected utility over the variety of possible bidders and bidder correlations that might use the auction in the future. To make the early design assumption meaningful, we do not allow the mechanism designer to adapt the auction to the known correlation once the bidders arrive, either directly from her own knowledge of the correlation or by having the bidders report the correlation. Instead, the bidders' messages are restricted to a willingness to pay based on their private value for the item being sold. Formally, we restrict the message space to have the same cardinality as the set of possible valuations. For tractability, we assume there are two ex ante identical bidders who have private values that can take on one of two values, low or high.

As in the Cremer-McLean model, we assume Bayesian-Nash as the behavioral principle. The spanning condition in the Cremer-McLean model ensures that any optimal Bayesian-Nash auction can be equivalently (in terms of expected payoffs) characterized as one that satisfies dominant strategy incentive constraints.[1] That is, the solution to the Bayesian-Nash auction design problem is not unique, and one of the solutions also satisfies dominant strategy incentive constraints.

In our model, the early design assumption makes the solution unique. Throughout most of the paper, we assume the seller's prior belief is that the correlation in the bidders' values is uniformly distributed on a given interval. In particular, p, the conditional probability that the bidders' values match, is uniformly distributed on $[\bar{p}-e, \bar{p}+e]$.[2] If the mechanism designer's prior is concentrated (e, the dispersion around the mean, is small), the optimal auction is a modified version of the Cremer-McLean auction. The modifications are that (i) the known correlation in the Cremer-McLean model is replaced by the lowest possible correlation and (ii) the unique solution does not satisfy dominant strategy incentive constraints. Under the modified Cremer-McLean auction, only the low-value bidders earn rents (for higher than the minimum correlation). If instead the dominant strategy version of the Cremer-McLean model were used, high-value bidders would also earn rents without reducing the rents earned by low-value bidders, making this a feasible but suboptimal modification. That only low-value bidders earn rents is in contrast to the Cremer-McLean result, in which no bidders earn rents, and in contrast to standard adverse selection models in which only high-type agents earn rents.

If instead the prior about the correlation in the bidders' environments is diffuse enough ($e > [1/4] \cdot [2\bar{p}-1]$), the uniquely optimal Bayesian-Nash auction must satisfy dominant strategy incentive constraints. More specifically, for intermediate levels of prior dispersion ($[1/4][2\bar{p}-1] \leq e < [1/2][2\bar{p}-1]$), the auction incorporates all information about the prior over the possible correlations in providing bidders with dominant strategy incentives to bid truthfully. With still higher dispersion ($e \geq [1/2][2\bar{p}-1]$), the optimal auction continues to provide dominant strategy incentives; in so doing, however, it does not incorporate any correlation information (is correlation free) and can be described as a modified Vickrey auction. Equivalently stated, if the prior is diffuse enough ($e > [1/4][2\bar{p}-1]$) and incorporates either perfect correlation or negative correlation as a possibility, the correlation-free dominant strategy auction is optimal. Under this second dominant strategy auction, only the high-value bidders earn rents, which can be interpreted as reproducing a standard adverse selection result. In contrast, under the first dominant strategy auction, both low- and high-value bidders earn rents. Low-value bidders earn rents unless the correlation is at its minimum. High-value bidders earn rents unless the correlation is at its maximum.

Under a uniform distribution over the possible correlations, it is optimal to induce all bidders with the same value to pool in the sense that they report the same message (regardless of their correlation). By means of a discrete version of our model (correlation with three possible realizations), we relax the uniform distribution assumption. As in the earlier uniform setup, the modified Cremer-McLean and dominant strategy auctions are the only ones that emerge as optimal. If a dominant strategy auction turns out to be optimal, it is again optimal to induce bidders with the same value to report the same message. However, when the modified Cremer-McLean auction turns out to be optimal, it is sometimes optimal to have high-value but low-correlation bidders pool with (all of the) low-value bidders—this happens if the probability distribution is sufficiently concentrated on the middle correlation level.

[1] Mookherjee and Reichelstein (1992) also study settings in which Bayesian-Nash incentive constraints can be replaced by dominant strategy incentive constraints. Key assumptions are risk neutrality, uncorrelated types, and the single crossing property. Mookherjee and Reichelstein's model can be thought of as studying settings in which second-best outcomes are to be implemented; the first-best is implemented in the Cremer-McLean model.

[2] Notice that p is closely related to the correlation in the bidders' valuations, which is $p^2-(1-p)^2 = 2p-1$.

The intuition for our results is as follows. The early design assumption (coupled with blocked communication to make the early design assumption meaningful) creates an inextricable information asymmetry between the early designer and the future users of the auction (including the seller herself). If this information asymmetry is small, the optimal approach is "close" to the standard (known correlation) solution for the correlated-type case and is a modified Cremer-McLean auction. (Assume for the moment it is also optimal to induce all bidders with the same value to pool.)

If the auction designer continues to use the modified Cremer-McLean auction while the lower bound on the correlation goes to zero, the transfers received from low-value bidders go to negative or positive infinity (depending on the other bidder's report). When the correlation is higher than the lower bound, these transfers result in low-value but high-correlation bidders earning large rents. To avoid these large rents, it becomes optimal for the seller to use the information about the correlation less aggressively. The optimal less aggressive use of the correlation is a dominant strategy auction. In this case, for intermediate levels of information asymmetry (moderate e-values), the optimal approach is to stay close to the standard solution for the uncorrelated-type case and use a dominant strategy auction. When the information asymmetry is sufficiently high (large e-values), any use of information about the correlation becomes infeasible, and the optimal approach is exactly the standard one for the uncorrelated-type case.

Finally, the assumed pooling behavior is optimal under a uniform distribution over the correlation or if the prior is sufficiently diffuse in the discrete version of our model. If instead the prior is sufficiently concentrated, it can be optimal to use the Cremer-McLean solution with a further modification: high-value but low-correlation bidders are induced to pool with all low-value bidders. These high-value but low-correlation bidders are not fully exploited, so the remaining high-value bidders can be. Again, a dominant strategy auction emerges as optimal if the prior is sufficiently diffuse. The optimal auction continues to be either close to the standard solution for the correlated-type case or close to the standard solution for the uncorrelated-type case. The broader message is that our approach to robustness seems to lead to mechanisms that are either modified versions of the fine-tuned ones or qualitatively different in a way that minimizes the importance of the variable(s) the mechanism has to be robust to (correlation in our model). Whether a qualitative middle ground emerges in other settings is an open question.

Existing approaches to robustness in the accounting literature include Reichelstein (1997) and Dutta and Reichelstein (2002), who study the design of optimal performance measures in investment settings. They use robustness as a way to choose between multiple optimal performance measures. In contrast, we build robustness into the objective function of the principal by assuming the contract has to be designed at an earlier point than is typically assumed.[3]

Bergemann and Morris (2005) study a general implementation model and require the mechanism to be robust to all possible beliefs that agents might have, which leads in "standard" cases to dominant strategy mechanisms. Bergemann and Morris (2008) study ex post implementation—they confine attention to Bayesian-Nash equilibria of games of incomplete information that are also Nash equilibria if the agents' information is instead complete (i.e., if the agents' knew each other's types). Although Bergemann and Morris focus on general settings and all possible beliefs, our focus is on a particular setting and optimal mechanism that allow for some but not complete variation in the agents' beliefs. That is, our approach can be viewed as something of an intermediate one between the standard (fine-tuned) approach and Bergemann and Morris's.

