RATIONAL IGNORANCE IN MECHANISM DESIGN: EQUAL PRIORITY AUCTIONS

ABSTRACT. We study the impact for mechanism design of the possibility that some participants are (rationally) uninformed about the rules associated with a trading mechanism. Since 'deviations' by the mechanism designer are not observed by these uninformed participants the nature of the 'equilibrium' of the design game changes, as do equilibrium mechanisms. We study the traditional independent private value auction environment and propose a method that makes it possible to characterize an interesting class of equilibrium outcomes for the game using standard reduced form direct mechanisms. We show that payoffs in the equilibrium where the seller's expected revenue is highest within this class can be characterized using a surprisingly simple mechanism called an *equal priority auction.* Equal priority means that with positive probability intermediate value buyers receive a fixed price offer that is independent of the of the values of the informed bidders, while high and low value buyers participate in an auction. The interesting new feature of this mechanism is that informed and uninformed buyers receive offers with the same probability despite the fact the seller believes that the informed will accept the offers for sure, while uninformed buyers might not.

1. Introduction

There is an acronym that floats around the internet - TLDNR - that explains why no one reads you email messages. It means “too long, didn’t read”. The long translation we adapt in this paper is “... there is undoubtedly information in your message, but its’ value to me isn’t likely to be as high as what I could get by reading something else”. We refer to this as ’rational ignorance’.

The message of this paper is that this kind of behavior can impact trading mechanisms. We aren’t the first to notice this. The marketing literature has documented buyers’ tendency to ignore information when they make purchase decisions. The simplest commitment of all is a price commitment. Dickson and Sawyer (1990) asked buyers in supermarkets about their price knowledge as they were shopping. Only 50% of all respondents to their in store survey claimed to know the price of the object they had just taken off the supermarket shelf to 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.

Of course, having buyers be pleasantly surprised to learn that a price is lower than they expected isn’t really a problem. The problem is the buyers who didn’t know the price was on special, and went somewhere else to buy it. If prices can’t influence buyer behavior, marketing has a problem.

We are interested in more than prices, we want to know how this kind of rational ignorance can impact trading mechanisms. We consider what is probably the best understood trading problem of all - the independent private value auction. We show that within a plausible class of equilibria - equilibria in which ignorant buyers convey no information to sellers - the revenue maximizing mechanism is something we call an *equal priority mechanism.*

The equal priority mechanism treats informed buyers and sellers with intermediate valuations in exactly the same way as ignorant buyers. When a mechanism chooses to attempt to trade with them, it makes them a take it or leave it price offer that is independent of any messages they may have sent. When buyers have very high or very low valuations, the seller treats messages as bids. If the seller decides to sell to one of these buyers, she will make an offer equal to the second highest bid she has received - much as she would in a standard auction.

Rationally ignorant buyers will be pooled with intermediate value buyers and receive a take it or leave it offer (which they might reject). In our formulation, this offer will be exactly the offer the buyers expected to receive. In other words, these buyers have rational expectations - there is nothing behavioral about them at all.

One appealing feature of the independent private value auction problem for mechanism design is that finding the revenue maximizing mechanism can be reduced to a problem of solving a maximization problem with a single parameter - the reserve price. The revenue maximizing mechanism with rationally ignorant buyers can be found by solving a problem with four parameters - harder, but still computationally tractable. The numerical solutions we have found in simple environments suggest that fixed priced trading is quite common. In fact, it is easy to show theoretically that if every buyer is equally likely to be informed or rationally ignorant, the trade will occur at a fixed price (with no auctions) much more than half the time. This may be another explanation for why auctions aren’t particularly common in many consumer trading platforms.

One well known trading platform on which auctions *are* used is eBay. The environment on eBay doesn't fit our model exactly because of timing issue, but the mechanisms used by eBay resemble the equal priority action we describe below. A seller can implement something very close to what we describe here by running an auction with a 'buy it now' option, then offering the same model separately at a fixed price. Since buy it now options disappear on eBay once a low valued buyer submits a bid, trade will occur at the fixed price (which should be the same as the buy it now price) a lot of the time, though auctions will continue to occur.

1.1. **Heuristic.** The formal derivation below is based on two arguments.

The first is that in an environment with uninformed buyers, standard auction mechanisms can't be supported as equilibrium even though the seller would much prefer to use them. The fault lies with the seller who can't resist the temptation of exploiting rationally ignorant buyers.

To see why, suppose the seller wants to use a second price auction with optimal reserve. This means that informed buyers read the auction rules, as they might on eBay, then realize they should bid their values. Uninformed buyers don't read the rules, so they only *anticipate* a second price auction. Acting on their expectations, they also bid their values.

What makes this break down is the fact that if the seller changes the auction rules, the uninformed won't realize it, and will continue to bid their values no matter what the seller does. A simple deviation can extract the surplus of the uninformed.

The seller can 'deviate' from the second price auction and ask for bidders to attach a coupon code to their bid. The coupon code isn't secret, it is plainly visible in the description of the bidding rules. A buyer who reads the new rules will see the coupon code and attach it to their bid.

A bidder who doesn't read won't add the code. The new mechanism commits to a second price auction for bids submitted with a code, but to treat bids with no code attached as if it were a first price auction. In other words, if the highest bid is submitted by an uninformed bidder, the seller will commit to make them an offer which is equal to their bid, instead of offering them the second highest bid.

The second argument involves how the seller should respond. The seller will want to sell to the uninformed buyers when informed buyers have low values. So the natural idea would be to have an auction, then if bids are too low, make an offer to the uninformed. The complication is that informed bidders don't have to bid. They can pretend to be uninformed. Since they are informed, they know when the seller will make an offer to the uninformed and what that offer will be. To prevent the informed from pretending to be uninformed, the seller has to keep the offer to the uninformed higher than she would like it to be, since the seller is never sure whether an uninformed buyer will accept the offer.

The seller then faces a trade off - keep the offer high and fully separate the informed from the uninformed, or lower the take it or leave it offer and allow some of the informed buyer to pool with the uninformed. We show that the latter is always revenue maximizing, which is where the equal priority phrase comes from in our title.

1.2. **Literature.** As mentioned above, the idea that consumers might not notice prices is an old one in the marketing literature, as in Dickson and Sawyer (1990) and references therein. 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.[1] 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.

What ignorant buyers does it to provide a type dependent outside options to informed buyers. This is one of the most basic problems in the literature on competing mechanisms. One example is the paper by McAfee (1993). His model had buyers whose outside option involved waiting until next period to purchase in a competing auction market just like the one in the current period. He imposed large market assumptions to ensure that the value of these outside options was independent of the reserve price that any seller in the existing market chose.

In our paper, the value of this outside option depends on the nature of the mechanism the seller chooses for the informed. This makes it resemble the later papers on competing mechanisms (at least in terms of outside options) like Virag (2010) who studies finite competing auction models where a seller who raises her reserve price increases congestion in other auctions, or Hendricks and Wiseman (2020) who study the same problem in a sequential auction environment.

With buyers potentially uninformed of the selling mechanism but nonetheless having rational expectations, the seller's commitment power is limited. There is an extensive literature on limited commitment (for example Bester and Strausz (2001), Kolotilin et al. (2013), Liu et al. (2014), or Skreta (2015)). To our knowledge, our model is the first to study commit with respect to a subset of traders involved in the same transaction.

A recent paper by Akbarpour and Li (2020) 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. The logic we described above explaining why the second price auction can't survive as an equilibrium is used in a similar way in their paper. The difference between their approach and ours is that they assume the credibility constraint applies to all buyers and describe mechanisms that are immune to this constraint. Here we assume that credibility is an issue only for some buyers and find optimal mechanisms.

Finally, our informed buyers can 'prove' they are informed in the same sense as Ben-Porath et al. (2014). The main difference is that they assume that the social choice function is known by all the players. This is not the case here. They also assume players have complete information about the state, which is not the case here.

2. The Model

There are n potential buyers of a single homogeneous good. Each buyer has a privately known valuation w that is independently drawn from the interval $[0,1]$. Assume for the moment, all valuations are distributed according to some distribution F with strictly positive density f. Buyers' payoff when they buy at price p is

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

given by $w - p$. The seller's cost is zero, so the profit from selling at price p is just p.

Each buyer is *informed* about the rules of the seller's mechanism with probability $(1 - \alpha)$. Otherwise a buyer is rationally ignorant and pays no attention to the rules of the mechanism.

