Unobserved Mechanisms

Li, Hao and Michael Peters[*]

Vancouver School of Economics

University of British Columbia

May 30, 2018

Abstract

Using a well known environment where buyers have identically and independently distributed private values for an object, we model mechanism design without full observability as a game of imperfect information in which some buyers may not be informed of the commitments made by the seller of the object. Informed buyers can pretend to be uninformed but uninformed cannot pretend to be informed. In equilibrium, the seller holds an auction among informed buyers with a reserve price that depends on how many informed buyers there are. This reserve price is kept secret from informed buyers when they report their values.

[*]Li: hao.li@ubc.ca. Peters: peters@econ.ubc.ca. We are grateful to Social Sciences and Humanities Research Council of Canada for financial support.

1 Introduction

An important lesson of mechanism design is the "power of commitment." An optimal auction, for example, requires the seller to set a reserve price. If no buyer chooses to bid that price, the seller is supposed to commit himself never to trade. Ex post, the seller will not want to carry out that commitment, as the large literature on auctions with resale explains, but the seller will benefit ex ante by convincing buyers that the commitment will be honored.

Commitment is something that is relatively easy to accomplish in digital markets. Price offers are generated by computer programs that cannot be easily re-written. Yet as pricing mechanisms begin to appear in digital markets, it is apparent that there is another important assumption required by mechanism design - one that is usually not discussed at all. Even if a seller can commit, this might be of limited benefit if buyers do not understand what the seller has committed to.

The fact that buyers do not notice sellers' commitments has been widely documented in the marketing literature. Dickson and Sawyer (1990), for example, asked buyers in supermarkets about their price knowledge as they were shopping. Only 50% of all respondents claimed to know the price of the object they had just put in their basket. Even when the item being placed in the basket had been specially marked down and heavily advertised, 25% of consumer did not even realize the good was on special.

Supermarket prices are staring consumers in the face as they look at the shelves. It would seem much harder for consumers to be aware of selling techniques on the internet. For example, many websites use *click stream pricing*, whereby the price offer that is made to a customer can be made to depend on exactly what the buyer does before getting to the offer page (i.e., the stream of clicks that leads them to the website). A higher price may be offered to a consumer who searches directly for a deluxe model of a product than to a consumer who searches for the standard model then asks how much it would cost to upgrade. Another example is airline ticket pricing. This propensity for being uninformed may be one reason why airlines, for example, make so little effort to explain the details of their pricing algorithms to buyers.

The point of this paper is to consider the implications of the possibility of buyers not knowing the selling mechanism for the behavior and payoffs of informed buyers as well as the seller. To this end, we choose the best known mechanism design environment of all - independent private value auctions (Myerson, 1981).[1] Our seller will realize that some buyers are uninformed, but will not be sure how many uninformed buyers there are, or who they are. Our buyers will be in a similar position. Those who are uninformed will understand that they do not understand the seller's mechanism. Informed buyers will know the mechanism and understand that there are uninformed buyers, but again, will not know the identities of the uninformed buyers.

We emphasize that our uninformed buyers are not behavioral - they have rational expectations in the equilibria we construct. They can be thought of as randomly attentive who understand their own inattention (Masatloglu, 2015; Masatloglu, Nakajima, and Ozbay, 2012). Formally, we treat the mechanism design process as a game of imperfect information in which uninformed buyers do not see part of the history - the part where the seller commits to a mechanism.

We first show that standard auction formats simply are not robust to the possibility that there are uninformed buyers. The key to this result is that when a seller "deviates" and changes the mechanism, informed buyers will understand the change and respond to it in the usual way, but uninformed buyers will not respond at all.

What this does is to give the seller the opportunity to extract surplus from the uninformed without losing any surplus from the informed. To see this consider a second price auction. Suppose the uninformed believe that the seller is holding a second price auction. In the usual way, a best response is for them to bid their values. The informed buyers know whether or not the seller is using a second price auction, so, of course, they also bid their values. The seller can then deviate, explaining to the (informed) buyers that if they want to bid in a second price auction, then they have to provide a certain password along with their bid.[2] If the high bidder in the auction has given the password, the bidder pays the second highest bid; if the high bidder does not give the password, the bidder pays the winning bid.

[1]For a full list of references, see, for example, Krishna (2010).

[2]The term "password" is just used by way of analogy. In practice, this is where the "click stream pricing" method, and other things like it would be used. Burying lower prices beneath a complex web of click streams is a way to provide a password since the informed buyers will discover the correct stream, while the uninformed will not.

The password analogy is one we will use again below. Of course, we do not necessarily mean this to be taken literally. In the click stream pricing example, realizing that there is an alternative path to a price quote is the same thing as knowing a password. Meet the competition (match lower prices elsewhere) requires buyers to provide information uninformed buyers do not have, again similar to a password. Coupons that offer price discounts must often be presented in some form at the time of a transactions, again in a manner similar to a password.

In a standard setting with full observability, an optimal mechanism will want to specialize the outcome for any aspect of a buyer's type. In our context this means that the seller in equilibrium will want to treat informed and uninformed buyers of the same value differently. This gives rise to a somewhat unusual incentive problem in the sense that informed buyers can "verify" they are informed by providing a password before the transaction. They do not have to do this unless it is in their interest. Uniformed buyers, on the other hand cannot pretend to be informed. This fact is the core argument in much of what follows below.

We show that equilibrium selling mechanisms discriminate against the uninformed in a special way. Uninformed buyers in our model always receive a take-it-or-leave-it offer. Since uniformed buyers do not understand how a seller is allocating a good, they cannot see deviations. In any outcome in which the uninformed believe that their price offer will depend on their value, the seller will be able to deviate and extract surplus from them. Indeed one implication of this is that in markets in which it is very unlikely that buyers are informed, the selling mechanism will appear to be a simple take-it-or-leave-it price with high probability.

Our model suggests that reserve prices will vary with the actual number of uninformed buyers. The reason is that the uninformed represent an outside option for the seller when holding an auction with the informed. For example, if there are no uninformed buyers, the seller wants to set the reserve price equal to the value of a buyer whose virtual value is zero, as is standard. However, if there is an uninformed buyer, the seller wants a higher reserve price such that the virtual value is equal to the seller's outside option of selling to the uninformed. Since there is a gap between these two reserve prices, the seller has to keep the actual reserve secret to prevent informed buyers whose values are in the gap from pretending to be uninformed. The fact that sellers will want to keep some aspects of their mechanism hidden from informed buyers is one of the main implications of the possibility that there are uninformed buyers.[3] This result also provides a simple explanation for why reserve prices are so often kept private in practice.

We show that uninformed buyers always reduce the seller's revenue relative to what they would be with full observability. This is because those who do not observe the seller's commitments nonetheless have rational expectations regarding the seller's mechanism. This means that the uninformed buyers create a constraint on the set of mechanisms that are available to the seller.

We also show that uninformed buyers are offered higher prices than they would be if there were no commitment at all, because the seller needs to discourage informed buyers from pretending that they are uninformed. This means that uninformed buyers with low values are worse off due to the presence of informed buyers. Those with high values, however, can be better off, because of reduced competition as informed buyers may not meet the reserve prices in the auctions. Indeed, at least some of the reserve prices are set higher than in standard auctions for the seller to exploit the option of making take-it-or-leave-it offers to uninformed buyers.

In contrast to uninformed buyers, the informed with low values are better off relative to what they get from the standard auction under full observability. This is because the reserve price for the auction when every buyer turns out to be informed is set lower than in the standard auction to prevent each individual informed buyer from pretending to be uninformed. As explained above, at least some of the other reserve prices are higher than in the standard auction, so informed buyers with high values can be worse off. However, the presence of uninformed buyers also means that informed buyers face reduced competition for the good as the former do not participate in any of the auctions. We show that when the total number of potential buyers becomes sufficiently large, the effect of reduced competition outweighs the effect of higher prices, so informed buyers with high enough values are better off relative to full observability.

[3] This result applies even if the seller uses a second price auction, where in the standard setting buyers' behavior would be independent of the reserve price.

1.1 Literature

As mentioned above, the idea that consumers might not notice prices in an old one in the marketing literature, as in Dickson and Sawyer (1990) and references there in. The approach had been used earlier in economics, as in, say Butters (1977), in which buyers randomly observe price offers in a competitive environment. In that literature, firms advertise prices which some buyers see, while others do not.[4] These papers considered the same problem that we do, which is how this unobservability would affect the prices that firms offer. The difference here is that we are interested in mechanisms, not prices.

Peters (2014) also uses the idea of unobservatility in a competitive setting in which the mechanism design problem is much simpler. The problem we consider here is not competitive. We are interested in the more basic issue of how unobservability affects mechanism design. As for mechanism design, our basic problem involves a standard mechanism design problem in which buyers have outside options that vary with their type. Such a problem was studied by McAfee (1993). His model had buyers whose outside option involved waiting until next period to purchase. Here, though the outside option does vary with type, it is endogenous because the seller can change it by modifying the offer made to uninformed buyers.

With buyers potentially uninformed of the selling mechanism but nonetheless having rational expectations, the seller's commitment power is limited. The limitation on the seller's commitment power in our unobserved mechanism problem is therefore endogenous, which differentiates our model from the existing literature on mechanism design problems under limited commitment (Bester and Strausz, 2001; Kolotilin, Li and Li, 2013; Liu, Mierendorff and Shi, 2014; Skreta, 2015). A recent paper by Akbarpour and Li (2018) provides another model of limited commitment. They assume that each individual buyer only observes the part of the seller's commitment in relation to the buyer's own report, and impose a "credibility" constraint that the seller does not wish to secretly alter other parts of the commitment. Insofar as the set of feasible deviations available to the seller with respect to each buyer depend on the extensive form of the mechanism, their approach to limited commitment remains exogenous.