The Cremer-McLean result on extracting the full surplus relies on a link between a bidder's value and his beliefs about the other bidders. In large type spaces, extracting the full surplus is no longer possible. For example, in Parreiras (2005), type includes not just the realization of a signal but also the precision of the system that generates the signal. Independence along the precision dimension of the type forces the seller to leave rents for the bidders. McLean and Postlewaite (2001) show that bidders with identical beliefs but different valuations will also earn rents. Neeman (2003) takes a simple auction (the English auction) as given and examines its effectiveness (expected price as a percentage of the highest valuation). Neeman studies how the effectiveness of the auction changes with assumptions about the seller's Bayesian sophistication. Lopomo (1998) also studies the English auction and shows it is optimal among "simple sequential auctions."

Perhaps closest to our paper is the Chung and Ely study (2007), which researches the foundations for dominant strategy auctions. They emphasize the role a maxmin objective function can play in making dominant strategy auctions optimal. They show that, under a regularity condition, a dominant strategy auction that completely ignores the correlation is the best way to guard against the worst potential buyer beliefs. They also study a Bayesian setting, but in a way

[3] Oddly, so-called generally accepted accounting principles (GAAP) are thought to be robust, though explicating and verifying this claim remains illusive.

that essentially reproduces their maxmin result. By allowing for subjective beliefs and assigning the seller a particular subjective belief that replicates the earlier maxmin payoff, they directly apply their maxmin result to the Bayesian setting.

In contrast, we maintain the standard expected utility maximization and common prior framework but assume the auction has to be designed earlier than is typically assumed. Furthermore, by restricting attention to a specific setting, we are able to characterize all the possible optimal auctions and the conditions under which each is optimal. Our approach yields the correlation-free dominant strategy auction only in the extreme cases (when the prior is sufficiently diffuse). All other times, either a modified Cremer-McLean auction (one that does not provide dominant strategy incentives) or a correlation-dependent dominant strategy auction is optimal.

The remainder of this paper is organized as follows. Section 2 presents the model. Section 3 presents the results: §3.1 presents the benchmark case in which the seller knows the correlation in the bidders' valuations at the time of auction design; §3.2 details the optimal auction under our early design assumption; §3.3 studies the discrete version of our model that allows for nonuniform distributions over the set of possible correlations. Section 4 concludes the paper.

2. Model

A risk-neutral seller has one unit of a good that must be sold to one of two risk-neutral bidders. The seller does not value the object, so her expected utility is the expected transfers received from the bidders. Bidder i, $i = 1, 2$, values the object being sold at V^i, and his expected utility is the probability of winning the auction times V^i less his expected transfer to the seller. Bidder i's value for the object takes on one of two values—V_L or V_H, $0 \le V_L < V_H$—and this value is privately known by bidder i.

It is common knowledge that V^i is equally likely to be V_L or V_H. The parameter of interest is the probability p that $V^i = V_k$ given $V^j = V_k$. That is, p is the conditional probability that the bidders' values match.

Before p is known, the seller designs a single auction to handle the variety of different possible correlations, knowing that p is uniformly distributed on the interval $[\bar{p} - e, \bar{p} + e]$ and that the average correlation is positive; i.e., $\bar{p} > 1/2$. The same auction is to be used for each correlation. This can be viewed as a seller who will conduct many sales, each with a different set of bidders, using the same auction mechanism. Designing an auction for each situation (either by waiting until the circumstance arises or by designing a complex menu of auctions) is prohibitively costly. To make the general-purpose auction design meaningful, we restrict the message space of each bidder to being binary. For ease of interpretation, we use $\{V_L, V_H\}$ as the message space.

Although the generic auction has to be designed before the correlation in the bidders' environments is known, by the time a specific application arises, the correlation is common knowledge among all players. For bidder i who observes V_k and p, his reporting strategy is denoted by $s^i(V_k, p)$. That is, s^i is mapping from bidder i's private information to his message space $\{V_L, V_H\}$.

The seller's objective is to maximize the expected transfers from the bidders, $t^1 + t^2$, while ensuring the following: the bidders' reporting strategies comprise a Bayesian-Nash equilibrium in their subgame (incentive compatibility); the equilibrium provides each bidder with an expected utility of at least zero (individual rationality); and the probability of the bidders winning the item, w^1 and w^2, sums to 1. This design problem is formulated in Program (P).

Program (P)

$$\max_{w^i, t^i, s^i} E_{p\in[\bar{p}-e,\bar{p}+e]} E_{(V^1,V^2)\in(V_k,V_m)_{k,m=L,H}} \left[t^1(s^1(V^1,p),s^2(V^2,p))+t^2(s^1(V^1,p),s^2(V^2,p))\right]$$

subject to

$$E_{V^{-i}\in\{V_L,V_H\}}\big[w^i(s^i(V^i,p),s^{-i}(V^{-i},p))V^i - t^i(s^i(V^i,p),s^{-i}(V^{-i},p)) \mid V^i,p\big] \ge 0 \quad \forall V^i,p,i$$

$$E_{V^{-i}\in\{V_L,V_H\}}\big[w^i(s^i(V^i,p),s^{-i}(V^{-i},p))V^i - t^i(s^i(V^i,p),s^{-i}(V^{-i},p)) \mid V^i,p\big]$$
$$\ge E_{V^{-i}\in\{V_L,V_H\}}\big[w^i(s'^i(V^i,p),s^{-i}(V^{-i},p))V^i - t^i(s'^i(V^i,p),s^{-i}(V^{-i},p)) \mid V^i,p\big] \quad \forall s'^i,V^i,p,i$$

$$w^1(s^1(V^1,p),s^2(V^2,p))+w^2(s^1(V^1,p),s^2(V^2,p))=1$$

$$w^i(\cdot)\ge 0 \quad \forall i$$

$$s^i(V^i,p)\in\{V_L,V_H\} \quad \forall V^i,p,i.$$

Throughout the paper, when convenient, we denote $w^i(V_k, V_m)$ and $t^i(V_k, V_m)$ by w_{km} and t_{km}, respectively. Because of the restriction on communication, the Revelation Principle cannot be directly applied to our setting. Instead, a modified version of the Revelation Principle applies in which the designer optimally partitions the type space, $\{V_L, V_H\} \times [\bar{p} - e, \bar{p} + e]$, into a two-element partition $\{P_1, P_2\}$ message space and conditions the transfers and probability of winning on the element of the partition reported. (Finding the optimal partition is a nontrivial problem.) As under the standard Revelation Principle, it is then without loss of generality to confine attention to an equilibrium that involves reporting strategies that are truthful in the sense that the bidders report the element of the partition that contains their underlying true type, (V^i, p).[4] If the correlation is uniformly distributed (which is assumed throughout most of our analysis), the optimal partition is to have all bidders with the same value report the same message. This claim is proved in the appendix and makes the $\{V_L, V_H\}$ message space particularly natural. In a later subsection of the paper, we provide a discrete version of our model in which another partition is sometimes optimal.

3. Results

3.1. Benchmark

Suppose the seller knows p at the time of designing the auction mechanism. In this case, despite V^i being private, the seller can extract the full surplus (Cremer and McLean 1988). An auction that accomplishes this is presented in Proposition 1. Because the proposition is a special case of Cremer and McLean, the proof is omitted.