Define

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

as the revenue function from a take-it-or-leave-it offer w to uninformed buyers. In what follows, we restrict attention to distribution functions such that $\pi(w)$ is strictly concave. Following the standard auction literature, we also define

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

as the virtual valuation function for informed buyers. We have $\phi(0) < 0$ and $\phi(1) = 1$, and so $\phi(w)$ crosses 0 at least once. Since $\pi'(w) = -\phi(w)f(w)$, concavity of $\pi(\cdot)$ implies that $\phi(w)$ crosses 0 only once. Let the crossing point be r^*; this is also the unique maximizer of $\pi(w)$. Furthermore, $\phi(w)$ is strictly increasing in v for $w \geq r^*$.[2] The valuation r^* represents the optimal reserve price in a standard auction, regardless of the number of buyers.[3] That is, when $\alpha = 0$, the seller's outside option is always 0, so the reserve price is such that the virtual valuation of the buyer with w at the reserve price is equal to the seller's outside option.

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 product of the set of buyer valuations and the interval $[0, 1]$. The message space $\mathcal{M}$ is common knowledge.

After processing all the buyers' messages, the seller's mechanism makes an *offer* to one of the buyers.[4] Our assumption is that it is common knowledge that all buyers understand and believe a take it or leave it price commitment.

This offer can be refused. For the moment, we'll assume that when it is, there is simply no trade at all. This makes for a cleaner analysis of the issues we are interested in. We'll explain later the sense in which the mechanism we describe will be part of an equilibrium even if multiple offers are allowed after the first one is rejected.

We'll also defer until later to show the sense in which in this environment, sellers won't offer a mechanism that commits buyers to accept offers ex post.

---

[2]At any $w \in (0, 1)$, if $f(w)$ is non-decreasing, then by definition $\phi(w)$ is strictly increasing; if $f(w)$ is strictly decreasing at w and if $\phi(w) \geq 0$, then $\phi(w)$ is strictly increasing in w, because concavity of $\pi(w)$ implies that $\phi(w)f(w)$ is strictly increasing in w.

[3]In much of the auction literature, the seller has the fixed outside option of keeping the good. The virtual valuation function $\phi(w)$ is assumed to be strictly increasing to simplify the analysis (the "regular case" in Myerson (1981)). In our model, the seller's outside option in an auction with informed buyers is to give it to an uninformed buyer with a take-it-or-leave-it offer, and is endogenous. We do not need to assume that $\phi(w)$ is strictly increasing for valuations below r^*.

[4]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.

Whether or not a buyer has taken the time to understand how a mechanism works is private information. In what follows we'll use the convention that a buyer who has figured out the mechanism is referred to as an informed buyer. One who hasn't is just called an uninformed buyer. So buyer i's type is given by the pair (v_i, ι_i), where $v_i \in [0,1]$ and $\iota_i \in \{\epsilon, \mu\}$, where ϵ means informed, and μ means uninformed. Each buyer has type μ with probability α that is independent of their valuation or the types of the other bidders.

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 designed for buyer i, while q_i maps the same set of messages into the probability with which the offer will actually be made to buyer i. The mapping q_i must satisfy

$$\sum_i q_i (b_1, \ldots, b_N) \leq 1$$

for every profile of messages. Let Γ be the set of feasible mechanisms.[5]

Definition. The imperfect information game $\mathcal{G}(\alpha)$ is defined to be the extensive form game of imperfect information in which the seller first commits to some $\gamma \in \Gamma$, the informed buyers send messages to the seller that depend on γ, and the uninformed send messages that are independent of γ. Allocations and final payoffs are determined by the mechanism γ, the realized messages, and the acceptance decisions of any buyer who receives an offer. The parameter α gives the common belief with which each buyer and the seller believes that each of the buyers is uninformed.

A strategy rule in $\mathcal{G}(\alpha)$ for a buyer is a function $\sigma_i : [0,1] \times \{\epsilon, \mu\} \times \Gamma \to \triangle(\mathcal{M})$ that specifies what message the buyer will send for each of the valuations conditional on whatever the buyer knows about the seller's mechanism. Since an uninformed 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'. 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, providing a coupon code that must be submitted with a bid. One of the stranger properties of equilibrium in this game is that pure strategy (for the seller) equilibrium typically won't exist. If such equilibrium did exist, uninformed bidders would guess the coupon code. This is just because of the fact that strategies are common knowledge in any Nash based equilibrium. This is just another way of saying that uninformed bidders have 'rational expectations'.

Once the uninformed guess the coupon code, they will submit informative bids and the seller won't be able to prevent himself from exploiting that information. To prevent the uninformed buyers from guessing this password, it has to be random. There is nothing secret about this password, it is freely available to anyone who takes the time to read the rules of the mechanism.

[5]We have thus restricted each mechanism γ in Γ to be a "pure" one, in that each p_i maps a profile of messages to a single price, instead of a distribution of prices. This can be easily dropped without affecting the formalism; for the equilibria we construct in the paper this restriction is without loss.

Refer to the seller's mixture as $\psi \in \triangle(\Gamma)$. Let $R\left(\gamma, \left(\sigma_i\left(\cdot,\cdot,\gamma\right)\right)_{i=1}^n\right)$ be the expected revenue for the seller from mechanism γ when an uninformed buyer i uses strategy $\sigma_i\left(\cdot,\mu\right)$ an informed buyer i uses $\sigma_i\left(\cdot,\epsilon,\gamma\right)$.

A perfect Bayesian equilibrium for this game 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 that satisfy the usual boilerplate conditions, a perfect Bayesian equilibrium should satisfy for each informed buyer i

$$\mathbb{E}_{v_{-i},\tau_{-i}}\left\{q_i\left(\sigma_i\left(w,\epsilon,\gamma\right), \sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot \max\left[\left(w-p_i\left(\sigma_i\left(w,\epsilon,\gamma\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\right),0\right]\right\} \geq$$

$$\mathbb{E}_{v_{-i},\tau_{-i}}\left\{q_i\left(b',\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot \max\left[\left(w-p_i\left(b',\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\right),0\right]\right\} \tag{2.1}$$

for all $w \in [0,1]$, $b' \in \mathcal{M}$ and realized γ from the mixture ψ; while for the uninformed the condition is

$$\mathbb{E}_{v_{-i},\tau_{-i},\gamma}\left\{q_i\left(\sigma_i\left(w,\mu\right), \sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot \max\left[\left(w-p_i\left(\sigma_i\left(w,\mu\right),\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\right),0\right]\right\} \geq$$

$$\mathbb{E}_{v_{-i},\tau_{-i},\gamma}\left\{q_i\left(b',\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\cdot \max\left[\left(w-p_i\left(b',\sigma_{-i}\left(v_{-i},\tau_{-i},\gamma\right)\right)\right),0\right]\right\} \tag{2.2}$$

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

$$\mathbb{E}_{\gamma}\left\{R\left(\gamma,\left(\sigma_i\left(\cdot,\cdot,\gamma\right)\right)_{i=1}^n\right)\right\} \geq R\left(\gamma',\left(\sigma_i\left(\cdot,\cdot,\gamma'\right)\right)_{i=1}^n\right) \tag{2.3}$$

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

The reason that the max operation appears when taking expectations is because a mechanism generates an offer instead of an outcome.

2.1. **Relationship to standard mechanism design.** This game has many equilibrium outcomes. One source of multiplicity comes from a continuum of coordination equilibrium in which uninformed buyers send informative messages.

A babbling equilibrium is one example, no uninformed buyer will send an uninformative message because he or she believes the seller's mechanism won't respond to it. Since the seller thinks uninformed buyers are babbling, there is no reason for their mechanism to respond to these messages.

There are equilibrium outcomes in which very low value uniformed buyers will separate from higher value buyers by saying they aren't interested in an offer. To support this, some of the uninformed who believe they will never accept a seller offer must nonetheless act as if they might accept an offer. If the set of buyers who inform the seller they aren't interested is constructed in the right way, the seller will commit not to make them offers (even when he has no better alternative) for incentive reasons.

We won't discuss these equilibrium outcomes in detail because they are straightforward extensions of what we describe below.

However, we haven't yet been able to rule out equilibrium in which uninformed buyers send much richer information with their messages, which are effectively cheap talk.

All of these alternative equilibrum where uninformed buyers convey information to the seller exhibit the same logic that we describe below. For example, informative cheap talk messages allow the seller to make offers to the buyers whose messages say they are most likely to accept them. This effect tends to raise the seller's expected revenues. Yet to maintain incentive compatibility the seller has to make fairly high offers so that the informed are incentivized to reveal that they are informed. This makes it hard to determine whether equilibria with more information conveyed by uninformed buyers actually benefit the seller.

Our game is also a kind of informed principal problem. The unusual part about it is that the seller has more information about the mechanism he is using than about the product itself. The equilibrium we describe has a kind of 'punishment by beliefs' aspect that is based on this.

For example, we restrict the game to one in which the seller makes just a single take it or leave offer after processing messages. The seller might like to make additional offers if he makes an offer to the uninformed which is rejected. However, if the uninformed believe there is only a single offer, they can support the equilibrium by simply disappearing if they don't get the first offer. If the seller understands this he won't find it profitable to alter his mechanism to introduce these second offers.