[4]See also Varian (1980), or Stahl (1989, 1994). Varian (1980) calls buyers informed if they see prices of all firms, and uninformed if they do not.

A recent literature has studied optimal mechanism design with exogenous communication constraints (Deneckere and S. Severinov, 2008; Mookherjee and Tsumagari, 2014). In our problem, uninformed buyers have the correct belief about the equilibrium mechanism, and as a result, they are uncommunicative. We show that the equilibrium mechanism corresponds to the solution to an optimal mechanism design problem in which informed buyers voluntarily reveal their values while uninformed buyers remain silent.

2 The Model

There are n potential buyers of a single homogenous good. Each buyer has a privately known value that is independently drawn from the interval $[0, 1]$. Assume for the moment, all values are distributed according to some distribution F with strictly positive density f. Buyers' payoff when they buy at price p is given by $v - p$. The seller's cost is zero, so the profit from selling at price p is just p.

There is a *common* message space $\mathcal{M}$ which is used by all buyers to communicate with the seller. For example, $\mathcal{M}$ might be the set of possible browsing histories for a buyer. A message from buyer i will be denoted b_i. We make no assumptions on $\mathcal{M}$ itself except that it is rich enough to embed the set of buyer values. and that the set of feasible messages is common knowledge.

The point of the seller's mechanism is to generate an *offer*.[5] It is important to note that even buyers who are uninformed can understand a take-it-or-leave-it offer. This is the one commitment that is understood by every buyer. To avoid complexities that have little to do with unobservability, as least as a first attempt at unobservability, we will assume that the seller's mechanism determines the identity of the buyer who receives the offer, and what this offer is. We assume that if this offer is refused, the game ends without trade. It would be possible, of course, to consider more complex interactions when an offer is rejected, but for the sake of simplicity we will not consider such alternatives here.

[5] In the standard mechanism design paradigm, a mechanism produces an allocation according to a mapping from messages. Representing the output of a mechanism as an offer instead of an allocation has no implications. This is no longer the case in our unobserved mechanism problem. See section 5 for additional comments on modeling the output of an unobserved mechanism as an algorithm.

What differentiates our paper from the standard mechanism design literature is that some buyers do not know what the seller's mechanism is. We model this by a game of *imperfect* information in which some buyers see the first move by the mechanism designer while others do not. We assume that the probability that a buyer does not see the commitment is α. Whether or not a buyer can see the mechanism is private information, and is independent of each other. Our solution concept is just perfect Bayesian equilibrium, so buyers who do not observe the mechanism will none the less correctly guess what it is in equilibrium.

With this preamble, we can describe the seller's mechanism more formally. A mechanism γ for the seller is a collection $\{\mathcal{M}, p_i, q_i\}_{i=1}^N$, where $\mathcal{M}$ is the common message space, p_i is a mapping from a profile of messages $(b_1, \ldots, b_N)$ to a take-it-or-leave-it offer to buyer i, and q_i is a mapping from $(b_1, \ldots, b_N)$ to a probability of the offer to each buyer i. Let Γ be the set of mechanisms.

Buyer i's type is a pair (v, τ), where v is the value for the object, and τ is a binary variable that is equal to ϵ when the buyer is informed, and μ when the buyer is uninformed. A strategy rule for a buyer is a function $\sigma_i : [0,1] \times \{\epsilon, \mu\} \times \Gamma \to \mathcal{M}$ that specifies what message the buyer will send for each of the values conditional on whatever the buyer knows about the seller's mechanism. Since an uniformed buyer never sees the mechanism a seller offers, we have the informational constraint

$$\sigma_i (v_i, \mu, \gamma) = \sigma_i (v_i, \mu, \gamma') = \sigma_i (v_i, \mu)$$

for all γ and γ'. Notice that this definition restricts buyers to pure strategies. We retain this assumption throughout the paper.

As mentioned above, informed buyers can pretend to be uninformed, but not conversely. This allows the seller to identify informed buyers by, for example, asking for a password along with a bid. To prevent the uninformed buyers from guessing this password, it has to be random. One way to understand this is to imagine that the seller randomizes over mechanisms in order to make this password unpredictable. Refer to this mixture as $\psi \in \triangle(\Gamma)$. Let $R(\gamma, (\sigma_i(\cdot, \cdot, \gamma))_{i=1}^n)$ be the expected revenue for the seller from mechanism γ when an uninformed buyer i uses strategy $\sigma_i(\cdot, \mu)$ an informed buyer i uses $\sigma_i(\cdot, \epsilon, \gamma)$.

A perfect Bayesian equilibrium for this problem is a mixture ψ for the seller, and pairs of strategy rules $\left(\sigma_i\left(\cdot,\epsilon,\gamma\right),\sigma_i\left(\cdot,\mu\right)\right)_{i=1}^N$ for informed and uninformed buyers respectively, which satisfy

$$\begin{aligned}&\mathbb{E}_{v_{-i},\tau_{-i}}\left[q_i\left(\sigma_i\left(v_i,\epsilon,\gamma\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot\max\left\{v_i-p_i\left(\sigma_i\left(v,\epsilon,\gamma\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right),\epsilon\right\}\right]\\ \geq&\mathbb{E}_{v_{-i},\tau_{-i}}\left[q_i\left(b_i,\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot\max\left\{v_i-p_i\left(b_i,\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right),0\right\}\right]\end{aligned}\tag{1}$$

for all $v_i\in[0,1]$, $b_i\in\mathcal{M}$ and realized γ from the mixture ψ;

$$\begin{aligned}&\mathbb{E}_{v_{-i},\tau_{-i},\gamma}\left[q_i\left(\sigma_i\left(v_i,\mu\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot\max\left\{v_i-p_i\left(\sigma_i\left(v_i,\mu\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right),\epsilon\right\}|\psi\right]\\ \geq&\mathbb{E}_{v_{-i},\tau_{-i},\gamma}\left[q_i\left(b_i,\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot\max\left\{v_i-p_i\left(b_i,\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right),0\right\}|\psi\right]\end{aligned}\tag{2}$$

for all $v_i\in[0,1]$, and $b_i\in\mathcal{M}$; and

$$\mathbb{E}_\gamma\left[R\left(\gamma,\left(\sigma_i\left(\cdot,\cdot,\gamma\right)\right)_{i=1}^n\right)|\psi\right]\geq R\left(\gamma',\left(\sigma_i\left(\cdot,\cdot,\gamma'\right)\right)_{i=1}^n\right)\tag{3}$$

for all $\gamma'\in\Gamma$.

In the expressions (1) and (2), the terms $\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)$ should be understood to mean that whenever one of the buyers other than i is uninformed, they use the strategy $\sigma_{-i}\left(v_{-i},\mu\right)$. The reason that the max operation appears when taking expectations is because a mechanism generates an offer instead of an outcome. It will be assumed throughout that this offer is take-it-or-leave-it. As mentioned above, if it is refused there is no trade at all.

2.1 Relationship to standard mechanism design

Once the outcome of the seller's randomization is realized, the equilibrium of the continuation game is a simple Bayesian equilibrium. So we could try to analyze it using direct mechanisms in which buyers' types consist of both their value and whether or not they are informed. However, there are some important issues to bear in mind.

The first is that information about whether or not a buyer is informed is quite different from the information about the value because the informed buyer can "verify" to the seller that knowledge but the uninformed buyer cannot.[6] Being informed in this problem is similar to being able to speak French, or being able to play the piano. We explained above how the seller could give the buyer the opportunity to reveal this part of the type. The seller randomizes over mechanisms, and includes in each mechanism a password - for example a real number selected in the interval $[0, 1]$ using a uniform distribution. The seller would publish this password as part of the mechanism. The informed buyers see it, the uninformed do not.

The second issue comes from the fact that a mechanism is just a move in an imperfect information game. Since uninformed buyers cannot see deviations by the seller, uninformed buyers with different values will receive the same offer in our equilibrium. Of course, it is possible to interpret a mechanism like this as a direct mechanism in which the seller commits to an outcome that is independent of buyer's value. This does not quite work here since if uninformed buyers do report their values, the seller can no longer commit not to use the information contained in those reports. As a consequence we need to retain, at least initially, much of the structure of the indirect mechanism.

Third, the outcomes that are produced by the seller in response to messages are not like outcomes in a standard auction. Here, an outcome is just an offer. It is possible that buyers in equilibrium will receive offers that they will not want to accept. As will be discussed, the standard reduced form functions have to be reinterpreted in order to make them work.

2.2 Example

As an example to illustrate how the imperfect information part of the story affects the analysis, suppose that the seller is using a second price auction in which bids are submitted in US dollars. The high bidder wins and pays the second highest US dollar bid. We would argue that this can never be an equilibrium mechanism.

To see why, suppose the seller deviates to a new mechanism in which bids can be either in Canadian or US dollars. The buyer who submits the highest bid (exchange rate adjusted) wins the auction. If the winning bid was submitted in Canadian dollars, the winner pays the second highest bid whether it is in US dollars or in Canadian dollars. If, on the other

[6]Verifiable information is discussed in Green and Laffont (1986) or Strausz (2016) among others.

hand, the winning bid is submitted in US dollars, the winner pays the bid.

The logic is that only informed buyers realize they can submit bids in Canadian dollars. If they choose to do so, they know that their bid will be treated in the usual second price way, so they will prefer to bid their value expressed in Canadian dollars. Uninformed buyers will not observe the deviation and erroneously believe that their US dollar bids will be treated as values in a second price auction. When they win they will be offered a trade at the value that they bid. This allows the deviating seller to extract all the surplus of the uninformed buyer when they win the auction. Moreover, the seller does so without sacrificing the revenue from informed buyers. Informed buyers do not have to bid in Canadian dollars, but they observe the deviation by the seller and understand that any bid they submit in US dollars will result in a strictly higher payment than the corresponding bid in Canadian dollars. So bidding in Canadian dollars is a dominant strategy for any informed buyer who wants to trade.

