Unobserved Mechanism Design: Equal Priority Auctions

Li, Hao and Michael Peters*

Vancouver School of Economics

University of British Columbia

September 9, 2023

Abstract

We study the impact for mechanism design of the possibility that some participants are unobservant of the rules associated with the trading mechanism but are otherwise rational. Since "deviations" by the mechanism designer are not observed by these participants the nature of the "equilibrium" of the design game changes, as do equilibrium mechanisms. We study the symmetric, regular case of the independent private value auction environment, and show how to characterize an interesting equilibrium outcome for the game by optimizing over reduced form direct mechanisms. This gives rise to a surprisingly simple mechanism that we call an *equal priority auction.* Observant bidders with intermediate valuations receive offers with the same probability as unobservant bidders, even though observant buyers will accept the offers for sure, while unobservant bidders might not.

*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. Gabriel Carroll, Vitor Farinha Luz, Alessandro Pavan, Sergei Severinov and Gabor Virag provided helpful comments and suggestions.

1 Introduction

There is an acronym that floats around the internet - TL;DR - that explains why no one reads your email messages. It means "too long, didn't read." The long translation we adapt in this paper is "there is no point in reading your message, because I already know everything that's relevant to me." We refer to this as "rationally unobservant." It sounds like an oxymoron, but is meant to capture the idea that one can be unobservant and rational at the same time.

The message of this paper is that this kind of behavior can impact trading mechanisms. We aren't the first to notice that traders are sometimes unobservant. The marketing literature has documented this behavior of buyers when they make purchase decisions. The simplest trading mechanism of all is a price commitment. Dickson and Sawyer [1990] asked buyers in supermarkets about their price knowledge as they were shopping. Even when the item being placed in their basket had been specially marked down and heavily advertised, 25% of consumers did not even realize the good was on special. Marketing has a problem when prices can't influence buyer behavior because buyers may be unobservant.

We are interested in more than prices; we want to know how trading mechanisms are impacted by buyers who are possibly unobservant but are always rational. We consider what is probably the best understood trading problem of all - the independent private value auction. Having rationally unobservant buyers who don't observe the seller's commitment turns the constrained optimization problem of mechanism design into a game of imperfect information where deviations by the seller may not be observed by some buyers.

We consider the symmetric, regular case of the independent private value auction environment in this paper. Our formal analysis is based on two arguments. The first is that in an environment with unobservant bidders, 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 unobservant bidders. To see why, suppose the seller wants to use a second price auction with optimal reserve. This means that observant bidders read the auction rules, as they might on eBay, then realize they should bid their valuations. Unobservant bidders don't read the rules, so they only *anticipate* a second price auction. Acting on their expectations, they also bid their valuations.

What makes this break down is the fact that if the seller changes the auction rules, the unobservant won't realize it, and will continue to bid their valuations no matter what the seller does. A simple deviation can extract the surplus of the unobservant. For example, the seller can ask 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 bidder 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 treats bids with no code attached as if it were a first price auction. In other words, if the highest bid is submitted by an unobservant bidder, the seller will commit to make them an offer equal to their bid, instead of offering them the second highest bid. This breaks the equilibrium, because it should be expected by rational unobservant bidders.

The second argument involves how the seller should respond. The seller will want to sell to the unobservant bidders when observant bidders have low valuations. So the natural idea would be to have an auction, then if bids are too low, make an offer to the unobservant. But observant bidders don't have to bid. They can pretend to be unobservant. Since they are observant, they know when the seller will make an offer to the unobservant and what that offer will be. To prevent observant bidders from pretending to be unobservant, the seller has to keep the offer to unobservant bidders higher than she would like it to be.

The seller then faces a trade off - keep the offer high and fully separate the observant from the unobservant, or lower the take it or leave it offer and allow some of the observant bidders to pool with the unobservant. We show that the latter is what happens in equilibrium, which is where the *equal priority* phrase comes from in our title. Intuitively, pooling sacrifices auction revenue from observant bidders, but this is minuscule when the seller is fully separating them from unobservant bidders. The revenue gain from lowering the offer to unobservant bidders then makes pooling profitable.