Like all Bayesian games, equilibrium strategies must constitute a fixed point. In its simplest form, uninformed buyers' expectations about the relationship between the messages they send and the offers they get must coincide with the actual relationship once the seller best replies to those expectations.

To make it easier to find this fixed point we would like to use the revelation principle and all the well known properties of reduced form mechanisms. It isn't obvious how to do this. The usual composition of the outcome function and the strategies can be used to create something that looks like a direct mechanism. Yet not all direct mechanisms that look incentive compatible correspond to equilibrium outcomes because uninformed buyers can only use a restricted set of communication strategies.

In addition, equilibrium outcomes in which uninformed buyers learn how to participate in a direct mechanism can't correspond to equilibrium outcomes.

So we are going to define a special kind of direct mechanism that we can use to characterize an important class of equilibrium outcomes. In particular, it will allow us to characterize equilibrium in which the seller uses a symmetric mechanism and uninformed buyers use uninformative messages (babble).

First a device that helps for this problem. In what follows the notation m always means the number of uninformed buyers.

Definition. Let $v \in [0,1]^n$. A mod m permutation $\rho_m(v)$ of v moves the i^{th} element of v to the $i+1^{st}$ mod $n-m$ element of $\rho_m(v)$ for $i \leq n-m$.

Shifting the elements of v around we can find the sum of the probabilities assigned to each element in a relatively simple way.

Definition. Define the k^{th} *composition* of ρ_m in such a way that $\rho_m^0(v) = v$ and $\rho_m^k(v) = \rho_m\left(\rho_m^{k-1}(n)\right)$

For example if $v = \{1,2,3,4,5\}$, then $\rho_2^1(v) = \rho_2(v) = \{3,1,2,4,5\}$. Furthermore,

$$\rho_2^2(v) = \rho_2\left(\rho_m^1(v)\right) = \{2,3,1,4,5\}.$$

Now we have

Definition. A direct mechanism δ is a collection of functions

$$\{(\tilde{q}_m^\epsilon, \tilde{p}_m^\epsilon)\}_{m=0}^{n-1}, \{(\tilde{q}_m^\mu, \tilde{p}_m^\mu)\}_{m=1}^{n}$$

where $\tilde{q}_m^\epsilon, \tilde{q}_m^\mu : [0,1]^n \to [0,1]$, $\tilde{p}_m^\epsilon, \tilde{p}_m^\mu : [0,1]^n \to [0,1]$ and

- $(\tilde{q}^{\epsilon}_m, \tilde{p}^{\epsilon}_m)$ and $(\tilde{q}^{\mu}_m, \tilde{p}^{\mu}_m)$ are invariant with respect to their last m arguments;
- $\sum_{i=1}^{n-m} \tilde{q}^{\epsilon}_m \left(\rho^{i-1}_m(v)\right) + m\tilde{q}^{\mu}_m(v) \leq 1$ for all v and for all m.

This is actually a very standard way or representing symmetric direct mechanisms. The function $\tilde{q}^{\epsilon}_m(v)$ gives the probability with which an offer is made to the buyer with the first value in v given that the other $n-m-1$ informed buyers have the values $v_{-1} = \{v_2, \dots, v_{n-m}\}$. Correspondingly, the function $\tilde{q}^{\mu}_m(v)$ gives the probability with which an offer is made to an uninformed buyer given that $n-m$ informed buyers have values $\{v_1, \dots, v_{n-m}\}$.

The composition $\rho^{i}_m(v)$ moves the 1^{st} element of v to its $i+1^{st}$ position and the $i+1^{st}$ element of v to the first position, the sum $\sum_{i=1}^{n-m} \tilde{q}^{\epsilon}_m \left(\rho^{i-1}_m(v)\right)$ gives the probability that the offer is made to one of the first $n-m$ elements of v. Then

$$\sum_{i=1}^{n-m} \tilde{q}^{\epsilon}_m \left(\rho^{i-1}_m(v)\right) + m\tilde{q}^{\mu}_m(v) \leq 1 \tag{2.4}$$

ensures that when the informed buyers have values given by the first $n-m$ values in v, the probability with which the good is offered to one of them plus the probability that it is offered to one of the uninformed buyers is less than or equal to 1.

We can use this to build something that looks exactly like a traditional reduced form mechanism.

The probability with which an informed buyer whose value is w receives an offer when there are m uninformed is

$$Q^{\epsilon}_m(w) = \mathbb{E}_v \left\{\tilde{q}^{\epsilon}_m(v) | v_1 = w\right\}.$$

Similarly

$$P^{\epsilon}_m(w) = \mathbb{E}_v \left\{\tilde{q}^{\epsilon}_m(v)\tilde{p}^{\epsilon}_m(v) | v_1 = w\right\}$$

is the expected price the informed bidder with value w would pay.

For each $m = 0, \dots, n-1$, let $B(m; n-1, \alpha)$ be the probability that there are m uninformed buyers among the $n-1$ others. This probability is given by

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

Now by taking expectations over m we have the usual reduced form functions;

$$Q^{\epsilon}(w) = \sum_{m=0}^{n-1} B(m; n-1, \alpha) Q^{\epsilon}_m(w).$$

$$P^{\epsilon}(w) = \sum_{m=0}^{n-1} B(m; n-1, \alpha) P^{\epsilon}_m(w).$$

At this point, we inherit all the usual results from mechanism design in iid environments for each of the informed buyers. In particular, if the mechanism δ is incentive compatible *with respect to valuations*, the payoff to an informed buyer with value w can be written as

$$U^{\epsilon}(w) = \int_0^w Q^{\epsilon}(x)\, dx,$$

with $Q^{\epsilon}(\cdot)$ non-decreasing. The (interim) payoff to an uninformed bidder with valuation w is

$$U^{\mu}(w) = \sum_{m=1}^{n} B(m; n-1, \alpha)\mathbb{E}_v \left\{ \tilde{q}^{\mu}_{m+1}(v) \max\left[w - \tilde{p}^{\mu}_{m+1}(v), 0 \right] \right\}.$$

Definition. The mechanism δ is incentive compatible for informed buyers if $Q^{\epsilon}(\cdot)$ is non-decreasing and

$$U^{\epsilon}(w) \geq U^{\mu}(w)$$

for every w.

From standard arguments and properties of the binomial distribution, it is straightforward to show that the seller's revenue from informed buyers (again from any incentive compatible mechanism) is given by

$$n(1-\alpha)\int_0^1 Q^{\epsilon}(w)\,\phi(w) f(w) dw.$$

The seller's revenue from uninformed buyers is given by

$$\sum_{m=1}^{n} B(m; n, \alpha)\mathbb{E}_v \left\{ m\tilde{q}^{\mu}_m(v)\pi\left(\tilde{p}^{\mu}_m(v)\right)\right\}.$$

The seller's total revenue $\Pi(\delta)$ from a direct mechanism δ is the sum of the above two expressions.

Theorem 1. *Suppose that uninformed buyers send uninformative messages in some symmetric equilibrium of the game $\mathcal{G}(\alpha)$. Then there is a direct mechanism δ^* having the property that*

$$\mathbb{E}_{\gamma}\left[R\left(\gamma, \left(\sigma_i(\cdot,\cdot,\gamma)\right)_{i=1}^{n}\right) |\psi\right] = \Pi\left(\delta^*\right)$$

where δ^* satisfies $\Pi(\delta^*) \geq \Pi(\delta)$ for every incentive compatible direct mechanism δ. Conversely, any incentive compatible direct mechanism that maximizes $\Pi(\delta)$ can be used to construct an equilibrium in the Bayesian game $\mathcal{G}(\alpha)$.

The assumption that the equilibrium has uninformative messages from buyers is used because direct mechanisms as we have defined them do not allow an allocation to depend on the values of the uninformed.

3. Equal-priority auctions

Our main result is that the revenue maximizing uninformative equilibrium for distributions for which $\phi(v)$ is non-decreasing and $\pi(v)$ is concave called an "equal priority auction".

We'll show this in two parts. First we'll describe the set of equal priority auctions and describe the priority auction that gives the seller the highest expected revenue. Later we'll show how to verify this is best for the seller among all mechanisms.

An equal priority auction is fully characterized by four numbers, a 'reserve price' r, a price offer t, and the upper and lower bound v_+ and v_- of an interval of buyer types. We'll assume throughout that $r \leq t \leq v_- \leq v_+$.

In what follows, there is some message that is treated as if the buyer who sent that message is uninformed. Each of the informed buyers sends a bid. A realized profile of messages and bids will then have m messages saying uninformed, and $n-m$ bids.