What we can learn from this example is that all equilibrium mechanisms will be discriminatory. Standard auction forms as they are usually studied can never be equilibrium mechanisms when there are uninformed buyers.

3 Unobserved Mechanism Design

For the remainder of the paper, we will focus on mechanisms that are symmetric with respect to buyer identity and assume that continuation equilibria are fully symmetric.[7] Given a mechanism γ, a profile of types $\{(v_1, \tau_1), \ldots, (v_n, \tau_n)\}$ then results in trader i receiving an offer with probability $q_i(\sigma(v_i, \tau_i, \gamma), \sigma(v_{-i}, \tau_{-i}, \gamma))$. The offer made in this case will be $p_i(\sigma(v_i, \tau_i, \gamma), \sigma(v_{-i}, \tau_{-i}, \gamma))$.

We assume that $q_i(\sigma(v_i, \tau_i, \gamma), \sigma(v_{-i}, \tau_{-i}, \gamma)) > 0$ if and only if there is a v_i' such that

$$\sigma(v_i', \tau_i, \gamma) = \sigma(v_i, \tau_i, \gamma)$$

and

$$p_i(\sigma(v_i', \tau, \gamma), \sigma(v_{-i}, \tau_{-i}, \gamma)) \leq v_i'.$$

[7] We drop the subscript from the strategy rule σ_i, but keep the subscripts in v_i and τ_i, and in the mapping p_i and q_i, so we can use p and q to represent the outcome profiles.

Any outcome function that does not satisfy this property is payoff equivalent to one that does. We think of $q_i\left(\sigma\left(v_i,\tau_i,\gamma\right),\sigma\left(v_{-i},\tau_{-i},\gamma\right)\right)$ as the probability with which i receives a "serious" offer. Similarly, using payoff equivalence, whenever

$$q_i\left(\sigma\left(v_i,\tau_i,\gamma\right),\sigma\left(v_{-i},\tau_{-i},\gamma\right)\right)=0,$$

we can assume

$$p_i\left(\sigma\left(v_i,,\tau_i,\gamma\right),\sigma\left(v_{-i},\tau_{-i},\gamma\right)\right)=1.$$

3.1 Uninformed buyers

Given an equilibrium $(\psi,\sigma(v,\mu),\sigma(v,\epsilon,\gamma))$, we can now define a reduced form probability as follows

$$Q\left(v_i,\tau_i\right)\equiv\mathbb{E}_{v_{-i},\tau_{-i},\gamma}\left[q_i\left(\sigma\left(v_i,\tau_i,\gamma\right),\sigma\left(v_{-i},\tau_{-i},\gamma\right)\right)|\psi\right].$$

For all v_i and τ_i such that $Q\left(v_i,\tau_i\right)>0$, define

$$\mathcal{P}\left(\tilde{p},v_i,\tau_i\right)\equiv\Pr\left\{p_i\left(\sigma\left(v_i,\tau,\gamma\right),\sigma\left(v_{-i},\tau_{-i},\gamma\right)\right)\leq\tilde{p}|\psi\right\}.$$

Finally,

$$P\left(v_i,\tau_i\right)=\begin{cases}\int\tilde{p}\ d\mathcal{P}\left(\tilde{p},v_i,\tau_i\right) & \text{if } Q\left(v_i,\tau_i\right)>0\\ \inf_{v_i':Q\left(v_i',\tau_i\right)>0}\int\tilde{p}\ d\mathcal{P}\left(\tilde{p},v_i',\tau_i\right) & \text{otherwise.}\end{cases}\tag{4}$$

By definition, $P\left(v_i,\tau_i\right)$ is either equal or close to some expected offer that is made with positive probability.

As in standard mechanism design, the payoff to a buyer of value v_i in the continuation equilibrium is

$$U\left(v_i,\tau_i\right)\equiv Q\left(v_i,\tau_i\right)\cdot\int\max\left\{v_i-\tilde{p},0\right\}\ d\mathcal{P}\left(\tilde{p},v_i,\tau_i\right).$$

By standard arguments, the function $U\left(v_i,\tau_i\right)$ is continuous in v_i.

The following result helps simplify the analysis.

Theorem 1. *In any symmetric equilibrium $P\left(v_i,\mu\right)$ is independent of v_i.*

Proof. Suppose the theorem is false and there is a pair v_i' and v_i such that $P(v_i', \mu) > P(v_i, \mu)$. We can assume $Q(v_i', \mu)$ and $Q(v_i, \mu)$ are both strictly positive. This is because the definition (4) forces $P(v_i, \mu)$ for an uninformed buyer with values who does not receive offers to be either the same as, or arbitrarily close to the expected price for one with values who does receive offers.

By construction, all offers in the support of $\mathcal{P}(\cdot, v_i', \mu)$ are accepted by an uninformed buyer with value v_i'' for whom $\sigma(v_i'', \mu) = \sigma(v_i', \mu)$, so we can choose v_i'' such that $U(v_i'', \mu) > 0$. Let

$$\beta(v_i') \equiv \inf\{v_i'' : \sigma(v_i'', \mu) = \sigma(v_i', \mu) \text{ and } U(v_i'', \mu) > 0\}.$$

We claim that $\tilde{p} \geq \beta(v_i')$ for almost all $\tilde{p}$ in the support of $\mathcal{P}(\cdot, v_i', \mu)$. If this were false, the seller could raise revenue by increasing his price offer on a set of positive measure. Since uninformed buyers would not observe this change in the price offer, their reporting behavior would not be affected. Uninformed buyers whose values are at least $\beta(v_i')$ would continue to accept these offers.

The implication of the above claim is that buyers whose values are close to $\beta(v_i')$ will have expected payoffs arbitrarily close to zero. Since by assumption $P(v_i, \mu) < P(v_i', \mu)$,

$$\int_0^{\beta(v_i')} d\mathcal{P}(\tilde{p}, v_i, \mu) > 0$$

which violates incentive compatibility since buyers whose values are close enough to $\beta(v_i')$ could profitably deviate by adopting the strategy $\sigma(v_i, \mu)$. ■

The proof of Theorem 1 illustrates why uninformed buyers are bad for sellers. The seller's best revenue occurs under the full commitment power. Uninformed buyers cannot see the commitment made by the seller directly. This creates a situation in which the seller can no longer commit not to use information revealed by the uninformed during communication. Even though sellers technically have full commitment power, uninformed buyers take that commitment power away from them.

Say that a mechanism is *non-discriminatory* if

$$(Q(v_i, \epsilon), P(v_i, \epsilon)) = (Q(v_i, \mu), P(v_i, \mu))$$

for all v_i. Theorem 1 implies that informed buyers must also receive a single take-it-or-leave-it offer if the mechanism is non-discriminatory. This cannot be an equilibrium since the seller can deviate and offer the informed a second price auction without affecting the uninformed and without violating incentive compatibility. A non-discriminatory mechanism can be supported in equilibrium only in the trivial case where there are no informed buyers.

Corollary 1. *If all buyers are uninformed for sure, then there is a unique equilibrium in which the seller randomly selects a buyer to make an offer that maximizes* $(1 - F(p))\, p$.

Corollary 1 illustrates our endogenous approach to limited commitment in mechanism. While full observability, with $\alpha = 0$, corresponds to the standard mechanism design problem under full commitment, full non-observability, with $\alpha = 1$, is different from no commitment. Indeed, if $n = 1$, the literature on the Coase Conjecture initiated by Coase (1972) establishes that without any commitment power the seller can do no better than offering the good at the price of 0 (Fudenberg, Levine, and Tirole, 1985; Gul, Sonnenschein, and Wilson, 1986),[8] while the equilibrium outcome given in Corollary 1 corresponds to what the seller could achieve with full commitment (Riley and Zeckhauser, 1983).[9]

3.2 Informed buyers

In what follows, we make the usual assumption that virtual value is increasing.

Assumption 1. *The virtual value, given by the function*

$$\phi(v) = v - \frac{1 - F(v)}{f(v)},$$

is strictly increasing in v.

[8] This requires stationarity of the equilibria. We have what is known in the literature as the no-gap case because the lowest value equals the seller's reservation value. For non-stationary equilibria, see Ausubel and Deneckere (1989).

[9] The conclusion that there is an equilibrium outcome with $n = 1$ and $\alpha = 1$ that is the same as the full commitment outcome does not rely on the present restriction to mechanisms with single offers. Without this restriction, it is an equilibrium in which the seller commits to a single take-it-or-leave-it offer that maximizes $(1 - F(p))\, p$, and the uninformed buyer accepts the offer if and only if the value is above the offer, with the belief that any subsequent offer will lead to an offer of 0 after a rejection. See section 5 for additional comments on unobserved mechanisms as producing algorithms.

Assumption 1 implies that there is a single value r^* where $\phi(r^*) = 0$. This value r^* represents the optimal reserve price in a standard auction, regardless of the number of buyers. That is, when $\alpha = 0$, the seller's outside option is always 0, so the reserve price is such that the virtual value of the buyer with v at the reserve price is equal to the seller's outside option. Since this optimal reserve price is independent of the number of buyers, r^* is also the optimal take-it-or-leave-it offer when the seller is dealing with a single buyer.

At the other extreme, when $\alpha = 1$, by Corollary 1 there is no informative communication from the n buyers, and the seller randomly chooses one of them to make a take-it-or-leave-it offer. For convenience, define

$$\pi(p) = (1 - F(p))p$$

as the revenue function from a take-it-or-leave-it offer p. Since the seller's outside option is 0, the optimal offer is again r^*. By Assumption 1, $\pi(p)$ increases for $p < r^*$ and decreases for $p > r^*$.