Proposition 1. *If the seller knows p at the time of designing the auction, an optimal auction is*

$$w_{LL} = 1/2 \qquad w_{HL} = 1$$
$$w_{LH} = 0 \qquad w_{HH} = 1/2$$

$$t_{LL} = \frac{V_H}{2} - \frac{p^2(V_H - V_L)}{2(2p-1)} \qquad t_{HL} = V_H - \frac{p^2(V_H - V_L)}{2(2p-1)}$$
$$t_{LH} = \frac{p(1-p)(V_H - V_L)}{2(2p-1)} \qquad t_{HH} = \frac{V_H}{2} + \frac{p(1-p)(V_H - V_L)}{2(2p-1)}.$$

Under Proposition 1's auction, the object is awarded to the highest bidder, with ties resolved by giving each bidder an equal chance of winning.[5] Because p is known and there is no blocked communication, the Revelation Principle can be applied to confine attention to direct revelation mechanisms under which truthful revelation is an equilibrium. Although the equilibrium concept is Bayesian-Nash, Cremer and McLean have shown that the combination of risk-neutral bidders and the correlation structure make it possible for the optimal Bayesian-Nash auction to both achieve the first-best solution (extract all the bidders' information rents) and be equivalently (in terms of all expected payoffs) characterized so that truthful bidding is also a dominant strategy equilibrium. The proposition presents this characterization.

Notice that Proposition 1's auction has the undesirable feature of each bidder having to make a transfer even when he is not awarded the object. Furthermore, even when a bidder is awarded the object, his transfer can exceed his value. That is, the transfers do not satisfy ex post individual rationality. Finally, as the correlation in the bidders' environments approaches zero, the transfers for a bidder become arbitrarily large when the other bidder reports a high value and arbitrarily small when the other bidder reports a low value. These observations are summarized in Corollary 1.

Corollary 1. *If p is known at the time of auction design:*

(i) *The optimal Bayesian-Nash auction can be characterized so as to also satisfy dominant strategy incentive constraints.*

(ii) *The bidders earn no rents under the truth-telling equilibrium.*

(iii) *The transfers do not satisfy ex post individual rationality.*

(iv) *As $p \to 1/2$, $t_{kL} \to -\infty$ and $t_{kH} \to +\infty$.*

3.2. Quasi-Robust Auctions

In this subsection, we characterize the optimal Bayesian-Nash auction under our original assumption that the auction is designed before p is known. The solution to Program (P) is presented in Proposition 2. The proof is provided in the appendix.

Proposition 2. *If at the time of designing the auction the seller does not know p but knows only that p is uniformly distributed on $[\bar{p} - e, \bar{p} + e]$, the optimal auction is*

$$w_{LL} = 1/2 \qquad w_{HL} = 1$$
$$w_{LH} = 0 \qquad w_{HH} = 1/2.$$

If $e < (2\bar{p} - 1)/4$ (BN auction)

$$t_{LL} = \frac{V_H}{2} - \frac{(\bar{p} - e)^2(V_H - V_L)}{2[2\bar{p} - 2e - 1]} \qquad t_{HL} = V_H$$
$$t_{LH} = \frac{(\bar{p} - e)(1 - \bar{p} + e)(V_H - V_L)}{2[2\bar{p} - 2e - 1]} \qquad t_{HH} = \frac{V_H}{2}.$$

If $(2\bar{p} - 1)/4 \le e < (2\bar{p} - 1)/2$ (DS1 auction)

$$t_{LL} = \frac{[\bar{p}^2 - e^2]V_L - [(1 - \bar{p})^2 - e^2]V_H}{2[2\bar{p} - 1]} \qquad t_{HL} = t_{LL} + \frac{V_H}{2}$$
$$t_{LH} = \frac{(\bar{p} - e)(1 - \bar{p} - e)(V_H - V_L)}{2[2\bar{p} - 1]} \qquad t_{HH} = t_{LH} + \frac{V_H}{2}.$$

[4] Equivalently, one can reformulate the restriction on blocked communication as a restriction on the outcome set. In this case, one allows for full communication and directly applies the Revelation Principle but under the added constraint that the mechanism uses at most four distinct outcomes.

[5] There are other equilibria in which the bidders are treated asymmetrically (for example, bidder 1 is always awarded the object in the event of a tie), but these asymmetric equilibria are revenue equivalent to treating the ex ante identical agents symmetrically. If instead the bidders were not ex ante identical, treating them asymmetrically (in some cases, even awarding the object to a bidder who does not have the highest value for the object) can be uniquely revenue maximizing (Maskin and Riley 2000).

If $e \geq (2\bar{p}-1)/2$ (DS2 *auction*)

$$t_{LL} = \frac{V_L}{2} \qquad t_{HL} = t_{LL} + \frac{V_H}{2}$$
$$t_{LH} = 0 \qquad t_{HH} = t_{LH} + \frac{V_H}{2}.$$

Given the agents are ex ante identical and that the item must be auctioned, the probability that each bidder is awarded the item can be characterized in a symmetric fashion as in the p-known setting: If agent i's bid is higher, he is assigned the item with probability 1, but in the event of a tie, he is assigned the item with probability 1/2. Also, under the auctions given in Proposition 2, a V_L-bidder finds it optimal to bid V_L and the V_H-bidder finds it optimal to bid V_H for all p-values; i.e., $s^i(V_k, p) = V_k$ for all p. (As we show later in §3.3, a diffuse prior about the correlation is important in ensuring an auction that induces this reporting behavior is optimal.)

The intuition behind the optimal auction is as follows. When the prior dispersion is limited (e is small), the Cremer-McLean auction is optimal, except for two modifications.[6] First, the known p is replaced by its lowest possible value $\bar{p} - e$. Second, it is no longer possible to characterize the optimal auction so that it satisfies dominant strategy constraints. With the first adjustment alone, the Cremer-McLean auction would have V_H-bidders earning rents whenever p is greater than $\bar{p} - e$. Instead, the optimal auction sets $t_{HL} = V_H$ and $t_{HH} = V_H/2$ to ensure that all V_H-bidders earn zero rents. This characterization is also one possible solution for the V_H type even when p is known (i.e., an alternative solution to the auction given in Proposition 1), but it becomes the unique solution when p is unknown. The only incentive compatibility constraint that binds is for a V_H-bidder with the lowest possible p of $\bar{p} - e$.

When the auction is to serve a wider range of correlations ($e \geq (2\bar{p} - 1)/4$), the optimal auction provides dominant strategy incentives. For intermediate levels of prior dispersion, the optimal such auction is (DS1). In this case, the incentive compatibility constraints bind for all V_H-bidders. The least expensive way to ensure that incentive compatibility just holds for V_H-bidders is to have them pay what they would have paid had they instead bid V_L plus $V_H/2$ for the extra half a chance they have of winning the object by increasing their bids from V_L to V_H. As a result, the optimal auction (DS1) satisfies dominant strategy incentive constraints. The solution is again unique. Note that DS1 depends on both $\bar{p}$ and e. That is, DS1 incorporates the common prior information about the correlation.