The allocation in an equal priority auction is determined in the following way:

- If there are uninformed bidders and the highest bid received from the informed bidders is no larger than v_+, then the seller makes an offer t to each uninformed bidder (m of them) and each informed bidder who bid in the interval $[v_-, v_+]$ (k of them) with probability $\frac{1}{m+k}$.
- otherwise, the seller makes an offer to the informed buyer who made the highest bid. Let v' be the second highest bid by an informed buyer. The offer to the high bidder is

$$\begin{cases} v' & v' > v_+ \\ r & m = 0; v' < r \\ v' & m = 0; v' \in (r, v_-) \\ \frac{(v_- + (m+k)v_+)}{(m+k+1)} & \text{otherwise.} \end{cases}$$

where m and k are defined in bullet point (3) above.

Offer probabilities and incentive compatibility. These rules constitute an indirect mechanism support some kind of Bayesian equilibrium in bidding strategies. Our main theorem is going to say that conditional on uninformative messages from the uninformed bidders, the revenue maximizing mechanism is going to be a special kind equal priority auction. To see what that means, and to understand how to find the best one, one bit of notation is required. Suppose for the moment, potentially counterfactually, that informed buyers bid their true values and that the parameters of the auction satisfy $r \leq t \leq v_- \leq v_+$.

If all bidders bid their values, the mechanism above has some special properies. Since t is no larger than v_-, whenever an informed bidder receives an offer, he or she will want to accept it. Then using the allocation rule in the indirect mechanism, we can calculate the probability with which each type of informed buyer trades. This probability of trading function is non-decreasing, from which it is well known that we can devise a set of transfers that would support a Bayesian equilibrium in which each buyer would want to bid their value.

This probability of trade function is given by

$$Q^{\epsilon}(w) = \tag{3.1}$$

$$\begin{cases} 0 & \text{if} \quad w < r \\ (1-\alpha)^{n-1} F^{n-1}(w) & \text{if} \quad w \in [r, v_-) \\ \chi(v_-, v_+) & \text{if} \quad w \in [v_-, v_+] \\ \sum_{m=0}^{n-1} B(m; n-1, \alpha) F^{n-1-m}(w) & \text{if} \quad w > v_+, \end{cases}$$

where the function χ gives the probability that a buyer whose value is in the pooling interval $[v_-, v_+]$ receives an offer equal to t. This function is given by

$$\chi(v_-, v_+) = \sum_{m=0}^{n-1} B(m; n-1, \alpha) \sum_{k=0}^{n-1-m} B_k^{n-1-m}(v_-, v_+)/(m+k+1)$$

where

$$B_k^{n-1-m}(v_-, v_+) = \binom{n-1-m}{k} (F(v_+) - F(v_-))^k F^{n-1-m-k}(v_-).$$

The logic in $\chi(v_-, v_+)$ is that an informed bidder has the same chance of receiving an offer as any of the uninformed buyers and informed buyers whose valuations are in the interval $[v_-, v_+]$ as long as none of the other informed bidders has value above v_+.

Lemma 2. *For* $w \geq v_+$*,*

$$Q^\epsilon(w) = ((1-\alpha)F(w) + \alpha)^{n-1}$$

and for $w \in [v_-, v_+]$*,*

$$\chi(v_-, v_+) = \frac{((1-\alpha)F(v_+) + \alpha)^n - ((1-\alpha)F(v_-))^n}{n(\alpha + (1-\alpha)(F(v_+) - F(v_-)))}$$

Proof. Appendix. □

Then mimicking standard mechanism design we could define an expected payoff $U^\epsilon(w)$ to an informed buyer as follows:

$$U^\epsilon(w) = \int_0^w Q^\epsilon(x)dx.$$

Lemma 3. *There is a Bayesian equilibrium in truthful bidding strategies if*

$$\int_0^{v_-} Q^\epsilon(x)dx = ((v_- - t))\,\chi(v_-, v_+). \tag{3.2}$$

Proof. Two arguments are needed. The first is to show that the transfers defined by (3) are the ones that make truthful bidding incentive compatible. The second is to show that when t satifies condition (3.2) no informed buyer can improve his or her payoff by pretending to be uninformed.

For values of v below v_-, a buyer's payoff is resolved by auction, so the trading probability is strictly increase and the buyer's expected payoff is a convex function. The derivative of this function at v_- is

$$(1-\alpha)^{n-1}F^{n-1}(v_-).$$

The probability of trade for buyers whose value is in the interval $[v_-, v_+]$ is

$$\chi(v_-, v_+) =$$
$$\sum_{m=0}^{n-1} B(m; n-1, \alpha) \sum_{k=0}^{n-1-m} B_k^{n-1-m}(v_-, v_+)/(m+k+1) >$$
$$B(0, n-1, \alpha)\, B_0^{n-1}(v_-, v_+) =$$
$$(1-\alpha)^{n-1}F^{n-1}(v_-).$$

The payoff function for buyers whose values lie in the interval $[v_-, v_+]$ is given by

$$U^\epsilon(v_-) + \chi(v_-, v_+)(v - v_-)$$

so that if (3.2) holds, all buyers whose valuations are below v_+ will weakly prefer to bid their valuations instead of pretending to be unininformed.

Finally by (3) buyers whose values exceed v_+ are always engaged in an auction with a 'reserve price' in the case that no other bidders bid above v_+. The limit of this payoff as $v \downarrow v_+$ is given by

$$\sum_{m=0}^{n-1} B(m; n-1, \alpha) \sum_{k=0}^{n-1-m} B_k^{n-1-m}(v_-, v_+)\left(v_+ - \frac{v_- + v_+(m-k)}{m+k+1}\right) =$$
$$\chi(v_-, v_+)(v_+ - v_-)$$

which ensures that biuyers whose values are above v_+ also prefer to bid truthfully. □

The following figure shows the Bayesian equilibrium payoffs to bidders with various values in an incentive compatible equal priority auction (in other words, one satisfying (3.2)). The green line represents the payoff each buyer type achieves by acting as an uninformed bidder. The red line represents the payoff to informed bidders - except that the payoff to informed bidders who bid in the interval $[v_-, v_+]$ coincides with the green line.

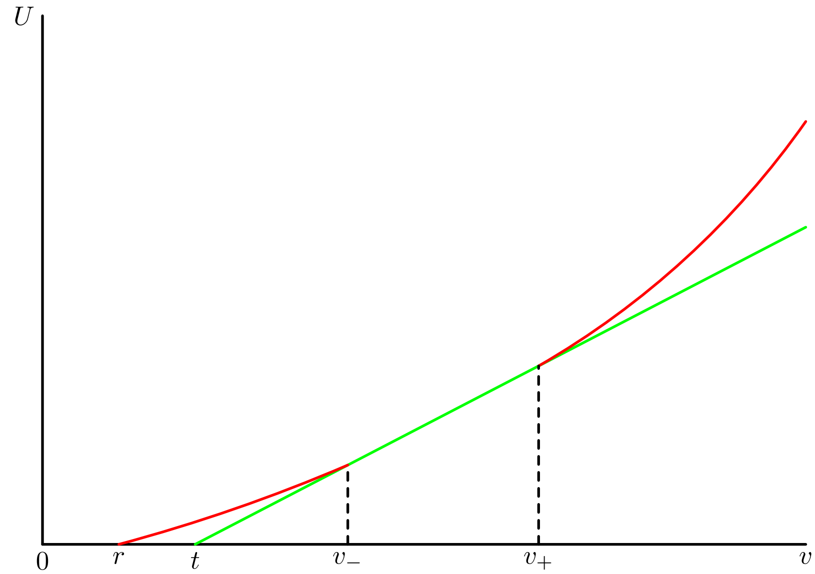

Seller revenues. This already looks like a direct mechanism, albeit one with very specific allocation rules. The seller's expected revenue from informed buyers is given by

$$n(1-\alpha)\int_r^1 Q^\epsilon(w)\phi(w)f(w)dw, \tag{3.3}$$

and the revenue from uninformed buyers is given by

$$\sum_{m=1}^{n} B(m;n,\alpha)\sum_{k=0}^{n-m} B_k^{n-m}(v_-,v_+)\frac{m}{m+k}\pi(t) = n\alpha\chi(v_-,v_+)\pi(t). \tag{3.4}$$

The revenue maximizing equal-priority auction $\{r, t, v_-, v_+\}$ maximizes the sum of (3.3) and (3.4) subject to

$$r \le t \le v_- \le v_+;$$

and

$$\int_r^{v_-}(1-\alpha)^{n-1}F^{n-1}(w)dw = \chi(v_-,v_+)(v_- - t).$$

Revenue Maximizing Equal Priority Auction.