When $\alpha \in (0, 1)$, symmetric equilibrium outcome functions turn out to depend on the number of informed and uninformed buyers, but not on their identities. Payments by informed buyers will be determined below using standard incentive conditions, together with the condition that they do not want to pretend to be uninformed. If the seller decides to make an offer to an uninformed buyer the seller must decide what that offer will be. By Theorem 1, this offer cannot depend on the uninformed buyer's value, but again it can depend on how many buyers are informed. Moreover, since the informed buyers can pretend to be uninformed, the take-it-or-leave-it offer to the uninformed will impact the incentive of the informed buyers. To characterize symmetric equilibrium outcomes, we make the following simplification.

Definition 1. *An equilibrium $(\psi, \sigma(\cdot, \mu), \sigma(\cdot, \epsilon, \gamma))$ involves pooling of uninformed buyers if for all v_i, v_i', we have $q_i(\sigma(v_i, \mu), \sigma(v_{-i}, \tau_{-i}, \gamma)) = q_i(\sigma(v_i', \mu), \sigma(v_{-i}, \tau_{-i}, \gamma))$ for all v_{-i} and τ_{-i}, and for all γ in the support of ψ.*

We will focus on equilibria that pool uninformed buyers in what follows. This amounts to an equilibrium refinement. When uninformed buyers are indifferent among a number of different messages only because all messages they send yield the same expected payoff, they could send different messages. In particular, uninformed buyers who get the same expected payoff of 0 because they have low values can send a message to the seller indicating that they are "not interested." This is ruled out by Definition 1.

Combined with Theorem 1, the restriction to pooling of uninformed buyers allows us to focus on equilibria in which uninformed buyers simply "keep silent." It then follows from a straightforward revelation-principle argument that without loss we can impose the standard incentive compatibility constraints for informed buyers to truthfully report their values, together with the new incentive condition in the present setup that informed buyers do not wish to keep silent. This allows us to present the characterization of equilibrium outcomes in terms of a constrained optimization problem (Theorem 2 below).[10] We stress that we are characterizing the outcomes, not the mechanisms; as we have explained previously, uninformed buyers have rational expectations about the equilibrium mechanism, and they are prevented from participating in it by the seller's use of randomization (i.e., password).

We can now drop σ from the outcome functions p and q, and for informed buyers, we also drop the dependence of their strategy on the realized mechanism γ. Let the realized number of uninformed buyers be m. For each $m = 0, \ldots, n-1$, denote the $(n-m)$-dimensional profile of reported values as $v(m)$. For each informed buyer i, we also write the profile as $v(m) = (v_i, v_{-i}(m))$. The seller's mechanism is

$$\{(q(v(m)), p(v(m)))_{m=0}^{n-1}, (t_m)_{m=1}^{n}\},$$

where $q(v(m))$, $m = 0, \ldots, n-1$, is the profile of probabilities that the seller assigns the good to the $n-m$ buyers who have reported the profile $v(m)$, with $p(v(m))$ the corresponding profile of price offers, and t_m, $m = 1, \ldots, n$, is the take-it-or-leave-it offer to a randomly selected buyer among those who have kept silent. For each $m = 0, \ldots, n-1$, let $b(m; n-1, \alpha)$ be the probability that there are m uninformed buyers among $n-1$ buyers, given by the binomial distribution:

$$b(m; n-1, \alpha) = \binom{n-1}{m} (1-\alpha)^{n-1-m} \alpha^m.$$

[10] For brevity, we will not mention the refinement of Definition 1 when we speak of equilibria of the game.

Then, the truth-telling payoff $U(v_i, \epsilon)$ of an informed buyer with value v_i is given by

$$\sum_{m=0}^{n-1} b(m; n-1, \alpha)\mathbb{E}_{v_{-i}(m)}\left[q_i(v_i, v_{-i}(m)) \max\left\{v_i - p_i(v_i, v_{-i}(m)), 0\right\}\right],$$

There is no loss in assuming $p_i(v_i, v_{-i}(m)) \leq v_i$, because otherwise we can always set $q_i(v_i, v_{-i}(m)) = 0$ and $p_i(v_i, v_{-i}(m)) = v_i$.[11] Then we have

$$U(v_i, \epsilon) = \sum_{m=0}^{n-1} b(m; n-1, \alpha)\left(Q_m(v_i, \epsilon)v_i - P_m(v_i, \epsilon)\right), \tag{5}$$

where

$$Q_m(v_i, \epsilon) = \mathbb{E}_{v_{-i}(m)}\left[q_i(v_i, v_{-i}(m))\right],$$

and

$$P_m(v_i, \epsilon) = \mathbb{E}_{v_{-i}(n-m)}\left[q_i(v_i, v_{-i}(m))p_i(v_i, v_{-i}(m))\right].$$

Theorem 2. *Any outcome of an equilibrium can be implemented by the seller holding a second-price auction among informed buyers with a reserve price r_m (if the realized number m of uninformed buyers is at most $n-1$) and making a take-it-or-leave-it offer t_m to a randomly selected uninformed buyer (if $m \geq 1$) when all bids are below r_m, where r_m and t_{m+1}, $m = 0, \ldots, n-1$, maximize the seller's revenue*

$$\sum_{m=0}^{n-1} b(m; n, \alpha)(n-m)\int_{r_m}^{1} \phi(v_i)F^{n-1-m}(v_i)f(v_i)dv_i + \sum_{m=1}^{n} b(m; n, \alpha)F^{n-m}(r_m)\pi(t_m), \tag{6}$$

subject to a "bidding constraint" for every v_i,

$$\begin{aligned} U(v_i, \epsilon) = &\sum_{m=0}^{n-1} b(m; n-1, \alpha)\int_{\min\{v_i, r_m\}}^{v_i} F^{n-1-m}(w)dw \\ \geq U(v_i, \mu) = &\sum_{m=0}^{n-1} \frac{b(m; n-1, \alpha)}{m+1}F^{n-1-m}(r_{m+1})\max\{v_i - t_{m+1}, 0\}. \end{aligned} \tag{7}$$

[11]We cannot do this with uninformed buyers, as the price offers do not depend on their value.

Proof. From (5), the standard incentive condition for truthful reporting by an informed buyer i with value v_i is

$$U(v_i, \epsilon) \geq \sum_{m=0}^{n-1} b(m; n-1, \alpha) \left(v_i Q_m(v_i', \epsilon) - P(v_i', \epsilon) \right),$$

for any $v_i' \in [0, 1]$. The necessary and sufficient conditions for an informed buyer with value v_i to truthfully report the value are the envelope condition

$$\frac{dU(v_i, \epsilon)}{dv_i} = \sum_{m=0}^{n-1} b(m; n-1, \alpha) Q_m(v_i, \epsilon), \tag{8}$$

and the monotonicity condition that the right-hand side of (8) above is non-decreasing in v_i.

If the informed buyer with value v_i deviates and keeps silent, he gets the same expected payoff $U(v_i, \mu)$ as an uninformed buyer with value v_i, given by

$$\sum_{m=0}^{n-1} \frac{b(m; n-1, \alpha)}{m+1} \mathbb{E}_{v(m+1)} \left[1 - \sum_j q_j(v_j, v_{-j}(m+1)) \right] \max\{v_i - t_{m+1}, 0\},$$

where $\sum_j q_j(v_j, v_{-j}(m+1))$ is the total probability that the seller makes a price offer to one of the buyers whose reports constitute the profile $v(m+1)$, with $\sum_j q_j(v(n)) = 0$ by definition. The bidding condition for an informed buyer j with value v_i not to keep silent is then

$$U(v_i, \epsilon) \geq U(v_i, \mu).$$

Since all offers to informed buyers are accepted for sure, the seller's expected total revenue from informed buyers is

$$\sum_{m=0}^{n-1} b(m; n, \alpha) \mathbb{E}_{v(m)} \left[\sum_j q_j(v(m)) p_j(v(m)) \right].$$

Since all buyers are ex ante symmetric, we can rewrite the above as

$$\sum_{m=0}^{n-1} b(m; n, \alpha)(n-m) \int_0^1 P_m(v_i, \epsilon) f(v_i) dv_i.$$

Using the identity

$$b(m; n, \alpha) = \frac{n}{n-m}(1-\alpha)b(m; n-1, \alpha)$$

for each $m = 0, \ldots, n-1$, and the payoff expression (5), we can further rewrite the revenue from informed buyers as

$$n(1-\alpha)\int_0^1 \left(\sum_{m=0}^{n-1} b(m; n-1, \alpha)Q_m(v_i, \epsilon)v_i - U(v_i, \epsilon)\right) f(v_i)dv_i,$$

which, by integration by parts and the envelope condition (8) together with $U(0, \epsilon) = 0$, becomes

$$n(1-\alpha)\int_0^1 \left(\sum_{m=0}^{n-1} b(m; n-1, \alpha)Q_m(v_i, \epsilon)\right) \phi(v_i)f(v_i)dv_i.$$

Using the binomial identity and symmetry again, we have the final formula for the revenue from informed buyers

$$\sum_{m=0}^{n-1} b(m; n, \alpha)\mathbb{E}_{v(m)}\left[\sum_j q_j(v(m))\phi(v_j)\right].$$

The seller chooses the profile of probabilities $q(v(m))$ for all $v(m)$ to maximize the sum of the above expression and the total expected revenue from uninformed buyers, given by

$$\sum_{m=1}^{n} b(m; n, \alpha)\left(1 - \mathbb{E}_{v(m)}\left[\sum_j q_j(v(m))\right]\right)\pi(t_m).$$

Aside from the bidding constraints, the only other constraint is a feasibility condition

$$\sum_j q_j(v(m)) \leq 1$$

for all $v(m)$.