If dominant strategy incentives are optimal (so $e \geq (2\bar{p} - 1)/4$), but the range of possible p-values is large and includes 1/2 or 1 (correlation of 0 or 1), the demands on the auction to be effective lead the seller to give up on the fine-tuned dominant strategy auction (DS1) and to instead use DS2. DS2 is correlation free (independent of $\bar{p}$ and e) and provides each bidder with dominant strategy incentives. DS2 can be thought of as a modified second-price (Vickrey) auction, where the modification is only to take the known binary set of values $\{V_L, V_H\}$ into account. If one bidder reports V_L and the other V_H, the winner does not just pay V_L (the Vickrey price). Instead, he pays V_L plus a premium $(V_H - V_L)/2$, which is arrived at from the seller's knowledge of the values of V_L and V_H. That is, if the bidding reveals that one bidder values the object more than the other, the seller knows how much more.

From an information rents perspective, the DS2 auction is familiar—only V_H-bidders earn rents. Under the Bayesian-Nash auction, it is only the V_L-bidders who earn rents. Under the DS1 auction, both V_L-bidders and V_H-bidders earn rents. V_L-bidders earn rents unless the correlation is at its minimum. V_H-bidders earn rents unless the correlation is at its maximum. A way to reconcile these results with standard adverse selection models is to recognize that a full representation of a bidder's type in our model is two dimensional (V^i, p).

For $\bar{p} = 2/3$, Figure 1 plots the seller's expected revenue as a function of e. The figure is plotted for normalized V-values of $V_L = 0$ and $V_H = 1$, or, equivalently, the y-axis in the figure can be viewed as the weight placed on V_H in calculating the seller's

Figure 1 Seller's Expected Revenue as a Function of *e*

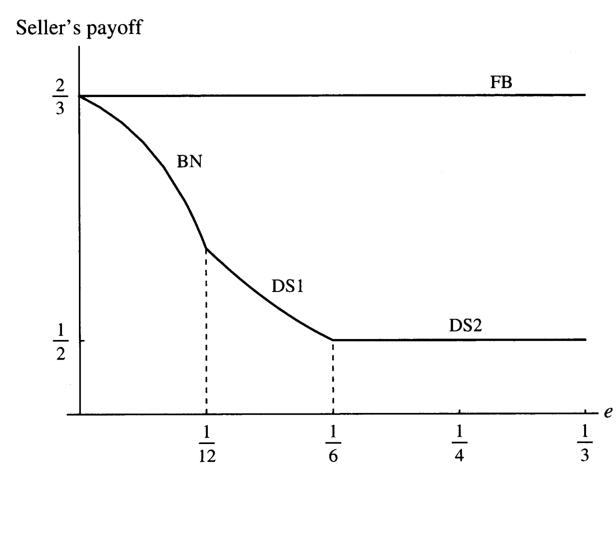

[6] With $e = 0$, the seller's expected revenue under the Bayesian-Nash auction is the same as under the Cremer-McLean auction in Proposition 1—the expected revenue in either case is $[\bar{p}/2]V_L + [1 - \bar{p}/2]V_H$. In effect, in the p-known case, there are multiple ways to characterize the optimal transfers, and the transfers we chose to present in Proposition 1 were those that provided dominant strategy incentives.

revenue. That is, the seller's revenue equals V_H times y-axis coordinate plus V_L times $(1-y)$-axis coordinate. In the p-known case, from Proposition 1, this is the first-best (FB) solution. In the p-unknown case, Proposition 2 applies.

The optimal auction in the p-unknown case has the following properties.

Corollary 2. *If at the time of designing the auction the seller does not know p but knows only that p is uniformly distributed on* $[\bar{p}-e, \bar{p}+e]$, *then*:

(i) *When* $e < (2\bar{p}-1)/4$, *the uniquely optimal Bayesian-Nash auction does not satisfy dominant strategy incentive compatibility constraints. When* $e > (2\bar{p}-1)/4$, *the uniquely optimal Bayesian-Nash auction satisfies dominant strategy incentive compatibility constraints.*

(ii) *When* $e < (2\bar{p}-1)/4$, *only* V_L*-bidders earn rents. When* $e > (2\bar{p}-1)/2$, *only* V_H*-bidders earn rents.*

(iii) *The transfers satisfy ex post individual rationality if and only if* DS2 *is optimal.*[7]

(iv) *The transfers are bounded from above and below.*

A special case of our model is when the seller knows only the lower bound on correlation, i.e., when p is uniformly distributed on $[p', 1]$. In this case, the optimal auction is BN if the seller's information is "good" (i.e., $p' > 2/3$) and DS2 otherwise. In effect, the requirement that the auction be effective even when the correlation is perfect ($p = 1$) necessitates that the seller provide dominant strategy incentives only via the correlation-free auction.

Corollary 3. *If at the time of designing the auction the seller does not know p but knows only that p is uniformly distributed on* $[p', 1]$, *the optimal auction is*

$$w_{LL} = 1/2 \qquad w_{HL} = 1$$
$$w_{LH} = 0 \qquad w_{HH} = 1/2\cdot$$

If $p' > 2/3$ (BN *auction*)

$$t_{LL} = \frac{V_H}{2} - \frac{p'^2(V_H - V_L)}{2[2p'-1]} \qquad t_{HL} = V_H$$
$$t_{LH} = \frac{p'(1-p')(V_H - V_L)}{2[2p'-1]} \qquad t_{HH} = \frac{V_H}{2}.$$

If $p' \le 2/3$ (DS2 *auction*)

$$t_{LL} = \frac{V_L}{2} \qquad t_{HL} = t_{LL} + \frac{V_H}{2}$$
$$t_{LH} = 0 \qquad t_{HH} = t_{LH} + \frac{V_H}{2}.$$

[7] If $e = 1-\bar{p}$, the (DS1) auction reduces to the (DS2) auction. Thus, the ex post individual rationality constraints are satisfied if and only if (i) $e \ge (2\bar{p}-1)/4$ and $e = 1-\bar{p}$ or (ii) $e \ge (2\bar{p}-1)/2$.

One issue we have suppressed thus far in the analysis is the multiple equilibria problem. Under BN, there is a pressing multiple equilibrium problem in that the bidders would both be better off playing an equilibrium that has them always reporting V_L (for all V and p). One could turn to an expanded message space in the hope of eliminating the undesirable equilibrium without creating new equilibria. If an approach similar to that of Ma et al. (1988) is feasible and were used, a large message space would be needed to deal with possible mixed strategy equilibrium, which is inconsistent with our blocked communication assumption. Under either DS1 or DS2, although the equilibrium the auction designer intends the bidders to play is again not unique, it can easily be made unique by reducing the V_H transfers t_{HL} and t_{HH} by an arbitrarily small amount. Hence, the multiple equilibria problem is another reason dominant strategy auctions may be favored, as Demski and Sappington (1984) argue. When the mechanism has to be designed early, turning to a (strict) dominant strategy mechanism will not only deal with the multiple equilibria problem, but it may also do so at little or even no cost to the seller, because such mechanisms can be optimal as general-purpose ones in Bayesian environments.

3.3. Nonuniform Distributions Can Make Pooling Optimal

Under nonuniform seller beliefs, the optimal auction no longer necessarily induces all V_H-bidders to bid V_H. To see this, consider the following discrete version of our model.