We show that under a plausible restriction - unobservant buyers convey no information to sellers - the equilibrium outcome is characterized by an *optimal* equal priority auction. Pooling happens at intermediate valuations, because it is too costly to include in the equal priority pool observant bidders with high valuations, and unnecessary to include those with low valuations. The auction treats observant bidders with intermediate valuations in exactly the same way as unobservant bidders. When the auction attempts to trade with them, it makes a take-it-or-leave-it price offer that is independent of any messages they may have sent. When bidders have very high or very low valuations, the seller treats messages as bids. If the seller decides to sell to one of these bidders, she will make an offer equal to what can be thought of as the second highest bid she has received - much as she would in a standard auction.

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 equilibrium mechanism with rationally unobservant bidders can be found by solving a problem with four parameters - a reserve price, two cutoff valuations that define observant bidders treated equally as unobservant bidders, and a price offer to them. This is a harder problem, but still computationally tractable. The numerical solutions we have found in simple environments suggest that fixed price trading is quite common. In fact, it is easy to see without any calculation that if every bidder is equally likely to be observant or unobservant, the trade will occur at a fixed price (with no auctions) more than half the time. This may be an explanation for why auctions aren't particularly common in many trading platforms. Even on trading platforms on which auctions *are* used, such as eBay, there are "buy it now" options where trading takes place at a fixed price.[1]

An equal priority auction is an indirect mechanism that we use to replicate the equilibrium outcome in the unobserved mechanism design game where unobservant bidders babble. Observant bidders with intermediate valuations and unobserved bidders are treated the same, but of course observant bidders know the auction rule while unobservant ones don't. If there is a hidden link on a website that gives access to the auction, then observant bidders are free to click on the link but unobservant ones wouldn't know how to do it. It doesn't matter to the seller or any bidder whether or not observant bidders with intermediate valuations click on the link to participate in the auction. This has an important implication. An econometrician who tries to recover distributions of valuations based on bids placed in the

[1]The environment on eBay doesn't fit our model exactly because bidders arrive randomly. Buy it now options disappear on eBay once a bidder with a low valuation submits a bid, while auctions continue to occur.

auction would get biased estimates, because intermediate valuations might be missing from the bidding data. Our model of unobserved mechanisms is admittedly crude, and we only look at babbling equilibria here, but equal priority auctions make empirically relevant points about the way rationally unobservant agents impact trading mechanisms.

2 Unobserved Mechanism Design

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 $[0, 1]$. We assume that all valuations are distributed according to some distribution F with strictly positive density f. Buyers' payoff when they buy at price p is given by $w - p$. The seller's cost is zero, so the profit from selling at price p is just p.

Define

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

as the revenue function from a take-it-or-leave-it offer w to a buyer. Following the standard auction literature, we also define

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

as the virtual valuation function.

Buyers are either *observant* or *unobservant*. We use τ_i as the "information type" of buyer i, with $\tau_i = \epsilon$ if i is observant, and $\tau_i = \mu$ if i is unobservant. Unobservant buyers communicate with the seller using a message space $\mathcal{M}^\mu$ - assumed to be a compact metric space which embeds $[0, 1]$, the set of values. Informed buyers have access to $\mathcal{M}^\mu$, and to a distinct message space $\mathcal{M}^\epsilon$. We'll assume that $\mathcal{M}^\epsilon$ is also compact and metric, and embeds $[0, 1]$. The important assumption here is that the seller can tell whether or not a message comes from $\mathcal{M}^\epsilon$. So when a buyer sends a message in $\mathcal{M}^\epsilon$, the seller knows that they are observant. If a buyer's messages comes from $\mathcal{M}^\mu$, the seller can't tell whether the buyer is observant or unobservant.[2]

[2] In section 4, we discuss a general framework of unobserved mechanism design where observant and

The seller writes an algorithm that processes the messages sent by all the buyers, then chooses which buyer to make an offer to. Unobservant buyers do not see the rules the seller is using to convert messages to offers. Informed buyers are fully aware of these rules. The seller and each of the buyers believes that each of the others is unobservant with probability $\alpha \in (0,1)$ independent of their valuation. All buyers, observant or unobservant, know that if the seller makes them an offer and they accept it, then the seller is committed to transact with them at the price. Since any offer can be rejected, this is quite different from standard mechanism design where a mechanism produces an allocation. It turns our problem into a game where the payoffs in the game are endogenously determined by the seller. Further, when an offer is rejected, the process ends without trade.[3]