Lemma 4. *Necessary conditions for an equal priority auction $\{r, t, v_-, v_+\}$ to be revenue maximizing among the class of equal priority auctions are*

$$(1-\alpha)(v_- - t)(\phi(v_+)-\phi(v_-))f(v_-) = \tag{3.5}$$

$$\alpha(\pi(t)-\phi(v_+)) + (1-\alpha)((\pi(v_-)-\pi(v_+)) - (F(v_+)-F(v_-))\phi(v_+));$$

$$\phi(r)f(r) + (\phi(v_+)-\phi(v_-))f(v_-) = 0; \tag{3.6}$$

$$\alpha\pi'(t) + (1-\alpha)(\phi(v_+) - \phi(v_-))f(v_-) = 0; \tag{3.7}$$

$$\int_r^{v_-} (1-\alpha)^{n-1}F^{n-1}(w)dw = \chi(v_-, v_+)(v_- - t). \tag{3.8}$$

The first three conditions are just interior solutions to the first order conditions, while the fourth condition is the incentive condition required by the restriction that an informed buyer whose valuation is v_- should be just indifferent between reporting his or her type and claiming to be uninformed. To establish that the optimal auction requires $r < t < v_- < v_+$ our proof (in the appendix) uses a variational argument.

One of the advantages of this parameterization is that it is simple to do comparative statics.

Corollary 5. *Let* $\overline{v}$ *satisfy* $\pi(r^*) = \phi(\overline{v})$*. The interval* $[v_-, v_+]$ *is degenerate only if*

$$\int_{r^*}^{\overline{v}} (1-\alpha)^{n-1}F^{n-1}(w)dw = \chi(\overline{v}, \overline{v})(\overline{v} - r^*).$$

The proof of this is quite simple so we'll do it here: Suppose the contrary. Then (i) condition (3.7) requires that $t = r^*$(in order that $\pi'(t) = 0$) and r^*satisfies $\phi(r^*) = 0$; (ii) (3.6) requires $r = r^*$; and (iii) (3.5) requires that $\alpha(\pi(t) - \phi(v_-)) = \alpha(\pi(t) - \phi(v_-)) = 0$. The result then follows from(3.8).

The implication is that the interval $[v_-, v_+]$ is typically non-degenerate. Briefly if the interval is degenerate, the seller can raise expected revenue by cutting the price t that he offers to the uninformed. The downside is that he loses revenue from the informed who are pooled together with the uninformed. A variational argument can be used more generally to show that the cutting the price offer to the uninformed has a first order impact on profits, while the loss from the informed is second order.

When all bidders are surely informed the revenue maximizing equal priority auction degenerates of an auction with reserve price r^*. As is readily seen in the first order conditions above, $t = r^*$ while$v_- = v_+ = \overline{v}$ as defined in Corollary 5.

In the opposite limit of $\alpha = 1$, and biddders are surely uninformed, auction converges to a fixed-price offer at the optimal price r^*. The lower-bound of the pooling interval becomes r^* to provide incentives for an unlikely informed buyer to participate in the equal-priority auction, while the upper-bound is again equal to $\overline{v}$.

3.1. **Equilibrium mechanisms.** We use Lagrangian relaxation to show that an equal-priority auction provides the seller the highest expected revenue among all direct mechanisms.

Recall that a direct mechanism δ consists of a series of functions $\{\tilde{q}^{\epsilon}_m, \tilde{p}^{\epsilon}_m\}_{m=0}^{n-1}$ and $\{\tilde{q}^{\mu}_m, \tilde{p}^{\mu}_m\}_{m=1}^{n}$. We first use the assumption that $\pi(\cdot)$ is strictly concave to simplify the optimal design problem.

Lemma 6. *If* $\pi(\cdot)$ *is strictly concave, then in any optimal direct mechanism,* $\tilde{p}^{\mu}_m(v)$ *is constant.*

Proof. Appendix. □

Then using Lemma 6, suppose the constant price offered to the uninformed is equal to t. As we showed above the revenue maximizing direct mechanism can be found by choosing a feasible mechanism δ that supports a trading probability for the uninformed Q^μ and a non-decreasing trading probability function $Q^\epsilon(\cdot)$ that maximizes

$$n(1-\alpha)\int_0^1 \{Q^\epsilon(w)\,\phi(w)f(w)dw\} + n\alpha Q^\mu \pi(t)$$

subject to

$$\int_0^w Q^\epsilon(\tilde{v})\,d\tilde{v} \geq Q^\mu \max[w-t,0]$$

for all w;

Let $\lambda(\cdot)$ be an arbitrary non-negative Lagrangian function from $[0,1]$ into $\mathbb{R}$. The relaxed problem is to maximize

$$n(1-\alpha)\int_0^1 \{Q^\epsilon(w)\,\phi(w)f(w)dw\} + \\ n\alpha Q^\mu \pi(t) + \\ \int_0^1 \lambda(w)\left\{\int_0^w Q^\epsilon(\tilde{v})\,d\tilde{v} - Q^\mu \max[w-t,0]\right\} dw$$

again choosing $Q^\epsilon(\cdot)$ and Q^μ to be feasible and non-decreasing. This problem has a different solution for every

It is well known that the solution to the relaxed problem is an upper bound on the solution to the full problem no matter what the Lagrangian function happens to be.[6]

The method of proof is to start (the hard part) by dreaming up a direct mechanism that you think might be a solution to the original problem. In our case this candidate is the equal priority auction. Then, instead of solving the relaxed problem, the method is to reverse engineer the Lagrangian function to try to find a function $\lambda(\cdot)$ such that the solution to the relaxed problem is an equal priority auction. Since the equal priority auction attains an upper bound for that Lagrangian function, it must be a solution to the full problem.

Now we just get on with it.

Denote the constant $\tilde{p}^\mu_m(v)$ given in Lemma 6 as t. To simplify notation, define

$$Q^\mu_m = \mathbb{E}_v\{\tilde{q}^\mu_m(v)\}$$

as the expected probability that an uninformed buyer receives the offer t when $n-m$ buyers are informed. Rewrite the seller's revenue from a direct mechanism δ as

$$n(1-\alpha)\int_0^1\left\{\sum_{m=0}^{n-1} B(m;n-1,\alpha)Q^\epsilon_m(w)\right\}\phi(w)f(w)dw + \sum_{m=1}^{n} B(m;n,\alpha)Q^\mu_m m\pi(t).$$

The incentive condition for informed buyers not to pretend to be uninformed becomes

$$U^\epsilon(w) \geq U^\mu(w) = \sum_{m=0}^{n-1} B(m;n-1,\alpha)Q^\mu_{m+1}\max[w-t,0].$$

[6]This even has a wikipedia entry under the term 'Lagrangian Relaxation'. The argument is intuitive. The solution to the original problem is feasible, so the integral in the payoff to the relaxed problem is non-negative no matter what the Lagrangian function λ happens to be. Yet the solution to the original problem will not generally be a solution to the relaxed problem, thus the upper bound.

The problem of optimal direct mechanism is to choose functions $\{\tilde{q}^{\epsilon}_m\}_{m=0}^{n-1}$ and $\{\tilde{q}^{\mu}_m\}_{m=1}^{n}$, and t to maximize the revenue subject to the above incentive condition, in addition to the incentive constraint with respect to valuations that $Q^{\epsilon}(\cdot)$ is weakly increasing, and the feasibility constraint (2.4).

Let $\lambda(\cdot)$ be a non-negative function from $[0,1]$ into $\mathbb{R}$. The relaxed problem is to maximize

$$n(1-\alpha)\int_0^1\left\{\sum_{m=0}^{n-1}B(m;n-1,\alpha)Q^{\epsilon}_m(w)\right\}\phi(w)f(w)dw+\sum_{m=1}^{n}B(m;n,\alpha)Q^{\mu}_m m\pi(t)+$$

$$\int_0^1\left\{\int_0^w\sum_{m=0}^{n-1}B(m;n-1,\alpha)Q^{\epsilon}_m(x)dx-\sum_{m=0}^{n-1}B(m;n-1,\alpha)Q^{\mu}_{m+1}\max[w-t,0]\right\}\lambda(w)dw,$$

subject to the incentive constraint with respect to valuations and the feasibility constraint. It is well known that the solution to the relaxed problem is an upper bound on the solution to the full problem no matter what the Lagrangian function.[7] The method of proof is to try to find a function $\lambda(\cdot)$ such that the solution to the relaxed problem is an equal priority auction. Since the equal priority auction yields an upper bound on the seller's payoff in the full problem, and since it satisfies all the constraints in the full problem, it must be a solution to the full problem.

Theorem 7. *Suppose that $\pi(\cdot)$ is strictly concave. Then, a revenue-maximizing equal priority auction is an optimal direct mechanism.*

Proof. Appendix. □

4. Variations

We have assumed that the objective of the seller is to make a single take it or leave it offer. If this offer is rejected, which it will sometimes be if it is made to an uninformed bidder, the game ends without trade. One question is how this might change if the seller could follow up a rejection by making an offer to one of the other bidders.

One natural way to proceed, might be to assume that uninformed buyers are unsure how many offers are made and how long it they should wait for one. They might then choose to give up in a manner that depends on their waiting costs. We won't persue this here because it is obviously far more complex.