The solution to the seller's problem takes the form of a reserve price when at least one buyer is uninformed, whereby the informed buyer who reports the highest value above the reserve price is assigned the good and a randomly selected uninformed buyer is selected for a take-it-or-leave-it offer when no informed buyer meets the reserve price. This is because for any other allocation rule, the seller can find a reserve price that weakly increases the revenue while keeping the bidding condition satisfied. We can therefore represent the solution by reserve prices r_m, $m = 0, \ldots, n-1$, together with take-it-or-leave-it offers t_m, $m = 1, \ldots, n$. For each $m = 0, \ldots, n-1$, we have

$$Q_m(v_i, \epsilon) = \begin{cases} F^{n-1-m}(v_i) & \text{if} \quad v_i \geq r_m, \\ 0 & \text{otherwise} \end{cases}$$

The theorem follows immediately. ■

The significance of the above theorem is to reduce the construction of the equilibrium outcome to solving a constrained optimization problem. The decision variables are $2n$ constants - the n reserve prices $\{r_0, \ldots, r_{n-1}\}$ for informed buyers, and the n take-it-or-leave-it offers $\{t_1, \ldots, t_n\}$ for the uninformed. We will refer to the solution to constrained optimization problem as the "equilibrium auction."

4 Equilibrium Auctions

The main difficulty in using Theorem 2 to characterize the equilibrium auctions is that there is a continuum of bidding constraints (7), and a priori we do not know which are binding. In standard mechanism design problems, the way to deal with a continuum of incentive constraints is to replace them with local conditions. This appears to be unsuitable for our problem here, because we naturally do not expect the bidding constraints to bind for all values of an informed buyer. Below we attempt to make some progress on this issue of bidding constraints by first studying two benchmarks.

4.1 Benchmarks

Imagine that the seller observes the realized number m of uninformed buyers, instead of having to infer it from buyers' reports. We maintain the assumption that the seller does not observe who are uninformed and who are informed. Let $\overline{r} > r^*$ be such that

$$\phi(\overline{r}) = \pi(r^*).$$

This is uniquely defined because the virtual value is strictly increasing, and $\phi(1) = 1 > \pi(r^*) > \phi(r^*) = 0$.

Proposition 1. *Suppose that the seller observes the realized number m of uninformed buyers. In the unique equilibrium, $r_m = \overline{r}$ for each $m = 1, \ldots, n-1$, and $r_0 = t_m = r^*$ for each $m = 1, \ldots, n$.*

Proof. When the seller observes the realized number m of uninformed buyers, the bidding constraints (7) can be satisfied by withholding the good when the number of buyers who keep silent exceeds m. Then, there is no incentive for an informed buyer to keep silent. The proposition follows immediately from the unconstrained solution to maximizing the objective function (6). ■

Proposition 1 provides a revenue upper-bound due to the absence of the bidding constraints. With at least one uninformed buyer, the seller has an outside option of making the best take-it-or-leave-it offer r^*. The value of this outside option is $\pi(r^*)$, independent of the realized number m of uninformed buyers, which means that the seller's best reserve price is $\overline{r}$, which equates the virtual value to $\pi(r^*)$. When $m = 0$, however, the value of the seller's outside option is 0 by assumption, so the best reserve price is r^*, which equates the virtual value to 0 as in the standard symmetric auction.

When the seller must commit to a mechanism that provides incentives for informed buyers to reveal the number of m, the mechanism given in Proposition 1 is not part of any equilibrium. The bidding constraints (7) are violated for informed buyers with value v_i just below $\overline{r}$, as

$$U(v_i, \epsilon) = \int_{r^*}^{v_i} (1-\alpha)^{n-1} F^{n-1}(w) dw,$$

which is smaller than

$$U(v_i, \mu) = \sum_{m=0}^{n-1} \frac{b(m; n-1, \alpha)}{m+1} F^{n-1-m}(\overline{r})(v_i - r^*).$$

An informed buyer with value v_i just below $\overline{r}$ wins the auction only when all other $n-1$ buyers are informed and have values below v_i, in which case the latter do not meet the reserve price $\overline{r}$ so the former would get the good with probability one at the lowest price of r^* by pretending to be the only one who is uninformed.

Indeed, if the seller wishes to set different reserve prices for the auction among informed buyers depending on whether or not there is the outside option of making a take-it-or-leave-it offer to uninformed buyers, informed buyers must be kept in dark about which reserve price will be used before they report to the seller. We next show that if instead informed buyers know the reserve price, the seller is prevented from exploiting the higher outside option when there is at least one uninformed buyer. As a result all reserve prices and all take-it-or-leave-it offers have to be the same.

Proposition 2. *Suppose that the seller is restricted to the same reserve price r in auctions. In the unique equilibrium, $r^* < r < \overline{r}$ and $t_m = r$ for each $m = 1, \ldots, n$.*

Proof. First, we claim that the bidding constraints (7) are satisfied if and only if $r \le t_m$ for all $m = 1, \ldots, n$. For the "only if" part, suppose that $r > t_m$ for some m. Then, the payoff $U(v_i, \epsilon)$ for an informed buyer with value v_i just below r is 0, but the buyer gets a strictly positive payoff by pretending to be uninformed, violating the bidding constraints. For the "if" part, from (7) we have $U(v_i, \epsilon) = U(v_i, \mu)$ for all $v_i \le r$, and

$$
\begin{aligned}
\frac{dU(v_i,\epsilon)}{dv_i} &= \sum_{m=0}^{n-1} b(m, n-1; \alpha) F^{n-1-m}(v_i) \\
&> \sum_{m=0}^{n-1} b(m, n-1; \alpha) \frac{F^{n-1-m}(r)}{m+1} \ge \frac{dU(v_i,\mu)}{dv_i}
\end{aligned}
$$

for all $v_i > r$. Thus, the bidding constraints are satisfied.

Next, we argue that $r > r^*$. Suppose instead $r \le r^*$. If $r < t_m$ for all $m = 1, \ldots, n$, we can increase r marginally without affecting the bidding constraints. The effect of this on the value of the objective (6) has the same sign as

$$
-n(1-\alpha)^n \phi(r) F^{n-1}(r) + \sum_{m=1}^{n-1} b(m; n-1, \alpha) \left(-\phi(r) + \pi(t_m)\right) F^{n-1-m}(r).
$$

Since $r \leq r^*$, the first term above is non-negative while the second term is strictly positive. This contradicts the constrained optimality of the seller's equilibrium mechanism. If $r = t_m$ for some $m = 1, \ldots, n$, we can increase r and all such t_m by the same marginal amount, keeping the bidding constraints satisfied. In addition to the above positive effect, there is a further non-negative effect on the revenue due to the increase in t_m, which is non-negative because $t_m = r \leq r^*$. This leads to a contradiction again.

Now, we can show that $t_m = r$ for all $m = 1, \ldots, n$. If $t_m > r$ for some m, the seller can marginally decrease t_m without affecting the bidding constraints. This however increases the seller's revenue because $t_m > r > r^*$, contradicting the constrained optimality of the seller's equilibrium mechanism.

Finally, we show that $t_m = r < \overline{r}$ for all $m = 1, \ldots, n$. If instead $t_m = r \geq \overline{r}$, the seller can marginally decrease all t_m and r by the same amount without affecting the bidding constraints. The effect of the decrease in the reserve price r on the seller's revenue (6) has the same sign as

$$n(1-\alpha)^n \phi(r) F^{n-1}(r) + \sum_{m=1}^{n-1} b(m; n-1, \alpha) \left(\phi(r) - \pi(r)\right) F^{n-1-m}(r),$$

which is strictly positive because $r \geq \overline{r}$ implies that $\phi(r) > 0$ and $\phi(r) > \pi(r)$. The effect of the decrease in the take-it-or-leave-it offer r on (6) is

$$-\sum_{m=1}^{n} b(m; n, \alpha) F^{n-m}(r) \frac{d\pi(r)}{dr},$$

which is again strictly positive because $r \geq \overline{r}$. This contradicts the constrained optimality of the seller's equilibrium mechanism. ■

The single-reserve price auction characterized by Proposition 2 provides a revenue lower bound for our equilibrium auction. The seller satisfies the bidding condition by setting all reserve prices and all take-it-leave-it offers to a single number r between r^* and $\overline{r}$. Compared to the auction characterized by Proposition 1, r is too high for the take-it-or-leave-it offers to uninformed buyers and for the auction participated by all n buyers, while at the same time too low for all the other auctions.

4.2 Price dispersion

The auction characterized by Proposition 2 is not part of an equilibrium. Starting from any common reserve prices and take-it-or-leave-it offers r, the seller could set a marginally lower reserve price r_0 for the auction when all buyers are informed. Since $r > r^*$ by Proposition 2, this would already increase the revenue from that auction while relaxing the bidding constraints. Further, the seller could marginally lower all take-it-or-leave-it offers to increase the revenue from uninformed buyers. This relaxes the bidding constraints, and allows the seller to increase all reserve prices when there is at least one uninformed buyer. The argument in the proof of Proposition 2 shows that r satisfies $\phi(r) < \pi(r)$, so the revenue in these auctions are further increased. We next proposition uses the constrained optimality result of Theorem 2 to establish necessary relations in reserve prices and take-it-or-leave-it offers in equilibrium auctions.

Proposition 3. *In any equilibrium, $r_0 < r^* < t_m < \max_{m'} r_{m'}$ for all $m = 1, \ldots, n$, and $\phi(r_m) < \pi(t_m)$ for all $m = 1, \ldots, n-1$.*

Proof. If $r_0 > r^*$, then by marginally decreasing it, the seller increases the revenue (when all buyers are informed) while relaxing the bidding constraints, as $U(v_i, \epsilon)$ is shifted up and $U(v_i, \mu)$ is unchanged. This is a contradiction to the constrained optimality of the equilibrium mechanism characterized by Theorem 2.

If $t_m < r^*$ for some m, then by marginally increasing it, the seller increases the revenue from uninformed buyers (when there are m of them) while relaxing the bidding constraints, as $U(v_i, \epsilon)$ is unchanged and $U(v_i, \mu)$ is shifted down, which contradicts the constrained optimality of the equilibrium mechanism.