Denote $\mathcal{M} = \mathcal{M}^{\mu} \cup \mathcal{M}^{\epsilon}$. Let $q_i : \mathcal{M}^n \to [0,1]$ be an integrable function that gives the probability with which an offer is made to bidder i for every possible profile of messages from buyers. A profile of these functions is feasible if

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

and $q_i(b) \geq 0$ for every profile $b = (b_1, \ldots, b_n) \in \mathcal{M}^n$. Next, let $\triangle[0,1]$ be a set of probability measures on the interval of values such that every bounded function is integrable, and let $P_i : \mathcal{M}^n \to \triangle[0,1]$ be an integrable function that describes the distribution of price offers buyer i receives conditional on receiving an offer. If we use the notation $\{P, q\} = \{(P_i)_{i=1}^n, (q_i)_{i=1}^n\}$, then the seller's mechanism or algorithm is just a feasible pair $\{P, q\}$. Let Γ be the set of all feasible mechanisms.

A strategy rule σ_i for buyer i is a pair of measurable functions $\{\sigma_i^{\epsilon}, \sigma_i^{\mu}\}$ with $\sigma_i^{\epsilon} : [0,1] \times \Gamma \to \mathcal{M}$ and $\sigma_i^{\mu} : [0,1] \to \mathcal{M}^{\mu}$ that specifies what messages the buyer will send for each of their valuations conditional on whether the buyer observes the seller's mechanism.[4] Write $\sigma = \{\sigma^{\epsilon}, \sigma^{\mu}\} = \{(\sigma_i^{\epsilon})_{i=1}^n, (\sigma_i^{\mu})_{i=1}^n\}$. We'll use the usual notation $\{\sigma_{-i}^{\epsilon}, \sigma_{-i}^{\mu}\}$ to refer to the

unobservant buyers use the same message space. We believe that any equilibrium outcome constructed here can be replicated in the general framework through randomization over mechanisms.

[3]In section 4, we discuss how to relax this assumption and allow the seller's algorithm to make multiple offers.

[4]To save notation, we consider only pure strategies by observant and unobservant buyers. The expressions and definitions introduced below can be easily extended to mixed strategies.

strategy rules used by the other players. These strategy rules depend on the other buyers' information types which buyer i doesn't know. When taking expectations, it is over both profiles of the other buyers' valuations v_{-i} and profiles of their information types τ_{-i}.

Let $\mathcal{R}(\gamma, \sigma)$ be the expected revenue for the seller from mechanism $\gamma = \{P, q\}$ when buyers use strategy rules given by σ. This is given by

$$\mathcal{R}(\gamma, \sigma) = \mathbb{E}_{v,\tau}\left[\sum_{i=1}^{n} q_i(\sigma)\int_{p_i \le v_i} p_i \; dP_i(p_i; \sigma)\right],$$

where the expectation is taken over profiles of buyers' valuations v and their information types τ.

The imperfect information game $\mathcal{G}$ is defined to be the extensive form game of imperfect information in which the seller first commits to some $\gamma \in \Gamma$, then the buyers send messages to the seller that depend on γ only if they are observant. The game $\mathcal{G}$ implicitly depends on the probability α with which buyers are unobservant. We omit the dependence of $\mathcal{G}$ on α for notational brevity.