Yet we can make one observation. There will be at least one equilibrium in such a game that resolves exactly to our equal priority auction. In this equilibrium we just imagine that all buyer, informed and uninformed believe that only a single offer will be made at some arbitrary point in time. If they don't receive the offer at that point, they just don't expect to get one at all.

Under those conditions, the seller can do no better than the equal priority auction because once an offer is rejected there won't be any buyers around to consider another one.

It is unclear whether alternative equilibrium with multiple offers might increase the seller's revenue. As always, multiple offer equilibrium improve the payoff to an informed buyer who wants to pretend to be uninformed. We defer this to future

[7]The solution to the original problem is feasible, so the integral in the payoff to the relaxed problem is non-negative. In turn, the solution to the original problem gives a lower payoff in the relaxed problem than the solution to the relaxed problem itself.

research since it is not clear at this point what is the best way to generalize to multiple offers.

Another sort of 'partial equilibrium' possibility is to assume that buyers wait around forever for offers. In this case the seller can do well by making a very high offer to each buyer in turn until all have rejected it. Then the seller could lower the offer slightly and continue to do this until some buyer accepts. This implements something that looks like a descending price auction.

The assumption that uninformed buyers will wait forever seems inconsistent at least with our view of why there are uninformed buyers in the first place - they have alternatives.

Notice also that this is not the same as implementing a descending clock auction. Uninformed buyers won't understand how to bid in a descending clock auction for the same reason we described above. If they expected a descending clock auction and stopped optimally, the seller would add a password then change the rules to make them a counter offer.

Finally we end with another question which we cannot fully answer. There are equilibrium outcomes in which uninformed buyers communicate with the seller. The equilibrium we can characterize has uninformed buyer with very low value to send a message to the seller saying that they are not interested in trading. The seller will commit not to make them an offer if he cannot find an alternative among the other buyers. This commitment will ensure that informed bidders have no incentive to pretend to be uninformed. Such equilibrium involve equal priority auctions just like the one we described above, except that the fixed price offer will more likely be accepted. As a consequence, seller revenue are higher in these equilibria.

This result is straightforward but tedious to show, so we leave it out.

More important, we have not been able to rule out more complicated equilibrium outcomes in which uninformed buyers receive different offers depending on the messages they send. So far we have made little progress understanding these outcomes if they exist at all.

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. We show that, assuming uninformed buyers don't communicate any useful information, the seller's revenue optimal equilibrium can be implemented with an equal priority auction. This mechanism is new as far as we know. It lies nicely between the extremes of pure auction, which is best when the seller is sure everyone is informed, and a simple take it or leave price offer to a buyer chosen randomly.

We don't include the proof here, but payoffs associated with the equal priority auction converge to payoffs in the straight auction when α goes to zero, and to the payoffs associated with the random take it or leave it offer as α goes to one.

One of the nice advantages of the equal priority auction is that it is parametric - all equal priority auctions can be described using only 4 parameters, which makes is easy to show existence. The parameters all lie in a compact set, and the payoff functions are integrals which depend continuously on the parameters.

The parametric representation makes it possible to do computations, and in principal, do empirical work. As we mentioned above, one of the implications of the the equal priority auctions is that the distribution of bids in the auction will be endogenous. In particular, it will be bi-modal with high and low bids while intermediate value bidders trade at a fixed price. This is something like what happens on eBay, though eBay auctions differ in many ways from what we have modeled here.

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 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 valuation, 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.

6. Appendix: Proofs

Proof of Lemma 1

Proof. Recall that stragies in $\mathcal{G}(\alpha)$ are given by some mixture ψ for the seller, and a collection of strategy rules $(\sigma_i(\cdot,\epsilon,\gamma), \sigma_i(\cdot,\mu))_{i=1}^N$ for informed and uninformed buyers respectively. These formulas can be reduced using the two assumptions in the theorem - symmetry and uninformative messages. Symmetry means that all informed buyers use the same strategy, call it $\sigma_\epsilon(\cdot,\cdot)$. Since uninformed buyers send uninformative messages, their messages won't affect any outcomes apart from possible randomization which can be implemented without these message. So we can ignore the message strategy of the uninformed entirely.

Symmetry also means all informed buyers that send the same message are treated the same way. Then, of course, it also follows that all uninformed buyers are treated the same way.

Using these ideas we can start to reduce the payoff functions. Once the mechanism γ adopted by the seller is realized, the payoff for informed bidders who send the message b_i can be written

$$\mathbb{E}_{v_{-i},\tau_{-i}} \int \cdots \int \{q_i(b_i, b_{-i}) \cdot \max[(w - p_i(b_i, b_{-i})), 0]\} \, d\sigma_{-i}(v_{-i}, \tau_{-i}, \gamma).$$

The seller's payoff must be the same for each realization γ in the support of his randomization. Furthermore uninformed buyers' strategies are independent of the realization of the seller's randomization. For that reason, we can characterize the seller's expected revenue in equilibrium by picking any γ in the support of ψ and calculating revenue for that realization.

Similarly for an informed buyer, his or her message is conditional on γ and is the same for any message b_i in the support of $\sigma(\cdot, \epsilon, \gamma)$. Then integrating out the randomizations of the other buyers, and using the fact that the payoff depends only on the number of uninformed buyers, we could write buyer payoffs for this realized mechanism γ as

$$\hat{U}(w, \epsilon) = \mathbb{E}_{v,m} \{q_\epsilon(\sigma_\epsilon^m(v)) \cdot \max[w - \hat{p}_m(\sigma_\epsilon(v)), 0] \, | b_1 = b_i\},$$

where

$$q_\epsilon(\sigma_\epsilon^m(v)) =$$

$$\int \cdots \int q_1\left(\sigma_1(v_1, \epsilon, \gamma), \ldots, \sigma_{n-m}(v_{n-m}, \epsilon, \gamma), \tilde{b}_{n-m+1}, \ldots, \tilde{b}_n\right) d\sigma_{n-m+1}(v_{n-m+1}, \mu) \ldots d\sigma_n(v_n, \mu).$$

In a symmetric equilibrium, $\sigma_i(\cdot, \epsilon, \gamma) = \sigma_{n-m}(\cdot, \epsilon, \gamma)$ for all $i \leq n - m$, while $\sigma_k(v_k, \mu)$ is independent of γ and v_k.

It is then easy to see that these formulas can be reduced to the direct mechanisms described in Definition 2.1. The fact that this mechanism derived from γ maximizes $\Pi(\delta)$ follows from the fact that it is the most profitable indirect mechanism.

The reverse direction of the argument in Theorem 1 follows from the token argument - the seller's description of his mechanism includes a randomly chosen token that must be submitted with the bid of an informed bidder. □

Proof of Lemma 2.

Proof. We have

$$\begin{aligned}
\chi(v_-, v_+) &= \sum_{l=0}^{n-1} \binom{n-1}{l} ((1-\alpha)F(v_-))^{n-1-l} \frac{1}{l+1} \sum_{k=0}^{l} \binom{l}{k} ((1-\alpha)(F(v_+) - F(v_-)))^k \alpha^{l-k} \\
&= \sum_{l=0}^{n-1} \binom{n-1}{l} ((1-\alpha)F(v_-))^{n-1-l} \frac{1}{l+1} ((1-\alpha)(F(v_+) - F(v_-)) + \alpha)^l \\
&= \frac{((1-\alpha)F(v_+) + \alpha)^n - ((1-\alpha)F(v_-))^n}{n((1-\alpha)(F(v_+) - F(v_-)) + \alpha)}. \qquad (6.1)
\end{aligned}$$

□

Proof of Lemma 4.

Proof. The critical bidding condition (3.8) binds at any optimal equal-priority auction.