If $\phi(r_m) > \pi(t_m)$ for some m, then by marginally decreasing r_m, the seller increases the revenue because the effect of the decrease has the same sign as

$$b(m; n-1, \alpha)\,(\phi(r_m) - \pi(t_m))\, F^{n-1-m}(r_m) f(r_m),$$

which is positive. At the same time, the bidding constraints are relaxed, as $U(v_i, \epsilon)$ is shifted up and $U(v_i, \mu)$ is shifted down, a contradiction to the constrained optimality.

Combining the above three claims, we have shown that $r_0 \leq r^* \leq t_m$ for all $m = 1, \ldots, n$, and $\phi(r_m) \leq \pi(t_m)$ for all $m = 1, \ldots, n-1$. If none of these inequalities is strict, we have the mechanism given by Proposition 1, which violates the bidding constraints. Thus, at least one of them is strict. If there is one that holds as an equality, the seller can make the same marginal change as above (decrease r_0 if $r_0 = r^*$, increase t_m if $t_m = r^*$, and decrease r_m if $\phi(r_m) = \pi(t_m)$) to strictly relax the bidding constraints with no effect on the revenue. But this would allow the seller to use any strict inequality to make the opposite marginal change to increase the revenue without violating the bidding constraints. This contradicts the constrained optimality of the equilibrium mechanism. Thus, $r_0 < r^* < t_m$ for all $m = 1, \ldots, n$, and $\phi(r_m) < \pi(t_m)$ for all $m = 1, \ldots, n-1$.

Finally, if $t_m \geq \max_{m'} r_{m'}$ for some m, then reducing t_m marginally keeps the bidding condition satisfied. This follows because the decrease does not affect $U(v_i, \epsilon)$ for any v_i, nor does it affect $U(v_i, \mu)$ for any $v_i < t_m$. At the same time, for all $v_i \geq \max_{m'} r_{m'}$ we continue to have $U(v_i, \epsilon) > U(v_i, \mu)$, as the proof of Proposition 2 implies that

$$\frac{dU(v_i, \epsilon)}{dv_i} > \frac{dU(v_i, \mu)}{dv_i}$$

regardless the take-it-or-leave-it offers. However, we have shown that $t_m > r^*$, so the seller's revenue is increased, a contradiction to the constrained optimality of the equilibrium mechanism. ■

An implication of Proposition 3 is dispersion in the reserve prices and the take-it-or-leave-it offers when the seller faces uncertainty about whether buyers are informed about the selling mechanism or not. When $\alpha = 0$, we have the standard auction for the n informed buyers with the reserve price r^*, while when $\alpha = 1$, we have the take-it-or-leave-it price offer of r^* to one randomly selected from the n uninformed buyers. There is no price dispersion in either extreme case, because the seller has the same zero value of keeping the good. However, away from the two extremes, the informed and uninformed are treated differently, as the former participate in auctions while the latter can only hope to get a take-or-leave-it offer. This means that the seller can give the good to a randomly selected uninformed buyer after failing to sell it through an auction among the informed buyers. Price dispersion emerges as a result. In particular, there are at least two distinct reserve prices for the auctions, $r_0 < r^*$

when all buyers turn out to be informed, and the highest reserve price $\max_{m'} r_{m'} > r^*$, that bracket all equilibrium take-it-or-leave-it offers.

The seller commits ex ante to reserve prices in auctions for informed buyers and take-it-or-leave-it price offers for the uninformed that are depend on the ex post extent of the knowledge about the commitments. This dependence turns out to have limited impact on equilibrium price dispersion. The reason is that marginal adjustments in the reserve prices or in the take-it-or-leave-it offers impact both the bidding constraints and the revenue, but the ratio of the two is independent of the realized number m of uninformed buyers. Instead, equilibrium dispersion in reserve prices and in take-it-or-leave-it offers is limited by the value distribution.

Our next result gives a sufficient condition on the distribution that eliminates any degree of dispersion in the take-it-or-leave-it price offers for the uninformed in equilibrium. Since by assumption $\phi(v)$ is increasing, there is a unique $\bar{t} \in (r^*, \bar{r})$ such that

$$\phi(\bar{t}) = \pi(\bar{t}).$$

Proposition 4 below shows that if $\pi(p)$ is strictly concave for $p \geq r^*$, then there exists some $\hat{t} \in (r^*, \bar{t})$ such that $t_m = \hat{t}$ for all $m = 1, \ldots, n$. This is helpful in characterizing the equilibrium auctions, because it implies that when the revenue function π is concave, a continuum of bidding constraints can be reduced to a single one, at some $\hat{v}$. Of course, the choice of $\hat{v}$ remains endogenous in any unobserved mechanism design problem. Nonetheless, the result that the bidding constraints bind only at one value allows us to reduce the amount of dispersion in equilibrium reserve prices for informed buyers.

Proposition 4. *Suppose $\pi(p)$ is strictly concave for $p \geq r^*$. In any equilibrium, there exist $\hat{t} \in (r^*, \bar{t})$ such that $t_m = \hat{t}$ for all $m = 1, \ldots, n$ and a unique $\hat{v} \in (\hat{t}, \bar{r})$ such that $U(\hat{v}, \epsilon) = U(\hat{v}, \mu)$. Furthermore,*

$$\alpha\phi(\hat{t})f(\hat{t}) + (1 - \alpha)\phi(r_0)f(r_0) = 0, \tag{9}$$

and if $r_m < \hat{v} < r_{m'}$ *for* $m, m' \geq 1$ *then*

$$(\phi(r_{m'}) - \phi(r_m))f(r_m) + \phi(r_0)f(r_0) = 0. \tag{10}$$

Proof. First, suppose that $t_m > t_{m'}$. Consider marginally decreasing t_m and simultaneously increasing $t_{m'}$ such that $U(v_i, \mu)$ for any v_i just above t_m stays unchanged, where

$$\frac{\partial U(v_i, \mu)}{\partial t_m} = -\frac{1}{m} b(m-1; n-1, \alpha) F^{n-m}(r_m),$$

with a similar expression for $\partial U(v_i, \mu)/\partial t_{m'}$. These marginal changes lower $U(w, \mu)$ for $w \in (t_{m'}, t_m)$, and have no other effects for any value $w \leq t_{m'}$ or $w \geq t_m$. Further, $U(w, \epsilon)$ remains unchanged. Thus, the bidding condition remains satisfied. To compute the effects on the seller's revenue R given by (6), note that

$$\frac{\partial R}{\partial t_m} = b(m; n, \alpha) F^{n-m}(r_m) \frac{d\pi(t_m)}{dt},$$

with a similar expression for $\partial R/\partial t_{m'}$. Since the ratio of $\partial R/\partial t_m$ to $\partial U(v_i, \mu)/\partial t_m$ is the same as that of $\partial R/\partial t_{m'}$ to $\partial U(v_i, \mu)/\partial t_{m'}$, the change in the seller's revenue has the same sign as

$$-\frac{d\pi(t_m)}{dt} + \frac{d\pi(t_{m'})}{dt}.$$

The above is positive because $\pi(p)$ is strictly concave for $p \geq r^*$. This contradicts the constrained optimality of the equilibrium mechanism characterized by Theorem 2. It follows immediately that there exists some $\hat{t}$ such that $t_m = \hat{t}$ for all $m = 1, \ldots, n$. By Proposition 3, since $t_m > r^*$ for all m, we have $\hat{t} > r^*$; since $\phi(r_m) < \pi(\hat{t})$ for all $m \geq 1$ and $\hat{t} < \max_{m'} r_{m'}$, and since $\phi(v) \geq \pi(v)$ for all $v \geq \bar{t}$ by the definition of $\bar{t}$, we also have $\hat{t} < \bar{t}$.

Next, since $U(v_i, \epsilon)$ is strictly convex for all $v \geq r_0$, the bidding constraints can bind only at a single value. It follows that there is a unique value $\hat{v}$ such that $U(\hat{v}, \epsilon) = U(\hat{v}, \mu)$. Since $U(v_i, \mu) = 0$ for all $v_i \leq \hat{t}$, we have $\hat{v} > \hat{t}$; and since $\phi(r_m) < \pi(\hat{t})$ for all m by Proposition 3, implying that $r_m < \bar{r}$, and since $dU(v_i, \epsilon)/dv_i > dU(v_i, \mu)$ for all $v_i \geq \max_{m'} r_{m'}$ from the proof of Proposition 2, we also have $\hat{v} < \bar{r}$.

Since $r_0 < r^* < \hat{t} < \hat{v}$, a marginal increase in r_0 reduces $U(\hat{v}, \epsilon)$ without changing $dU(\hat{v}, \epsilon)/dv_i$ or $U(\cdot, \mu)$, with

$$\frac{\partial R/\partial r_0}{\partial U(\hat{v}, \epsilon)/\partial r_0} = \frac{-n(1-\alpha)^{n-1}\phi(r_0)F^{n-1}(r_0)f(r_0)}{-(1-\alpha)^{n-1}F^{n-1}(r_0)}$$
$$= n(1-\alpha)\phi(r_0)f(r_0).$$

From the first part of this proof, a marginal increase in t_m for any $m = 1, \ldots, n$ from $\hat{t}$ reduces $U(\hat{v}, \mu)$ without changing $dU(\hat{v}, \mu)/dv_i$ or $U(\cdot, \epsilon)$, with

$$\frac{\partial R/\partial t_m}{\partial U(\hat{v}, \mu)/\partial t_m} = n\alpha\phi(\hat{t})f(\hat{t}).$$

Thus, a necessary condition for equilibrium is (9).