Our solution concept uses a refinement of Bayesian Nash equilibrium. Neither perfect Bayesian nor sequential equilibrium work in our context because the seller can use mechanisms which preclude any kind of sequential rationality. For example, the seller could deviate to a mechanism in which all observant bidders are asked to submit bids with an offer with price 0 made to the buyer who submits the highest bid strictly less than 1. This would be a silly deviation. Yet no matter what beliefs the players hold about each other or what strategies they play, either some buyer will have a profitable deviation, or some buyers will not be able to find best replies.[5]

In order to describe the refinement, we need the following definitions:

Definition 1 *The* continuation game $\mathcal{G}(\gamma, \sigma^{\mu})$ *is the Bayesian game played by all the observant buyers where the seller's uses mechanism* $\gamma = \{P, q\}$ *and the unobservant buyers use strategy* σ^{μ}*. A profile of strategies* $\zeta_i : [0, 1] \to \mathcal{M}$ *used by each observant bidder* i *is called a*

[5] We thank a referee for pointing this out to us.

continuation equilibrium *of* $\mathcal{G}(\gamma, \sigma^{\mu})$ *if for all* i *with* $\tau_i = \epsilon$, $v_i \in [0,1]$, $b_i \in \mathcal{M}$,

$$\mathbb{E}_{v_{-i},\tau_{-i}} \left[q_i \left(\zeta_i(v_i), \zeta_{-i}(v_{-i}), \sigma^{\mu}_{-i}(v_{-i}) \right) \int \max \{ v_i - p_i, 0 \} \, dP_i \left(p_i; \zeta_i(v_i), \zeta_{-i}(v_{-i}), \sigma^{\mu}_{-i}(v_{-i}) \right) \right]$$
$$\geq \mathbb{E}_{v_{-i},\tau_{-i}} \left[q_i \left(b_i, \zeta_{-i}(v_{-i}), \sigma^{\mu}_{-i}(v_{-i}) \right) \int \max \{ v_i - p_i, 0 \} \, dP_i \left(p_i; b_i, \zeta_{-i}(v_{-i}), \sigma^{\mu}_{-i}(v_{-i}) \right) \right].$$

Using the above continuation idea, we can give a simple definition of Bayesian equilibrium.

Definition 2 *The mechanism* $\gamma = \{P, q\}$ *along with strategies* $\{\sigma^{\epsilon}, \sigma^{\mu}\}$ *constitutes a* Bayesian equilibrium *for the game* $\mathcal{G}$, *if* $\mathcal{R}(\gamma, \sigma^{\epsilon}, \sigma^{\mu}) \geq \mathcal{R}(\tilde{\gamma}, \sigma^{\epsilon}, \sigma^{\mu})$ *for all* $\tilde{\gamma} \in \Gamma$; $\sigma^{\epsilon}(\cdot, \gamma)$ *is a continuation equilibrium for* $\mathcal{G}(\gamma, \sigma^{\mu})$; *and for all* i *such that* $\tau_i = \mu$, $v_i \in [0,1]$, $b_i \in \mathcal{M}^{\mu}$,

$$\mathbb{E}_{v_{-i},\tau_{-i}} \left[q_i \left(\sigma^{\mu}_i (v_i), \sigma^{\epsilon}_{-i}(v_{-i}, \gamma), \sigma^{\mu}_{-i}(v_{-i}) \right) \int \max \{ v_i - p_i, 0 \} \, dP_i \left(p_i; \sigma^{\mu}_i (v_i), \sigma^{\epsilon}_{-i}(v_{-i}, \gamma), \sigma^{\mu}_{-i}(v_{-i}) \right) \right]$$
$$\geq \mathbb{E}_{v_{-i},\tau_{-i}} \left[q_i \left(b_i, \sigma^{\epsilon}_{-i}(v_{-i}, \gamma), \sigma^{\mu}_{-i}(v_{-i}) \right) \int \max \{ v_i - p_i, 0 \} \, dP_i \left(p_i; b_i, \sigma^{\epsilon}_{-i}(v_{-i}, \gamma), \sigma^{\mu}_{-i}(v_{-i}) \right) \right].$$

As usual this isn't a very restrictive solution concept since strategy rules used by the observant don't have to be a continuation equilibrium away from the equilibrium path after γ is offered. As we can't use solution concepts that impose sequential rationality off the equilibrium path, we use the following refinement:

Definition 3 *The triple* $\{\gamma, \sigma^{\epsilon}, \sigma^{\mu}\}$ *is a U-equilibrium if it is a Bayesian equilibrium and in addition there does not exist an alternative mechanism* $\tilde{\gamma}$ *and a continuation equilibrium* ζ *for* $\mathcal{G}(\tilde{\gamma}, \sigma^{\mu})$ *such that*

$$\mathcal{R}(\tilde{\gamma}, \zeta, \sigma^{\mu}) > \mathcal{R}(\gamma, \sigma^{\epsilon}, \sigma^{\mu}) \tag{1}$$

Since this is an unusual equilibrium concept, a few comments are in order. First, note that the concept of a continuation equilibrium depends on fixed behavior of the unobservant. This is because the unobservant don't know the mechanism that is being used off the equilibrium path.