Suppose $p \in \{\bar{p}-(1-\bar{p}), \bar{p}, \bar{p}+(1-\bar{p})\} = \{2\bar{p}-1, \bar{p}, 1\}$ with probabilities $\{(1-z)/2, z, (1-z)/2\}$. The parameter z captures how diffuse the prior is; $z = 1/3$ is our earlier uniform prior. In this setting, if z is large enough (so the prior is concentrated), then the optimal auction always entails pooling, in that V_H-bidders with the lowest p-value pool with all V_L-bidders.

In particular, let $\widetilde{\text{BN}}$ denote the auction that is the same as BN except that $\bar{p} - e$ is replaced by $\bar{p}$, so the optimal transfers are $t_{LL} = V_H/2 - \bar{p}^2(V_H - V_L)/(2[2\bar{p} - 1])$, $t_{LH} = \bar{p}(1 - \bar{p})(V_H - V_L)/(2[2\bar{p} - 1])$, $t_{HL} = V_H$, and $t_{HH} = V_H/2$. Under this auction, the V_H-bidders with p of $2\bar{p}-1$ pool with all V_L-bidders. To reinterpret the problem as one in which the modified Revelation Principle discussed at the end of §2 is applied, we can interpret the binary message space as {"My value is V_L or my p is $2\bar{p} - 1$," "My value is V_H and my p is $\bar{p}$ or greater"}.

The optimal auction then takes the following form:

For $\bar{p} \le 5/6$:

If $z \le (3\bar{p}^2 - 4\bar{p} + 2)/((2-\bar{p})\bar{p})$, DS2 is optimal; else, $\widetilde{\text{BN}}$ is optimal.

For $\bar{p} > 5/6$:

If $z \le (4 - 22\bar{p} + 31\bar{p}^2 - 12\bar{p}^3)/((-6 + 11\bar{p} - 4\bar{p}^2)\bar{p})$, BN is optimal; else, $\widetilde{\text{BN}}$ is optimal.

Note that the form of the optimal auction remains unchanged (BN or DS2), but the reporting (pooling) strategies change when the prior is less diffuse. Intuitively, a prior more concentrated on $p = \bar{p}$ makes the $p = 2\bar{p} - 1$ case less significant for the seller, and the optimal auction lets the V_H-bidder with $p = 2\bar{p} - 1$ earn rents by pooling with the V_L-bidders. In the proof of Proposition 2, we allow for arbitrary pooling strategies, but the diffuse (uniform) prior assumption ensures that having V_L-bidders bid V_L and V_H-bidders bid V_H is optimal.

4. Concluding Remarks

This paper addresses the design of mechanisms when the characteristics of the environments in which the mechanisms are to be employed are not (yet) precisely nailed down. In particular, the paper studies the design of a generic auction to be used for a variety of possible bidder pairs and correlations in the bidders' valuations. When there is sufficient uncertainty about the correlation, the seller finds it optimal to provide dominant strategy incentives to the bidders. The early design assumption creates an information asymmetry between the early designer and the future bidders who will arrive to use the in-place auction format. It is this information asymmetry that leads to the optimality of dominant strategy incentives in a fully Bayesian setting. Our search for robust mechanisms can be viewed as part of the Wilson Program:

> This brings me to a point I wish to emphasize: The optimal trading rule for a direct revelation game is specialized to a particular environment. For example, the rule generally depends on the agents' probability assessments about each other's private information. If left in this form, therefore, the theory is mute on one of the most basic problems challenging theory. I refer to the problem of explaining the prevalence of a few simple trading rules in most of the commerce conducted.... The rules of these markets are not changed daily as the environment changes; rather they persist as stable, viable institutions. As a believer that practice advances before theory, and that the task of theory is to explain how it is that practitioners are (usually) right, I see a plausible conjecture: These institutions survive because they employ trading rules that are efficient for a wide class of environments. (Wilson 1987, pp. 36–37)

Correlation is only one of many possible aspects of the environment that might not be as well understood at the time of mechanism design as is typically assumed. In this sense, our auctions can be viewed as "quasi-robust" rather than robust. Allowing for a greater variety of possible environments (e.g., the set of possible values may be uncertain, the number of bidders may be indeterminate, and private values may be correlated with each other as well as with commonly observed macro variables) reveals interesting and important avenues that are left out of our analysis. The implied underutilization of information in the face of robustness considerations also surfaces in Evans' (2008) contracting setting, where unbounded renegotiation possibilities are introduced. In broader terms, organizational arrangements, including sourcing, human resource policies, and accounting itself all have a robustness dimension, a dimension that to date largely escapes the mechanism design orientation.

We expect our approach to robust mechanism design—expected utility maximization under blocked communication—to lend itself quite naturally to a variety of settings. Furthermore, we conjecture that the optimal robust mechanism will continue to be either a modified version of the standard solution or a qualitatively different one that minimizes the importance of the information asymmetry (although not necessarily a dominant strategy mechanism).[8] Close to the present paper is the principal-multiagent model of Demski and Sappington (1984), which again has binary cost/value parameters but allows for agent risk aversion. From numerical examples, similar results to those obtained in this paper appear to arise in the modified Demski-Sappington model.

Glover et al. (2006) study the robustness problem in a moral hazard setting under agent risk neutrality and limited liability. The information asymmetry is about the informativeness of the performance measures. When robustness is a minor concern, the solution is a modified version of the standard one: a bang-bang contract in which a large bonus payment is made when the best realization of the performance measures occurs; otherwise, the bonus payment is zero. The main result is that it can be optimal to ignore performance measures that are informative but subject to a larger robustness concern. The standard condition involving likelihood ratios is replaced by one involving both likelihood ratios and the agent's marginal productivity in the performance measures. Under agent risk aversion, this revised condition leads to nonmonotonic contracts when they would be monotonic in the standard model.

Acknowledgments

The authors thank conference participants at the 2005 Accounting Research Workshop in Bern, the 2005 Carnegie Mellon Accounting Mini Conference, the 2005 University of Chicago–University of Minnesota Accounting Theory Conference, and the 2006 Symposium on Accounting Research at the University of Houston; workshop participants at the Catholic University of Portugal in Lisbon and the Norwegian School of Economics and Business Administration; and Rick Antle, Tim Baldenius, Ron Dye, John Fellingham, Hans Frimor, Thomas Hemmer, Yuji Ijiri, Carolyn Levine, Brian Mittendorf, Paul Newman, David

[8] Demski and Dye's (1999) result that it is always optimal to fine-tune a linear contract by adding nonlinear terms in a LEN-style moral hazard setting begs a parallel robustness issue. (LEN stands for linear contracts, exponential utility, and normally distributed noise terms.)

Sappington, Doug Schroeder, Jack Stecher, Shyam Sunder, and particularly Stefan Reichelstein and two anonymous referees for helpful comments. The first author gratefully acknowledges assistance from the John J. Gerlach Chair. The third author gratefully acknowledges financial support from the Dean's Summer Fund.

Appendix

Proof of Proposition 2. We first characterize the optimal auction among those that have each bidder reporting V_L when his type is V_L and V_H when his type is V_H. We then verify that no auction that has the bidders reporting differently increases the seller's expected revenue.