Finally, suppose that $r_m < \hat{v} < r_{m'}$ for $m, m' = 1, \ldots, n-1$. Consider a marginal change in r_m. At the same time, marginally change $r_{m'}$ in the opposite direction of r_m so that the function $U(v_i, \mu)$ stays the same, and marginally change r_0 in the opposite direction of r_m so that the value of $U(\hat{v}, \epsilon)$ stays the same. By construction, the bidding constraints remain binding at the same single value $\hat{v}$. A necessary condition for equilibrium is then

$$1 = \frac{\partial R/\partial r_{m'}}{\partial U(\hat{v}, \mu)/\partial r_{m'}} \frac{\partial U(\hat{v}, \mu)/\partial r_m}{\partial R/\partial r_m} + \frac{\partial R/\partial r_0}{\partial U(\hat{v}, \epsilon)/\partial r_0} \frac{\partial U(\hat{v}, \epsilon)/\partial r_m}{\partial R/\partial r_m}.$$

We have

$$\frac{\partial R}{\partial r_m} = b(m; n, \alpha)(n-m)F^{n-1-m}(r_m)f(r_m)(-\phi(r_m) + \pi(\hat{p})),$$

and

$$\frac{\partial U(\hat{v}, \mu)}{\partial r_m} = \frac{n-m}{m} b(m-1; n-1, \alpha)F^{n-1-m}(r_m)f(r_m)(\hat{v} - \hat{t}),$$

with similar expressions for m', while

$$\frac{\partial U(\hat{v}, \epsilon)}{\partial r_m} = -b(m; n-1, \alpha)F^{n-1-m}(r_n),$$

because $r_m < \hat{v}$ and $\partial U(\hat{v}, \epsilon)/\partial r_{m'} = 0$ because $r_{m'} > \hat{v}$. Using the above expressions, we can rewrite the necessary equilibrium condition as

$$1 = \frac{-\phi(r_{m'}) + \pi(\hat{t})}{-\phi(r_m) + \pi(\hat{t})} + \frac{f(r_0)(-\phi(r_0))}{f(r_m)(-\phi(r_m) + \pi(\hat{t}))},$$

which reduces to (10). ■

An immediate implication of (10) in Proposition 4 is that, if $r_m, r_{m'} > \hat{v}$ for $m, m' = 1, \ldots, n-1$ then $r_m = r_{m'}$. Therefore, high reserve prices—those that may bind for informed buyers with sufficiently high values who strictly prefer participating in bidding —should not be discriminatory depending on the realized number of informed buyers, if take-it-or-leave-it offers are the same. On the other hand, we do not have a similar result for low reserve prices that bind for informed buyers with low values. This is because changes to such low reserve prices can affect the equilibrium payoff of informed buyers with a critical value who are just indifferent between participating in bidding and keeping silent. Low reserve prices can have allow for some dispersion. However, as we have already mentioned, such dispersion is entirely pinned down by the value distribution function F. In particular, if $F(v)$ is weakly decreasing, then (10) implies that all low reserve prices are the same as well.[12]

The two necessary conditions (9) and (10) in Proposition 4 are independent of the total number of buyers n. Thus, how the seller sets the reserve prices and the take-it-or-leave-it offers in equilibrium is not sensitive to how many potential buyers there are, at least when π is concave. Combining (9) and (10), we find that the dispersion among the reserve prices in auctions for informed buyers disappears, as all reserve prices (other than r_0) converge to some critical value $\hat{v}$, when the probability α that a given buyer is uninformed goes to zero. Further, when α is arbitrarily close to 0, from (9) we have that r_0 converges to r^*. This makes sense, because when each buyer is almost surely informed, the seller's revenue is almost completely determined by the auction corresponding to $m = 0$, so the reserve price r_0 in this auction needs to be arbitrarily close to the one used in the standard auction.

[12]Thus, for the uniform value distribution, there are at most four reserve prices: r_0, $\hat{v}$, a reserve price lower than $\hat{v}$ and another one higher than $\hat{v}$.

4.3 Equilibrium welfare

The presence of uninformed buyers makes the seller worse off. We have modeled uninformed buyers as agents who do not observe the seller's commitments but who nonetheless have rational expectations regarding the seller's mechanism that prevents the seller from exploiting their lack of knowledge. This effectively makes it harder for the seller to use standard devices to elicit information about their values from those buyers. Theorem 1 shows that in equilibrium it is as if the seller has no commitment with respect to uninformed buyers. A simple revealed-preference type of argument then shows that the seller is weakly worse off in equilibrium whenever buyers might be uninformed compared to the standard environment of when all buyers are informed of the seller's mechanism. Further, since the seller could have used auctions instead of making take-it-or-leave-it offers to uninformed buyers, the seller is strictly worse off in equilibrium than in the standard setting.

If there are no informed buyers, uninformed buyers have an equal chance of receiving a take-it-or-leave-it offer equal to r^*. The resulting payoff, as a function of the value v_i, is

$$U^*(v_i, \mu) = \frac{1}{n} \max\{v_i - r^*, 0\}.$$

We now compare the above to the equilibrium payoff function of the uninformed, $U(v_i, \mu)$ as given in (7).

Since $t_m > r^*$ for all $m = 1, \ldots, n$ by Proposition 3, the uninformed buyer with a value v_i just above r^* is worse off in equilibrium than when there are no informed buyers around. This is due to an "incentive effect" because the seller in equilibrium raises all take-it-or-leave-it offers to the uninformed above r^* in order to discourage informed buyers from pretending to uninformed. For uninformed buyers with high values that clear all take-it-or-leave-it offers, however, the "outside option effect" that the seller can give the good to an uninformed buyer when no informed buyer wins the auction instead of having to keep it, allows the seller to raise at least some of the reserve prices above r^*. This potentially reduces the probability that the good is sold to informed buyers and thus improves the chance that it goes to a given uninformed buyer.

When the total number of buyers n is sufficiently large, the outside option effect can be shown to be negative in aggregate on the payoff of the uninformed. From Proposition 3, we have $r_m < \overline{r}$ for all $m = 1, \ldots, n-1$ because $\phi(r_m) < \pi(t_m)$ and $t_m > r^*$. Thus, the equilibrium payoff of the uninformed buyer with the highest value of 1 is bounded from above by

$$\begin{aligned} U(1,\mu) &< \sum_{m=0}^{n-1} \frac{b(m; n-1, \alpha)}{m+1} F^{n-1-m}(\overline{r})(1-r^*) \\ &= \frac{1}{n\alpha}\left((\alpha + (1-\alpha)F(\overline{r}))^n - ((1-\alpha)F(\overline{r}))^n\right)(1-r^*). \end{aligned}$$

The above is less than $U^*(1,\mu)$ when n is sufficiently large. What happens is that when n is large, averaging across the realized number of uninformed buyers m, the probability that no informed buyer meets the reserve price r_m is sufficiently small so that the reduced competition in the lottery for the uninformed from n to m becomes relatively unimportant. Since $U^*(v_i,\mu)$ is linear for values above r^* and $U(v_i,\mu)$ is convex, an uninformed buyer with any value is worse off in equilibrium than when there are no informed buyers.

When all n buyers are informed, they face the standard optimal reserve price of r^*, with the payoff function

$$U^*(v_i,\epsilon) = \int_{\min\{v_i, r^*\}}^{v_i} F^{n-1}(w)dw.$$

We now compare the above to the equilibrium payoff function of the informed, $U(v_i,\epsilon)$ as given in (7).

Since $r_0 < r^*$ by Proposition 3, the informed buyer with a value v_i between r_0 and r^* is better off in equilibrium than when there are no uninformed buyers around. Thus, the same incentive effect that makes uninformed buyers worse off in the presence of informed buyers now motivates the seller to lower the reserve price r_0 below r^*, making informed buyers with low values better off in the presence of uninformed buyers. For informed buyers with sufficiently high values, again their payoff is impacted by the outside option effect in opposing directions. So long as some buyers turn out to be uninformed, i.e. $m \geq 1$, the equilibrium reserve price r_m in the auction for informed buyers is higher than r^*, but at the same time an informed buyer competes against $n-1-m$ instead of $n-1$ other informed buyers.

In contrast to our discussion of the welfare of the uninformed, the net effect is positive on the payoff of the informed buyers with high values when the total number of buyers n is sufficiently large. Since $r_m < \overline{r}$ for all $m = 1, \ldots, n-1$, we have

$$\begin{aligned} U(1,\epsilon) &> \sum_{m=0}^{n-1} b(m; n-1, \alpha) \int_{\overline{r}}^{1} F^{n-1-m}(w)dw \\ &= \int_{\overline{r}}^{1} ((\alpha + (1-\alpha)F(w))^{n-1} dw. \end{aligned}$$

The above is greater than $U^*(1,\epsilon)$ when n is sufficiently large.[13] What happens is that when n is large, averaging across the realized number of uninformed buyers m, the reduced competition in the auctions for informed buyers from n to $n-m$ becomes relatively more important than the higher reserve prices they face. The comparison with the standard auction for informed buyers with intermediate values remains ambiguous.

4.4 An example

Suppose that $n = 2$ and that the revenue function $\pi(p)$ is strictly concave for $p > r^*$. By Proposition 4, there is a single equilibrium take-it-or-leave-it offer $\hat{t} \in (r^*, \overline{t})$. It follows that the single value $\hat{v}$ at which the bidding constraints bind coincides with the only other reserve price r_1 above r^*. The choice variables for the constrained optimization problem given in Theorem 2 are then $\hat{t}$, r_0 and r_1, and the continuum of bidding constraint reduces to a single one, given by

$$(1-\alpha)\int_{r_0}^{r_1} F(w)dw \geq \left(\alpha \frac{1}{2} + (1-\alpha)F(r_1)\right)(r_1 - \hat{t}). \tag{11}$$

Let λ be the Lagrange multiplier for the above constraint. We can use λ to rewrite the first order necessary conditions in Proposition 4 for an equilibrium mechanism $(\hat{t}, r_0, r_1)$,

[13] Using integration by parts, for n large enough, it suffices to show that

$$(1-\alpha)\int_{\overline{r}}^{1} ((\alpha + (1-\alpha)F(w))^{n-2} f(w)wdw < \int_{r^*}^{1} F^{n-2}(w)f(w)wdw.$$

The above is true for large enough n by using another integration by parts.

equations (9) and (10), as

$$
\begin{aligned}
2\alpha\phi(\hat{t})f(\hat{t}) &= \lambda, \\
2(1-\alpha)\phi(r_0)f(r_0) &= -\lambda, \\
2(1-\alpha)\alpha(\pi(\hat{t})-\phi(r_1))f(r_1) &= \lambda\left(\alpha\frac{1}{2}+(1-\alpha)(r_1-\hat{t})f(r_1)\right),
\end{aligned} \tag{12}
$$

together with (11) holding with equality.