Second, the seller's deviations to alternative mechanisms as described in (1) are restricted to those for which some continuation equilibrium exists. This avoids the problem when the seller offers a mechanism for which there is no continuation equilibrium among the observant.

If the message spaces $\mathcal{M}^{\epsilon}$ and $\mathcal{M}^{\mu}$ are both finite, then we could use perfect Bayesian equilibrium as part of our solution concept since a continuation equilibrium always exists. In this case, the on-path strategies in a U-equilibrium would always be part of a perfect Bayesian equilibrium. The sense in which our solution concept is stronger is that it selects out the seller-optimal perfect Bayesian equilibrium.

2.1 Direct mechanisms

We do not have a full characterization of all U-equilibria. However, we can characterize a special U-equilibrium called *babbling equilibrium*, where unobservant buyers send messages that are uninformative of their valuations, that is, $\sigma_i^{\mu}(w) = \sigma_i^{\mu}(\tilde{w})$ for all i and valuations $w, \tilde{w}$.[6] In any U-equilibrium, the behavior of the unobservant is known and fixed, and the rest of the equilibrium can be found by finding the seller's best reply to this behavior. Since the seller's commitment is seen by the observant buyers we can find this best reply using the revelation principle and solving for an optimal mechanism. The definition of U-equilibrium then requires the behavior of the unobservant to be a best reply to the optimal mechanism. This is generally a difficult fixed-point problem. For babbling equilibrium, however, the problem can be solved by restricting to direct mechanisms that ignore messages from unobservant buyers.

To do so we need to add some notation to describe a symmetric direct mechanism. In what follows the notation m always means the number of unobservant buyers (i.e., buyers who send messages in $\mathcal{M}^{\mu}$). We reorder n buyers such that the first $n-m$ of them are observant; the orders among the observant and among the unobservant are arbitrary. For each $v = (v_1, \ldots, v_n) \in [0,1]^n$, and for each $i = 1, \ldots, n-m$, let

$$\rho_m^i(v) = (v_i, v_2, \ldots, v_{i-1}, v_1, v_{i+1}, \ldots, v_{n-m}, v_{n-m+1}, \ldots, v_n\};$$

that is, $\rho_m^i(v)$ switches the positions of v_1 and v_i. Now we have

[6] We haven't been precise enough about the space of feasible mechanisms to prove existence. We partially address this issue in Theorem 1 below.

Definition 4 *A symmetric direct mechanism δ is a collection of functions*

$$\left\{ \left(q_m^{\epsilon}, p_m^{\epsilon}\right)_{m=0}^{n-1}, \left(q_m^{\mu}, p_m^{\mu}\right)_{m=1}^{n} \right\}$$

where $q_m^{\epsilon}, p_m^{\epsilon} : [0,1]^n \to [0,1]$ for each $m = 0, \ldots, n-1$, and $q_m^{\mu}, p_m^{\mu} : [0,1]^n \to [0,1]$ for each $m = 1, \ldots, n$, satisfying

- *$(q_m^{\tau}(v), p_m^{\tau}(v))$, $\tau = \epsilon, \mu$, are invariant to $(v_{n-m+1}, \ldots, v_n)$;*
- *$(q_m^{\epsilon}(v), p_m^{\epsilon}(v))$ are invariant to permutations of $(v_2, \ldots, v_{n-m})$, and $(q_m^{\mu}(v), p_m^{\mu}(v))$ are invariant to permutations of $(v_1, \ldots, v_{n-m})$;*
- *for all v and for all m,*

$$\sum_{i=1}^{n-m} q_m^{\epsilon}\left(\rho_m^i(v)\right) + m q_m^{\mu}(v) \leq 1. \tag{2}$$