Given the symmetric agent setup and the fact that the item must be auctioned, the probability of each bidder being awarded the item can be characterized in a symmetric fashion: if agent i's bid is higher, he is assigned the item with probability 1, and in the event of a tie, he is assigned the item with probability 1/2. To determine the optimal transfers among the auctions we initially restrict attention to, we reformulate (P) in two ways. First, we replace $s^i(V_k, p)$ with V_k. This yields the familiar individual rationality (IR_k) and incentive compatibility (IC_k) constraints for the k-value bidder. Second, we delete redundant constraints using the fact that satisfying the (IR_k) and (IC_k) constraints for $p=\bar{p}-e$ and $p=\bar{p}+e$ ensures that the constraints hold for all other p. The resulting program (P) is given below:

$$\max_t \quad 2\times\left(\frac{\bar{p}}{2}t_{LL}+\frac{1-\bar{p}}{2}t_{LH}+\frac{1-\bar{p}}{2}t_{HL}+\frac{\bar{p}}{2}t_{HH}\right)$$

$$\text{s.t.}\quad [\bar{p}-e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}-e)](0-t_{LH})\geq 0 \quad \mathrm{IR}_L(\bar{p}-e)$$

$$[\bar{p}-e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}-e)](0-t_{LH})$$
$$\geq[\bar{p}-e](V_L-t_{HL})+[1-(\bar{p}-e)]\left(\frac{V_L}{2}-t_{HH}\right) \quad \mathrm{IC}_L(\bar{p}-e)$$

$$[\bar{p}+e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}+e)](0-t_{LH})\geq 0 \quad \mathrm{IR}_L(\bar{p}+e)$$

$$[\bar{p}+e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}+e)](0-t_{LH})$$
$$\geq[\bar{p}+e](V_L-t_{HL})+[1-(\bar{p}+e)]\left(\frac{V_L}{2}-t_{HH}\right) \quad \mathrm{IC}_L(\bar{p}+e)$$

$$[\bar{p}-e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}-e)](V_H-t_{HL})\geq 0 \quad \mathrm{IR}_H(\bar{p}-e)$$

$$[\bar{p}-e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}-e)](V_H-t_{HL})$$
$$\geq[\bar{p}-e](0-t_{LH})+[1-(\bar{p}-e)]\left(\frac{V_H}{2}-t_{LL}\right) \quad \mathrm{IC}_H(\bar{p}-e)$$

$$[\bar{p}+e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}+e)](V_H-t_{HL})\geq 0 \quad \mathrm{IR}_H(\bar{p}+e)$$

$$[\bar{p}+e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}+e)](V_H-t_{HL})$$
$$\geq[\bar{p}+e](0-t_{LH})+[1-(\bar{p}+e)]\left(\frac{V_H}{2}-t_{LL}\right) \quad \mathrm{IC}_H(\bar{p}+e).$$

Denoting the objective function of (P) by $f(t)$, the left-hand side less the right-hand side of the ith constraint by $g_i(t)$, and the Lagrange multiplier on the ith constraint by λ_i, the Lagrangian is $f(t)+\sum_{i=1}^{8}\lambda_i g_i(t)$. The first-order condition (FOC) of the Lagrangian with respect to each of the four transfers are as follows:

$$\bar{p}-(\lambda_1+\lambda_2)(\bar{p}-e)-(\lambda_3+\lambda_4)(\bar{p}+e)+\lambda_6(1-\bar{p}+e)$$
$$+\lambda_8(1-\bar{p}-e)=0 \quad \mathrm{FOC}-t_{LL}$$
$$(1-\bar{p})-(\lambda_1+\lambda_2)(1-\bar{p}+e)-(\lambda_3+\lambda_4)(1-\bar{p}-e)$$
$$+\lambda_6(\bar{p}-e)+\lambda_8(\bar{p}+e)=0 \quad \mathrm{FOC}-t_{LH}$$
$$(1-\bar{p})+\lambda_2(\bar{p}-e)+\lambda_4(\bar{p}+e)-(\lambda_5+\lambda_6)(1-\bar{p}+e)$$
$$-(\lambda_7+\lambda_8)(1-\bar{p}-e)=0 \quad \mathrm{FOC}-t_{HL}$$
$$\bar{p}+\lambda_2(1-\bar{p}+e)+\lambda_4(1-\bar{p}-e)-(\lambda_5+\lambda_6)(\bar{p}-e)$$
$$-(\lambda_7+\lambda_8)(\bar{p}+e)=0 \quad \mathrm{FOC}-t_{HH}.$$

The optimal solution is one that (i) satisfies the eight constraints in (P), (ii) satisfies the four first-order conditions, (iii) satisfies the eight complementary slackness conditions $\lambda_i g_i(t)=0$, and (iv) prescribes nonnegative multipliers.

The solution listed below satisfies (i)–(iv) for $e<(2\bar{p}-1)/4$; hence, for these e-values, it is optimal. Under this solution, denoted BN, the four transfers are obtained by solving the following four (binding) constraints in (P): $\mathrm{IR}_L(\bar{p}-e)$, $\mathrm{IR}_H(\bar{p}-e)$, $\mathrm{IC}_H(\bar{p}-e)$, and $\mathrm{IR}_H(\bar{p}+e)$.

The BN solution

$$t_{LL}=\frac{V_H}{2}-\frac{(\bar{p}-e)^2(V_H-V_L)}{2[2\bar{p}-2e-1]} \qquad t_{HL}=V_H$$
$$t_{LH}=\frac{(\bar{p}-e)(1-\bar{p}+e)(V_H-V_L)}{2[2\bar{p}-2e-1]} \qquad t_{HH}=\frac{V_H}{2}$$
$$\lambda_1=\frac{2\bar{p}-e-1}{2\bar{p}-2e-1} \qquad \lambda_2=0 \qquad \lambda_3=0 \qquad \lambda_4=0$$
$$\lambda_5=\frac{2\bar{p}-4e-1}{2[2\bar{p}-2e-1]} \qquad \lambda_6=\frac{e}{2\bar{p}-2e-1} \qquad \lambda_7=1/2 \qquad \lambda_8=0.$$

Notice that under BN the condition $\lambda_5\geq 0$ yields the upper bound on e. Thus, when $e\geq(2\bar{p}-1)/4$, the optimal solution changes. For $(2\bar{p}-1)/4\leq e<(2\bar{p}-1)/2$, the solution listed below is optimal—it satisfies (i)–(iv). Under this solution, denoted DS1, the four transfers are obtained by solving the following four (binding) constraints in (P): $\mathrm{IR}_L(\bar{p}-e)$, $\mathrm{IC}_H(\bar{p}-e)$, $\mathrm{IR}_H(\bar{p}+e)$, and $\mathrm{IC}_H(\bar{p}+e)$.