This example illustrates that the equilibrium mechanism is discontinuous at $\alpha = 0$ and $\alpha = 1$. In the two extremes, the reserve price for auctions and the price for take-it-or-leave-it offers are both r^*. In the limits, however, only the price relevant to the seller revenue - r_0 in the case of α going to 0 and $\hat{t}$ in the case of α going to 1 - converges to r^*, with the other prices pinned down by the incentive effect and the outside option effect that do not exist in the extremes.

For α arbitrarily close to 0, the first order conditions with respect to $\hat{t}$ and r_0 (the first two equations in 12) together imply that r_0 is close to r^*. At the same time, the first order conditions with respect to $\hat{t}$ and r_1 (the first and the last equations in 12) imply that if $\hat{t}$ converges to r^* as well, then r_1 converges to $\overline{r}$. This violates the constraint (11). Thus, for α converging to 0, $\hat{t}$ is bounded away from r^* and r_1 is in turn bounded away from $\hat{t}$.

For α arbitrarily close to 1, the first order conditions with respect to $\hat{t}$ and r_0 imply that $\hat{t}$ is close to r^*, which implies r_1 is close to r^* as well from the binding constraint (11). Then, the first order conditions with respect to r_0 and r_1 (the last two equations in 12) imply that the limit of r_0 satisfies

$$-\frac{1}{2}\phi(r_0)f(r_0) = \pi(r^*)f(r^*).$$

Thus, for α converging to 0, r_0 is bounded away from r^*.

This example can also be used to make the welfare comparisons for the uninformed and informed buyers away from the case of a large number of buyers. For the uninformed buyers, we can show when α is close to 0, in contrast to what happens when n is large, those with high values can be better off in equilibrium than in the absence of informed buyers.[14] To see

[14]Of course, when α is close 1, the equilibrium payoff function $U(v_i, \mu)$ for the uninformed becomes close to $U^*(v_i, \mu)$.

this, using the binding constraint (11) and integration by parts, we have

$$\begin{aligned}U(1,\mu) &= \left(\alpha\frac{1}{2}+(1-\alpha)F(r_1)\right)(r_1-\hat{t})+\left(\alpha\frac{1}{2}+(1-\alpha)F(r_1)\right)(1-r_1)\\&=(1-\alpha)\left(F(r_0)(1-r_0)+\int_{r_0}^{r_1}(1-w)f(w)dw\right)+\alpha\frac{1}{2}(1-r_1).\end{aligned}$$

We already know that r_0 converges from below to r^* when α is arbitrarily close to 0, while r_1 is bounded below away from r^*. Therefore, for any value distribution with the property that $F(r^*) \geq \frac{1}{2}$, including the uniform distribution, the above is strictly greater than $U^*(1,\mu)$. Further, since r_0 converges to r^*, the uninformed buyer is better off in equilibrium for almost all values.

For the informed buyers, when α is close to 1, consistent with what happens when n is large, those with high values are always better off in equilibrium than in the absence of uninformed buyers. This is simply because r_1 converges to r^* when α goes to 1.[15] Since r_0 is bounded away from r^*, the informed buyer is better off in equilibrium for all values.

5 Concluding Remarks

In this paper we have considered a traditional mechanism design problem and modified it by assuming some buyers do not know the mechanism the seller is using. Our main result is that to induce the informed to reveal themselves, the seller has to hide certain features of the mechanism from them. In this simple problem this results in a simple explanation for the fact that sellers hide their reserve prices in auctions - something that is observed frequently in practise. At the point where informed buyers report their values the reserve price is a random variable.

We have taken a number of shortcuts in our model - in particular, we assume that messages lead to a single offer. For the auction among the informed buyers this is without loss, since the winner of the auction always wants to accept the offer when they win the auction. For the uninformed this assumption is unrealistic. Once the seller learns who the

[15] In fact, when the value distribution is uniform, informed buyers with high values are always better off regardless of the value of α. Further, it can be shown that $U(r_1,\epsilon) < U^*(r_1,\epsilon)$ so that the informed buyer with intermediate values are worse off than when all buyers are informed.

uninformed buyers are, the seller is likely to approach them in sequence with offers. A general approach to unobserved mechanisms is to model the output of a mechanism as an "algorithm," which is a sequence of take-it-or-leave-it offers and the identities of the buyers to whom the offers are made. As in the present paper, the seller first makes a commitment in terms of how a particular algorithm is chosen in response to the messages sent by the buyers, who however may not observe it. It is straightforward to generalize the analysis in the present paper to the case in which algorithms are restricted to at most one take-it-or-leave-it offer for each buyer. The main insights are intact - an uninformed buyer receives an expected offer independent of the buyer's value, while informed buyers face a secret reserve price when they bid in an auction. We leave the characterization of unrestricted equilibrium algorithms to future research.

Perhaps a more restrictive assumption we use is that buyers are either fully informed or fully uninformed. A more reasonable assumption might be that buyers have partial information about commitments. For example, we could assume that some buyers may only be able to understand commitments to actions based on their own messages, but not commitments that depend on the messages of others. If all buyers have this type of partial information, then there is an equilibrium in which the seller implements the optimal auction of Myerson (1981) through a first-price sealed bid auction. This corresponds to the main result of Akbarpour and Li (2018), who frame the issue of partial observability in terms of limited commitment by the seller. When buyers have differential information about the seller's commitments - for example, if buyers either fully observe the seller's commitment or only observe the part based on their own message - we nonetheless believe that our basic insight could be extended to this kind of assumption. Yet we are reluctant to pursue without a better model of what buyers can and cannot understand.

Finally, our model is too stylized to make strong claims about what it predicts. However, it does have a fairly simple prediction. When the seller attempts to sell to an uninformed buyer, the seller does so at a fixed price, and this attempted sale will often fail. When the good is sold to an informed buyer the transaction is much more likely to succeed but prices will vary with realized values. The probability with which these two events occur depends on the probability with which buyers are informed. For this reason our model predicts that price variability will be larger in markets where buyers are likely to be informed. Conversely, prices will be stable and transaction variability much larger in markets where buyers are unlikely to be informed. This suggests a potential method for measuring the degree of understanding of mechanisms in auction markets.

References

[1] Akbarpour, M., and S. Li (2018): "Credible Mechanisms, " Discussion paper, Stanford University.

[2] Bester, H., and R. Strausz (2001): "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case, " *Econometrica*, 69(4),1077–1098.

[3] Butters, G. (1977): "Equilibrium Distributions of Sales and Advertizing Prices," *Review of Economic Studies* 44(3), 465–491.

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

[5] Deneckere, R., and S. Severinov (2008): "Mechanism design with partial state verifiability," *Games and Economic Behavior* 64(2), 487–513.

[6] Dickson, P. R., and A. G. Sawyer (1990): "The Price Knowledge and Search of Supermarket Shoppers," *Journal of Marketing*, 54(3), 251–334.

[7] Fudenberg, D., D. Levine, and J. Tirole (1985): "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game-Theoretic Models of Bargaining*, ed. by A. Roth, Cambridge University Press.

[8] Green, J., and J. J. Laffort (1986): "Mechanism Design with Partially Verifiable Information," *Review of Economic Studies*, 53(3), 447–456.

[9] Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory* 39(1), 155–190.

[10] Kolotilin, A., H. Li, and W. Li (2013): "Optimal Limited Authority for Principal," *Journal of Economic Theory* 148(6), 2344–2382.

[11] Krishna, V. (2010): *Auction Theory*, second edition, Academic Press.

[12] Liu, Q., Mierendorff, K., and X. Shi, (2014). "Auctions with Limited Commitment, " Discussion paper, Columbia University.

[13] Masatloglu, Y. (2015): "A Random Inattention Model," Discussion paper, University of Michigan.

[14] Masatloglu, Y., D. Nakajima, and E. Ozbay (2012): "Revealed Attention," *American Economic Review* 102(5), 2183–2205.

[15] McAfee, P. (1993): "Mechanism Design by Competing Sellers," *Econometrica* 61(6), 1281–1312.

[16] Mookherjee, D., and M. Tsumagari (2014): "Mechanism Design with Communication Constraints," *Journal of Political Economy* 122(5), 1094–1129.

[17] Myerson, R. B. (1981): "Optimal Auction Design, " *Mathematics of Operations Research* 6(1), 58–63.

[18] Peters, M. (2014): "Can Mechanism Designers Exploit Market Information," Discussion paper, Vancouver School of Economics.

[19] Riley, J., and R. Zeckhauser (1987): "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *The Quarterly Journal of Economics* 98(2), 267–289.

[20] Skreta, V. (2015). "Optimal Auction Design under Non-commitment," *Journal of Economic Theory*, 159 (Part B), 854–890.

[21] Stahl, D. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review* 79(4), 700–712.

[22] Stahl, D. (1994): "Oligopolistic Pricing and Advertising," *Journal of Economic Theory* 64(1), 162–177.

[23] Strausz, R. (2016): "Mechanism Design with Partially Verifiable Information," Discussion paper, Humboldt University.

[24] Varian, H. (1980): "A model of Sales," *American Economic Review* 70(4), 651–659.