The function $q_m^{\mu}(v)$ gives the probability with which an offer $p_m^{\mu}(v)$ is made to an unobservant buyer given that there are m unobservant buyers and the profile of valuations is $v = \{v_1, \ldots, v_n\}$. The function $q_m^{\epsilon}(v)$ gives the probability with which an offer $p_m^{\epsilon}(v)$ is made to buyer 1 given that there are m unobservant buyers and the valuation profile of buyers $i = 2, \ldots, n$ is $v_{-1} = \{v_2, \ldots, v_n\}$. Since unobservant buyers babble, we require the allocation and the offer functions of both the observant and the unobservant to be independent of the valuations of the latter. Symmetry requires the allocation and the offer functions of unobservant buyers to be invariant to permutations of the valuation profile of the observant, and the allocation and the offer functions of each observant buyer to be invariant to permutations of the valuation profile of the other observant buyers. Since $\rho_m^i(v)$ switches the positions of the first element of v and its i-th element, the sum $\sum_{i=1}^{n-m} q_m^{\epsilon}\left(\rho_m^i(v)\right)$ gives the probability that the offer is made to one of the first $n-m$ elements of v. Then (2) ensures that when the observant buyers have valuations given by the first $n-m$ valuations 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 unobservant buyers is less than or equal to 1.

We can use the above definitions to build something that looks exactly like a traditional reduced form mechanism. The probability with which an observant buyer whose valuation is w receives an offer when there are m unobservant is

$$Q_m^{\epsilon}(w) = \mathbb{E}_v\left[q_m^{\epsilon}(v)|v_1 = w\right].$$

Similarly

$$P_m^{\epsilon}(w) = \mathbb{E}_v\left[q_m^{\epsilon}(v)p_m^{\epsilon}(v)|v_1 = w\right]$$

is the expected price the observant bidder with valuation w would pay. Note that we have assumed that in any direct mechanism an observant buyer accepts the offer he receives with probability one. The is no max operator for observant buyers. This assumption is justified because observant buyers know the mechanism.

For each $m = 0, \ldots, n-1$, let $B(m; n-1, \alpha)$ be the probability that there are m unobservant 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:

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

We then have

$$U^{\epsilon}(w) = wQ^{\epsilon}(w) - P^{\epsilon}(w).$$

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

$$U^{\epsilon}(w) = \int_0^w Q^{\epsilon}(x)\,dx, \tag{3}$$

with $Q^{\epsilon}(\cdot)$ non-decreasing.[7]

The (interim) payoff to an unobservant bidder with valuation w is

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

Definition 5 *The mechanism δ is incentive compatible for observant buyers if (3) holds, $Q^{\epsilon}(\cdot)$ is non-decreasing and*

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

for every $w \in [0,1]$.

From standard arguments and properties of the binomial distribution, it is straightforward to show that the seller's revenue from any incentive compatible direct mechanism δ is given by

$$R(\delta) = n(1-\alpha) \int_0^1 Q^{\epsilon}(w)\, \phi(w) f(w) dw + \sum_{m=1}^{n} B(m; n, \alpha) \mathbb{E}_v \left[m q^{\mu}_m(v) \pi \left(p^{\mu}_m(v) \right) \right], \qquad (4)$$

where the first term is the revenue from observant buyers, and the second term is the revenue from the unobservant buyers. The following result provides a two-way relationship between the optimal direct mechanism and a U-equilibrium of the unobserved mechanism design game with babbling by unobservant buyers.

Theorem 1 *Fix a game of unobserved mechanisms $\mathcal{G}$. For any babbling equilibrium $\{\gamma, \sigma\}$, there is an incentive compatible and symmetric direct mechanism δ^*, with $R(\delta^*) = \mathcal{R}(\gamma, \sigma)$ and $R(\delta^*) \geq R(\delta)$ for every incentive compatible direct mechanism δ. Conversely, any incentive compatible and symmetric direct mechanism δ^* that maximizes $R(\delta)$ can be used to construct a babbling equilibrium (γ, σ) such that $\mathcal{R}(\gamma, \sigma) = R(\delta^*)$.*