The DS1 solution

$$t_{LL}=\frac{[\bar{p}^2-e^2]V_L-[(1-\bar{p})^2-e^2]V_H}{2[2\bar{p}-1]} \qquad t_{HL}=t_{LL}+\frac{V_H}{2}$$
$$t_{LH}=\frac{(\bar{p}-e)(1-\bar{p}-e)(V_H-V_L)}{2[2\bar{p}-1]} \qquad t_{HH}=t_{LH}+\frac{V_H}{2}$$
$$\lambda_1=\frac{2\bar{p}+2e-1}{2\bar{p}-1} \qquad \lambda_2=0 \qquad \lambda_3=0 \qquad \lambda_4=0$$
$$\lambda_5=0 \qquad \lambda_6=1/2 \qquad \lambda_7=\frac{[2\bar{p}-1]-2e}{2\bar{p}-1} \qquad \lambda_8=\frac{4e-[2\bar{p}-1]}{2[2\bar{p}-1]}.$$

Under DS1, the conditions $\lambda_7\geq 0$ and $\lambda_8\geq 0$ yield the upper and lower bounds on e, respectively. Thus, when

$e \geq (2\bar{p}-1)/2$, $\lambda_7 < 0$, and the solution again changes. For these values of e, the solution listed next, denoted DS2, is optimal—it satisfies (i)–(iv). Under this solution, the four transfers are obtained by solving the following four (binding) constraints in (P): $\text{IR}_L(\bar{p}-e)$, $\text{IR}_L(\bar{p}+e)$, $\text{IC}_H(\bar{p}-e)$, and $\text{IC}_H(\bar{p}+e)$.

The DS2 solution

$$t_{LL} = \frac{V_L}{2} \qquad t_{HL} = t_{LL} + \frac{V_H}{2}$$
$$t_{LH} = 0 \qquad t_{HH} = t_{LH} + \frac{V_H}{2}$$
$$\lambda_1 = \frac{2\bar{p}+2e-1}{2e} \quad \lambda_2 = 0 \quad \lambda_3 = \frac{2e-[2\bar{p}-1]}{2e} \quad \lambda_4 = 0$$
$$\lambda_5 = 0 \quad \lambda_6 = 1/2 \quad \lambda_7 = 0 \quad \lambda_8 = 1/2.$$

Having encompassed the entire range of e-values, we have identified the solution assuming inducing truthful reporting by the bidders is optimal. We still need to verify that there is no auction that induces a bidder to report V_L when his value is V_H or vice versa for some realizations of p that improves the seller's expected revenue.

The preferences are such that if a V_k-bidder finds it optimal to report V_k for p and V_m for $p' > p$, then he will also prefer to report V_m for all $p'' > p'$. Similarly, if a V_k-bidder finds it optimal to report V_k for p and V_m for $p' < p$, then he will also prefer to report V_m for all $p'' < p'$. Hence, one can confine attention to cutoff pooling strategies. Because the pressing incentive problem is one in which a V_H-bidder wants to mimic the V_L-bidders for all p below a cutoff, we focus on that case providing the optimal contracts for each parameter region and arguing that each is dominated by one of the contracts we have presented above. We have also derived the optimal contracts for the other pooling cases and verified that they are dominated by one of the above contracts but have not included these other cases in the proof. These other pooling cases are less interesting than the one we present next in the sense that the solutions are more obviously dominated. (The solutions for these other cases are available from the authors upon request.)

Suppose the V_L-bidder reports V_L for all p, and the V_H-bidder reports V_L for $p \in [\bar{p}-e, \tilde{p})$ and reports V_H for $p \in [\tilde{p}, \bar{p}+e]$. For any given $\tilde{p}$, the optimal transfers that induce such reporting are determined by solving program $(\tilde{P})$, presented next; the program uses the fact that satisfying (IR_k) and (IC_k) for $p = \bar{p}-e$, $\bar{p}+e$, and $\tilde{p}$ ensures the constraints hold for all other p. Also, in $(\tilde{P})$, $\hat{p} = (\bar{p}+e+\tilde{p})/2$, i.e., $\hat{p}$ is the mean p-value over the interval $p \in [\tilde{p}, \bar{p}+e]$ where the bidders report truthfully:

$$\max_t \quad 2\times\left(\left[\frac{\tilde{p}-(\bar{p}-e)}{2e}\right]t_{LL} + \left[\frac{(\bar{p}+e)-\tilde{p}}{2e}\right] \cdot\left[\frac{\hat{p}}{2}t_{LL}+\frac{1-\hat{p}}{2}t_{LH}+\frac{1-\hat{p}}{2}t_{HL}+\frac{\hat{p}}{2}t_{HH}\right]\right)$$

s.t.
$$\frac{V_L}{2}-t_{LL} \geq 0 \quad \text{IR}_L(\bar{p}-e)$$
$$\frac{V_L}{2}-t_{LL} \geq V_L - t_{HL} \quad \text{IC}_L(\bar{p}-e)$$
$$[\bar{p}+e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}+e)](0-t_{LH}) \geq 0 \quad \text{IR}_L(\bar{p}+e)$$
$$[\bar{p}+e]\left(\frac{V_L}{2}-t_{LL}\right)+[1-(\bar{p}+e)](0-t_{LH}) \geq [\bar{p}+e](V_L-t_{HL})+[1-(\bar{p}+e)]\left(\frac{V_L}{2}-t_{HH}\right) \quad \text{IC}_L(\bar{p}+e)$$
$$\tilde{p}\left(\frac{V_L}{2}-t_{LL}\right)+[1-\tilde{p}](0-t_{LH}) \geq 0 \quad \text{IR}_L(\tilde{p})$$
$$\tilde{p}\left(\frac{V_L}{2}-t_{LL}\right)+[1-\tilde{p}](0-t_{LH}) \geq \tilde{p}(V_L-t_{HL})+[1-\tilde{p}]\left(\frac{V_L}{2}-t_{HH}\right) \quad \text{IC}_L(\tilde{p})$$
$$\frac{V_H}{2}-t_{LL} \geq 0 \quad \text{IR}_H(\bar{p}-e)$$
$$\frac{V_H}{2}-t_{LL} \geq V_H - t_{HL} \quad \text{IC}_H(\bar{p}-e)$$
$$[\bar{p}+e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}+e)](V_H-t_{HL}) \geq 0 \quad \text{IR}_H(\bar{p}+e)$$
$$[\bar{p}+e]\left(\frac{V_H}{2}-t_{HH}\right)+[1-(\bar{p}+e)](V_H-t_{HL}) \geq [\bar{p}+e](0-t_{LH})+[1-(\bar{p}+e)]\left(\frac{V_H}{2}-t_{LL}\right) \quad \text{IC}_H(\bar{p}+e)$$
$$\tilde{p}\left(\frac{V_H}{2}-t_{HH}\right)+[1-\tilde{p}](V_H-t_{HL}) \geq 0 \quad \text{IR}_H(\tilde{p})$$
$$\tilde{p}\left(\frac{V_H}{2}-t_{HH}\right)+[1-\tilde{p}](V_H-t_{HL}) \geq \tilde{p}(0-t_{LH})+[1-\tilde{p}]\left(\frac{V_H}{2}-t_{LL}\right) \quad \text{IC}_H(\tilde{p}).$$

For any given $\tilde{p}$, the above program is linear; thus, the solution to $(\tilde{P})$ can be characterized using the standard Lagrangian approach as before. In presenting this solution, we make use of the critical multiplier values of C_1 and C_2. In particular, let $C_1 = (e^2-3\bar{p}+\bar{p}^2+e(5+2\bar{p}-8\tilde{p})+(3-\tilde{p})\tilde{p})/(4e(1-2\tilde{p}))$ and $C_2 = (\bar{p}-\tilde{p}-e(3-4\tilde{p}))/(2e(\bar{p}+\tilde{p}+e-1))$.