Define

$$\Delta U(v_-) = \int_r^{v_-} (1-\alpha)^{n-1} F^{n-1}(w)dw - \chi(v_-, v_+)(v_- - t).$$

We have:

$$\frac{\partial \Delta U(v_-)}{\partial r} = -(1-\alpha)^{n-1} F^{n-1}(r);$$

$$\frac{\partial \Delta U(v_-)}{\partial t} = \chi(v_-, v_+);$$

$$\frac{\partial \Delta U(v_-)}{\partial v_-} = (1-\alpha)^{n-1} F^{n-1}(v_-) - \chi(v_-, v_+) - \frac{\partial \chi(v_-, v_+)}{\partial v_-}(v_- - t);$$

and

$$\frac{\partial \Delta U(v_-)}{\partial v_+} = -\frac{\partial \chi(v_-, v_+)}{\partial v_+}(v_- - t).$$

Next, letting R be the revenue from the equal-priority auction, given by the sum of (3.3) and (3.4), we have

$$\frac{\partial R}{\partial r} = -n(1-\alpha)^n F^{n-1}(r)\phi(r)f(r);$$

$$\frac{\partial R}{\partial t} = n\alpha\chi(v_-, v_+)\pi'(t);$$

$$\frac{\partial R}{\partial v_-} = n(1-\alpha)((1-\alpha)^{n-1}F^{n-1}(v_-) - \chi(v_-, v_+))\phi(v_-)f(v_-)$$
$$+ n((1-\alpha)(\pi(v_-) - \pi(v_+)) + \alpha\pi(t))\frac{\partial \chi(v_-, v_+)}{\partial v_-};$$

and

$$\frac{\partial R}{\partial v_+} = n(1-\alpha)\left(\chi(v_-, v_+) - ((1-\alpha)F(v_+) + \alpha)^{n-1}\right)\phi(v_+)f(v_+)$$
$$+ n((1-\alpha)(\pi(v_-) - \pi(v_+)) + \alpha\pi(t))\frac{\partial \chi(v_-, v_+)}{\partial v_+}.$$

It is straightforward to use Lemma 2 to verify that

$$\frac{\partial \chi(v_-, v_+)}{\partial v_-} = \frac{(1-\alpha)f(v_-)}{(1-\alpha)(F(v_+) - F(v_-)) + \alpha}\left(\chi(v_-, v_+) - ((1-\alpha)F(v_-))^{n-1}\right);$$

and

$$\frac{\partial \chi(v_-, v_+)}{\partial v_+} = \frac{(1-\alpha)f(v_+)}{(1-\alpha)(F(v_+) - F(v_-)) + \alpha}\left(((1-\alpha)F(v_+) + \alpha)^{n-1} - \chi(v_-, v_+)\right).$$

Now we are ready to derive the three optimality conditions stated in the lemma. First,

$$\frac{\partial R/\partial v_-}{\partial \Delta U(v_-)/\partial v_-} = \frac{\partial R/\partial v_+}{\partial \Delta U(v_+)/\partial v_+}.$$

Using the expressions for $\chi(v_-, v_+)$, $\partial\chi(v_-, v_+)/\partial v_-$ and $\partial\chi(v_-, v_+)/\partial v_+$, straightforward algebra lead us to the first-order condition (3.5) for an optimal equal-priority auction with respect to v_- and v_+. Note that (3.5) implies that

$$\frac{\partial R/\partial v_+}{\partial \Delta U(v_-)/\partial v_+} = -n(1-\alpha)(\phi(v_+) - \phi(v_-))f(v_-).$$

Second, from

$$\frac{\partial R/\partial t}{\partial \Delta U(v_-)/\partial t} = \frac{\partial R/\partial v_+}{\partial \Delta U(v_+)/\partial v_+},$$

we have the first order condition (3.7) with respect to t and v_+.

Third,

$$\frac{\partial R/\partial r}{\partial \Delta U(v_-)/\partial r} \geq \frac{\partial R/\partial v_+}{\partial \Delta U(v_+)/\partial v_+},$$

and $r \geq 0$, with complementary slackness. This gives the first-order condition

$$-\phi(r)f(r) \leq (\phi(v_+) - \phi(v_-))f(v_-),$$

and $r \geq 0$, with complementary slackness. Note that $-\phi(0)f(0) = 1$. By (3.5), we have $\phi(v_+) < \pi(t) < \pi(r^*) < r^*$, while $v_- > t > r^*$. Thus,

$$(\phi(v_+) - \phi(v_-))f(v_-) = (\phi(v_+) - v_-)f(v_-) + 1 - F(v_-) < 1.$$

It follows that the optimal r is interior and so (3.6) holds. □

Proof of Lemma 6.

Proof. Fix a direct mechanism $\{\tilde{q}^\epsilon_m, \tilde{p}^\epsilon_m\}_{m=0}^{n-1}$ and $\{\tilde{q}^\mu_m, \tilde{p}^\mu_m\}_{m=1}^{n}$. Define $t \in [0,1]$ to be the expected offer to uninformed buyers, given by

$$\sum_{m=0}^{n-1} B(m; n-1, \alpha)\mathbb{E}_v\{\tilde{q}^\mu_{m+1}(v)(t - \tilde{p}^\mu_{m+1}(v))\} = 0.$$

Since $\tilde{p}^\mu_m(v) \in [0,1]$ for all v,

$$\sum_{m=0}^{n-1} B(m; n-1, \alpha)\mathbb{E}_v\{\tilde{q}^\mu_{m+1}(v)\} \max[w - t, 0]$$

$$= \sum_{m=0}^{n-1} B(m; n-1, \alpha)\mathbb{E}_v \left\{\tilde{q}^\mu_{m+1}(v) \max[w - \tilde{p}^\mu_{m+1}(v), 0]\right\}$$

for all $w \leq \min \tilde{p}^\mu_m(v)$ and for all $w \geq \max \tilde{p}^\mu_m(v)$. Since $U^\mu(w)$ is convex in w, we have

$$U^\mu(w) \geq \sum_{m=0}^{n-1} B(m; n-1, \alpha)\mathbb{E}_v\{\tilde{q}^\mu_{m+1}(v)\} \max[w - t, 0]$$

for all w. Thus, replacing each all functions $\{\tilde{p}^\mu_m(\cdot)\}_{m=1}^n$ with t reduces the deviation payoff of an informed buyer from pretending to be uninformed. The seller's revenue from uninformed buyers is

$$\sum_{m=1}^{n} B(m; n, \alpha)\mathbb{E}_v\left\{m\tilde{q}^\mu_m(v)\pi\left(\tilde{p}^\mu_m(v)\right)\right\} = n\alpha \sum_{m=0}^{n-1} B(m; n-1, \alpha)\mathbb{E}_v\left\{\tilde{q}^\mu_{m+1}(v)\pi\left(\tilde{p}^\mu_{m+1}(v)\right)\right\}.$$

The lemma then follows from the strict concavity of $\pi(\cdot)$. □

Proof of Theorem 7.

Proof. Suppose that $\{r, t, v_-, v_+\}$ is a revenue-maximizing equal priority auction. By Lemma 4, the first order conditions (3.5)-(3.8) are satisfied. We construct a non-negatively valued multiplier function $\lambda(w)$ for all $w \in [0,1]$ such that $\{r, v_-, v_+, t\}$ solves the Lagrangian relaxation.

Using integration by parts, we can rewrite the Lagrangian as

$$\sum_{m=0}^{n-1} B(m; n-1, \alpha) \int_0^1 \left\{ n(1-\alpha)\phi(w)f(w) + \int_w^1 \lambda(x)dx \right\} Q_m^{\epsilon}(w)dw$$
$$+ \sum_{m=0}^{n-1} B(m; n-1, \alpha) \left(n\alpha\pi(t) - \int_0^1 \lambda(w)\max\{(w-t), 0\}dw \right) Q_{m+1}^{\mu}.$$

For each $w \in [0,1]$, denote

$$K^{\epsilon}(w) = n(1-\alpha)\phi(w) + \int_w^1 \lambda(x)dx / f(w);$$

$$K^{\mu} = n\alpha\pi(t) - \int_0^1 \lambda(x)\max[x-t, 0]dx.$$

We can then further rewrite the Lagrangian as

$$(1-\alpha)^{n-1}\int_0^1 K^{\epsilon}(w)Q_0^{\epsilon}(w)f(w)dw + \alpha^{n-1}K^{\mu}\tilde{q}_n^{\mu}$$
$$+ \sum_{m=1}^{n-1}\left(\int_0^1 B(m; n-1, \alpha)K^{\epsilon}(w)Q_m^{\epsilon}(w)f(w)dw + B(m-1; n-1, \alpha)K^{\mu}Q_m^{\mu} \right),$$

where $Q_0^{\epsilon}(w)$ is the probability that an informed buyer with valuation w gets the good when all buyers are informed, and $\tilde{q}_n^{\mu}$ is the probability that each uninformed buyer gets the good when all buyers are uninformed.

Now we construct $\lambda(\cdot)$ as follows. Let $\lambda(w) = 0$ for all $w \notin [v_-, v_+]$, and let

$$\lambda(w) = n(1-\alpha)\frac{d}{dw}(f(w)(\phi(w) - \phi(v_+))) = n(1-\alpha)(2f(w) + f'(w)(w - \phi(v_+)))$$

for all $w \in (v_-, v_+)$, with $\lambda(v_-)$ and $\lambda(v_+)$ given by the corresponding limit from above and from below. Since by assumption $\pi(\cdot)$ is strictly concave, $f(w)\phi(w)$ is strictly increasing in w, and thus $\lambda(w) > 0$ at any $w \in [v_-, v_+]$ such that $f'(w) \leq 0$. By (3.5) we have $\phi(v_+) < \pi(t) < \pi(r^*) < r^*$. Since $w \geq v_- > t > r^*$, we have $\lambda(w) > 0$ at any $w \in [v_-, v_+]$ such that $f'(w) > 0$. Thus, $\lambda(w)$ as constructed is non-negative for any w.