If $C_1 > 0$ and $C_2 > 0$ ($\widetilde{\text{BN}}$ auction)

$$t_{LL} = \frac{V_H}{2} - \frac{\tilde{p}^2(V_H-V_L)}{2[2\tilde{p}-1]} \qquad t_{HL} = V_H$$
$$t_{LH} = \frac{\tilde{p}(1-\tilde{p})(V_H-V_L)}{2[2\tilde{p}-1]} \qquad t_{HH} = \frac{V_H}{2}.$$

If $C_1 \leq 0$ and $C_2 > 0$ ($\widetilde{\text{DS1}}$ auction)

$$t_{LL} = \frac{[\tilde{p}(\bar{p}+e)]V_L - [(1-\bar{p}-e)(1-\tilde{p})]V_H}{2[\bar{p}+e+\tilde{p}-1]} \qquad t_{HL} = t_{LL}+\frac{V_H}{2}$$
$$t_{LH} = \frac{\tilde{p}(1-\bar{p}-e)(V_H-V_L)}{2[\bar{p}+e+\tilde{p}-1]} \qquad t_{HH} = t_{LH}+\frac{V_H}{2}.$$

If $C_2 \leq 0$ ($\widetilde{\text{DS2}}$ auction)

$$t_{LL} = \frac{V_L}{2} \qquad t_{HL} = t_{LL}+\frac{V_H}{2}$$
$$t_{LH} = 0 \qquad t_{HH} = t_{LH}+\frac{V_H}{2}.$$

As expected, the above reduces to the solution under (P) for $\tilde{p} = \bar{p} - e$. In fact, this boundary choice of $\tilde{p}$ is optimal for the seller in each of the three cases listed above. To see this, consider the first case. Substituting $\widetilde{\text{BN}}$ into the objective function of $(\tilde{P})$ and taking the derivative of this payoff with respect to $\tilde{p}$ yields

$$\frac{[\bar{p}^2+e^2-2e[1-\bar{p}-6\tilde{p}(1-\tilde{p})]-2\bar{p}[1+2\tilde{p}(1-\tilde{p})]+\tilde{p}[4(1-\tilde{p}^2)-\tilde{p}]][V_H-V_L]}{8e[2\tilde{p}-1]^2}.$$

Tedious algebra then verifies that the term in the numerator before $[V_H - V_L]$ is negative for $\tilde{p} \in (\bar{p}-e, \bar{p}+e]$ and valid parameter values. Because the other terms are all positive, the overall expression is negative. Thus, the sellers' payoff is decreasing in $\tilde{p}$, so the principal optimally sets $\tilde{p} = \bar{p} - e$.

Next consider the second case. Substituting $\widetilde{\text{DS1}}$ into the objective function of $(\tilde{P})$ and taking the derivative of this payoff with respect to $\tilde{p}$ yields

$$\frac{[-\bar{p}(1-\bar{p})-4e^3-e^2(8\bar{p}-5)-2\bar{p}\tilde{p}+(2-\tilde{p})\tilde{p}-e(1-6\bar{p}+4\bar{p}^2+2\tilde{p})][V_H-V_L]}{4e[\bar{p}+e+\tilde{p}-1]^2}.$$

Again, the first term in the numerator is negative for $\tilde{p} \in (\bar{p}-e, \bar{p}+e]$ and valid parameter values, and all other terms are positive. Thus, the sellers' payoff is again decreasing in $\tilde{p}$, so the principal optimally sets $\tilde{p} = \bar{p} - e$ in this case as well.

Finally, consider the last case. Substituting $\widetilde{\text{DS2}}$ into the objective function of $(\tilde{P})$ and taking the derivative of the function with respect to $\tilde{p}$ yields

$$-\frac{V_H - V_L}{4e}.$$

Clearly, the above term is negative, so the principal sets $\tilde{p} = \bar{p} - e$. Thus, the auction in Proposition 2 that has V_L-bidders reporting V_L and V_H-bidders reporting V_H for all p is revenue maximizing when p is uniformly distributed. □

Proof of Corollary 2. The proof follows directly from the characterization of transfers provided in Proposition 2. □

Proof of Corollary 3. Define $\bar{p} = (1 + p')/2$ and $e = (1 - p')/2$. Because $\bar{p} > 1/2$, these choices are permissible in our base model. With these choices, the lower bound on p, $\bar{p} - e$, is simply p', whereas the upper bound on p, $\bar{p} + e$, is just 1. This is precisely the setup in Corollary 3, so substituting $\bar{p} = (1+p')/2$ and $e = (1-p')/2$ in Proposition 2 yields the result.

In particular, notice that with the above choice of $\bar{p}$ and e, (i) the transfers in DS1 reduce to those in the correlation-free DS2, and (ii) the condition $e \leq (2\bar{p} - 1)/4$ is equivalent to the condition $p' \geq 2/3$. Hence, the principal offers the BN auction if $p' \geq 2/3$ and the DS2 auction otherwise. □

References

Bergemann, D., S. Morris. 2005. Robust mechanism design. *Econometrica* **73** 1521–1534.

Bergemann, D., S. Morris. 2008. Ex post implementation. *Games Econom. Behav.* **63** 527–566.

Chung, K., J. Ely. 2007. Foundations of dominant strategy mechanisms. *Rev. Econom. Stud.* **74** 447–476.

Cremer, J., R. P. McLean. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* **56** 1247–1257.

Demski, J., R. Dye. 1999. Risk, return, and moral hazard. *J. Accounting Res.* **37** 27–55.

Demski, J., D. Sappington. 1984. Optimal incentive contracts with multiple agents. *J. Econom. Theory* **33** 152–171.

Dutta, S., S. Reichelstein. 2002. Controlling investment decisions: Depreciation and capital charges. *Rev. Accounting Stud.* **7** 253–281.

Evans, R. 2008. Simple efficient contracts in complex environments. *Econometrica* **76** 459–491.

Glover, J., Y. Ijiri, C. Levine, P. Liang. 2006. Usefulness lost: Aggregating information with differing levels of verifiability. Working paper, Carnegie Mellon University, Pittsburgh.

Laffont, J.-J., J. Tirole. 1993. *A Theory of Incentives in Procurement and Regulation.* MIT Press, Cambridge, MA.

Lopomo, G. 1998. The English auction is optimal among simple sequential auctions. *J. Econom. Theory* **82** 144–166.

Ma, A., J. Moore, S. Turnbull. 1988. Stopping agents from "cheating." *J. Econom. Theory* **46** 355–372.

Maskin, E., J. Riley. 2000. Asymmetric auctions. *Rev. Econom. Stud.* **67** 413–438.

McLean, R. P., A. Postlewaite. 2001. Efficient auction mechanisms with interdependent valuations and multidimensional signals. Working paper, University of Pennsylvania, Philadelphia.

Mookherjee, D., S. Reichelstein. 1992. Dominant strategy implementation of Bayesian incentive compatible allocation rules. *J. Econom. Theory* **56** 378–399.

Neeman, Z. 2003. The effectiveness of English auctions. *Games Econom. Behav.* **43** 214–238.

Parreiras, S. O. 2005. Correlated information, mechanism design and informational rents. *J. Econom. Theory* **123** 210–217.

Reichelstein, S. 1997. Investment decisions and managerial performance evaluation. *Rev. Accounting Stud.* **2** 157–180.

Wilson, R. 1987. Game-theoretic analysis of trading processes. T. Bewley, ed. *Advances in Economic Theory: Fifth World Congress.* Cambridge University Press, Cambridge, UK, 33–70.