For any $w \in [v_-, v_+]$, by construction $K^{\epsilon}(w)$ is constant, because

$$\int_w^1 \lambda(x)dx = n(1-\alpha)f(w)(\phi(v_+) - \phi(w)).$$

Further, using integration by parts, we have

$$\int_0^1 \lambda(w)\max\{(w-t), 0\}dw =$$

$$-\int_{v_-}^{v_+} (w-t)d\left(\int_w^1 \lambda(x)dx\right) =$$

$$n(1-\alpha)\left((v_- - t)f(v_-)(\phi(v_+)-\phi(v_-)) + \int_{v_-}^{v_+} f(w)(\phi(v_+)-\phi(w))dw\right) =$$

$$n(1-\alpha)\left((v_- - t)f(v_-)(\phi(v_+)-\phi(v_-)) + \phi(v_+)(F(v_+)-F(v_-)) - (\pi(v_-)-\pi(v_+))\right).$$

It then follows from (3.5) that

$$\frac{B(m;n-1,\alpha)}{n-m}K^{\epsilon}(w) = \frac{B(m-1;n-1,\alpha)}{m}K^{\mu}.$$

For all $w > v_+$, since $\pi(\cdot)$ is strictly concave,

$$K^{\epsilon}(w) = n(1-\alpha)\phi(w) > n(1-\alpha)\phi(v_+) = K^{\epsilon}(v_+),$$

and so

$$\frac{B(m;n-1,\alpha)}{n-m}K^{\epsilon}(w) > \frac{B(m-1;n-1,\alpha)}{m}K^{\mu}.$$

For all $w < v_-$,

$$\begin{aligned}K^{\epsilon}(w) &= n(1-\alpha)\phi(w) + \int_{v_-}^{v_+}\lambda(x)dx/f(w)\\ &= n(1-\alpha)(\phi(w) + f(v_-)(\phi(v_+)-\phi(v_-))/f(w)).\end{aligned}$$

We claim that

$$\phi(w) + \frac{f(v_-)(\phi(v_+)-\phi(v_-))}{f(w)} < \phi(v_+)$$

for all $w < v_-$, and thus $K^{\epsilon}(w) < K^{\epsilon}(v_+)$ and thus

$$\frac{B(m;n-1,\alpha)}{n-m}K^{\epsilon}(v) \le \frac{B(m-1;n-1,\alpha)}{m}K^{\mu}.$$

To establish the claim, recall that in showing that the constructed multiplier function $\lambda(v)$ is positive for $v \in [v_-, v_+]$, we have proved that $f(w)(\phi(w)-\phi(v_+))$ is strictly increasing in w for all $w \ge \phi(v_+)$. This immediately implies that the claim holds for any $w \in [\phi(v_+), v_-)$. For $w < \phi(v_+)$, we have

$$f(w)(\phi(w)-\phi(v_+)) = f(w)(w-\phi(v_+)) - (1-F(w)) < -(1-F(w)) < -(1-F(r^*)),$$

where the last inequality follows because $\phi(v_+) < \pi(t) < \pi(r^*) < r^*$, while

$$f(v_-)(\phi(v_+)-\phi(v_-)) < f(r^*)\phi(v_+) < f(r^*)r^*,$$

where the first equality comes from $f(w)(\phi(v)-\phi(v_+))$ being strictly increasing in w for all $w \ge \phi(v_+)$. The claim then follows from the definition of r^*.

Now we show that the equal-priority auction $\{r, v_-, v_+, t\}$ with r satisfying (3.6) is a point-wise maximizer of the Lagrangian. For this purpose, we disaggregate $Q_m(w)$ and write the Lagrangian as

$$(1-\alpha)^{n-1}\int_0^1 K^{\epsilon}(w)Q_0^{\epsilon}(w)f(w)dw + \alpha^{n-1}K^{\mu}q_n^{\mu} +$$

$$\sum_{m=1}^{n-1}\mathbb{E}_v\left\{\frac{B(m;n-1,\alpha)}{n-m}\sum_{i=1}^{n-m}K^{\epsilon}(v_i)\tilde{q}_m^{\epsilon}(\rho_m^{i-1}(v)) + B(m-1;n-1,\alpha)K^{\mu}\tilde{q}_m^{\mu}(v)\right\}.$$

Fix any realized number m of uninformed buyers such that $1 \le m \le n-1$, and consider the last term in the above objective function. Suppose that for some realized valuation profile v we have $v_i > v_+$ for some $i = 1,\ldots,n-m$, but $\tilde{q}_m^{\mu}(v) > 0$.

By (2.4), we can decrease $\tilde{q}^{\mu}_m(v)$ marginally by $d\tilde{q}^{\mu}_m(v) > 0$ and increase $\tilde{q}^{\epsilon}_m(\rho^{i-1}_m(v))$ by $md\tilde{q}^{\mu}_m(v)$. Since

$$\frac{m}{n-m}B(m; n-1, \alpha)K^{\epsilon}(v_i) > B(m-1; n-1, \alpha)K^{\mu},$$

the effect on the seller's revenue is strictly positive. Therefore, $\tilde{q}^{\mu}_m(v) = 0$ for any v such that $v_i > v_+$ for some $i = 1, \ldots, n-m$. Further, since $K^{\epsilon}(w)$ is strictly increasing for $w > v_+$, we have $\tilde{q}^{\epsilon}_m(\rho^{i-1}_m(v)) = 1$ for $v_i = \max[v_1, \ldots, v_{n-m}]$. By similar arguments, we can show that there is a maximizer of the Lagrangian such that $\tilde{q}^{\epsilon}_m(\rho^{i-1}_m(v)) = 0$ whenever $v_i < v_-$, and $\tilde{q}^{\epsilon}_m(\rho^{i-1}_m(v)) = \tilde{q}^{\mu}_m(v)$ if $v_i \in [v_-, v_+]$.

For $m = 0$ and the first term in the Lagrangian, the strict concavity of $\pi(\cdot)$ implies $K^{\epsilon}(w)$ for $w < v_-$ crosses 0 at most once and only from below. Thus, for r that satisfies (3.6), it is point-wise maximizing to set $\tilde{q}^{\epsilon}_0(\rho^{i-1}_0(v)) = 1$ if $v_i = \max[v_1, \ldots, v_n]$ and $v_i > v_+$, or if $v_i = \max[v_1, \ldots, v_n]$ and $v_i \in [r, v_-)$; set $\tilde{q}^{\epsilon}_0(\rho^{i-1}_0(v)) = 1/k$ if $v_i \in [v_-, v_+]$, $\max[v_1, \ldots, v_n] \in [v_-, v_+]$ and $\#\{j : v_j \in [v_-, v_+]\} = k$; and set $\tilde{q}^{\epsilon}_0(\rho^{i-1}_0(v)) = 0$ otherwise.

For $m = n$ and the second term in the Lagrangian, we have $q^{\mu}_n = 1/n$ because $K^{\mu} > 0$. □

References

Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020. doi: 10.3982/ECTA15925. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15925.

Elchanan Ben-Porath, Eddie Dekel, and Barton L. Lipman. Optimal allocation with costly verification. *The American Economic Review*, 104(12):3779–3813, 2014. ISSN 00028282. URL http://www.jstor.org/stable/43495357.

Helmut Bester and Roland Strausz. Contracting with imperfect commitment and the revelation principle: The single agent case. *Econometrica*, 69(4):1077–1098, 2001. doi: 10.1111/1468-0262.00231. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00231.

G. Butters. Equilibrium distributions of sales and advertizing prices. *Review of Economic Studies*, 44:465–491, 1977.

Peter R. Dickson and Alan G. Sawyer. The price knowledge and search of supermarket shoppers. *Journal of Marketing*, 54(3):251–334, July 1990.

K. Hendricks and T. Wiseman. How to sell (or procure) in a sequential auction market*. 2020.

A. Kolotilin, H. Li, and W. Li. Optimal limited authority for principal. *Journal of Economic Theory*, 148(6):2344â2382, 2013.

Q. Liu, K. Mierendorff, and X. Shi. Auctions with limited commitment. Columbia University working paper, 2014.

Preston McAfee. Mechanism design by competing sellers. *Econometrica*, 61(6): 1281–1312, November 1993.

Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6: 55–73, 1981.

V. Skreta. Optimal auction design under non-commitment. *Journal of Economic Theory*, 159:854â–890, 2015.

D. Stahl. Oligopolistic pricing and advertising. *Journal of Economic Theory*, 64 (1):162–â177, 1994.

H. Varian. A model of sales. *American Economic Review*, 70(4):651–â659, 1980.

Gabor Virag. Competing auctions: Finite markets and convergence. *Theoretical Economics*, 5:241–274, 2010.