The proof of this is straightforward. We provide a sketch of the argument here. Fix a babbling equilibrium $\{\gamma, \sigma\}$ in the game $\mathcal{G}$ with message spaces $\mathcal{M}^{\epsilon}$ and $\mathcal{M}^{\mu}$. The continuation equilibrium $\sigma^{\epsilon}(\cdot, \gamma)$ in the game $\{\gamma, \sigma^{\mu}\}$ on the equilibrium path is just an equilibrium

[7]See, for example, Myerson (1981). We have assumed $U^{\epsilon}(0) = 0$ for simplicity. This is usually not part of requirement for incentive compatibility, but clearly necessary for any revenue maximizing direct mechanism.

of a standard Bayesian game among the observant buyers. The seller doesn't actually care what the unobservant buyers say in an equilibrium in which their messages don't convey information about their types - all he needs to keep track of is whether or not a buyer's message was in $\mathcal{M}^\mu$. So it doesn't matter here whether unobservant buyers use asymmetric strategies. On the equilibrium path, some observant buyers may use a strategy $\sigma_i^\epsilon(v_i, \gamma)$ that mimics unobservant buyers for some valuations v_i. But by the standard revelation principle, there is an incentive compatible direct mechanism δ in which observant buyers report their information type and valuations truthfully, and gives the same expected revenue as γ. This direct mechanism δ might not be symmetric. However, it is well known that in the symmetric, independent private values environment, an asymmetric mechanism can't produce a higher expected revenue than a symmetric one. The definition of U-equilibrium allows the seller to choose the continuation equilibrium for the fixed strategy of unobservant buyers σ^μ. This means that there is a symmetric incentive compatible mechanism δ^* that achieves the equilibrium revenue $\mathcal{R}(\gamma, \sigma^\epsilon, \sigma^\mu)$ and is an optimal incentive compatible mechanism with respect to observant buyers.

The reverse direction follows by construction. Fix any message $b^\mu \in \mathcal{M}^\mu$. Let $\sigma_i^\mu(v_i) = b^\mu$ for all i and all $v_i \in [0,1]$. By assumption $\mathcal{M}^\epsilon$ embeds $[0,1]$ so we can find a subset of $\mathcal{M}^\epsilon$ and a bijection β^ϵ between this subset and $[0,1]$. For each i and $v_i \in [0,1]$, let $\sigma_i^\epsilon(v_i, \delta^*) = \beta^\epsilon(v_i)$, and $\sigma_i^\epsilon(v_i, \gamma) = b^\mu$ for all $\gamma \neq \delta^*$. Then, $\{\delta^*, \sigma^\epsilon, \sigma^\mu\}$ is a U-equilibrium of $\mathcal{G}$.

3 Equal Priority Mechanisms

For the remainder of the paper, we restrict attention to distribution functions such that $\pi(w)$ is strictly concave. 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 $w \geq r^*$.[8] The valuation r^* represents the optimal reserve price

[8]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.

in a standard auction, regardless of the number of buyers.[9] 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.

Our main result is that for valuation distributions such that $\pi(\cdot)$ is concave, the outcome of a U-equilibrium of the game $\mathcal{G}$ where unobservant buyers babble corresponds to an optimal "equal priority mechanism." We'll establish the main result in two parts. First we'll describe the set of equal priority mechanisms, and then the one that gives the seller the highest expected revenue. Later we'll show how to verify that the seller cannot do strictly better among all direct mechanisms.

An equal priority mechanism is fully characterized by four numbers, a "reserve price" r, a take-it-or-leave-it offer t, and the upper and lower bound $\overline{w}$ and $\underline{w}$ of an interval of buyer valuations, satisfying $r \leq \underline{w} \leq \overline{w}$. Unobservant buyers keep silent. Let m be the number of buyers who keep silent, and k be the number of reported valuations in the interval $[\underline{w}, \overline{w}]$. The mechanism treats the m unobservant buyers and the k observant buyers with the same allocation priority; we refer to them as "equal priority pool." Priorities of observant buyers with reported valuations above $\overline{w}$ and those with valuations below $\underline{w}$ are equal to the valuations themselves, with the former all higher and the latter all lower than the $m + k$ buyers in the equal priority pool. In words, the allocation and offers in an equal priority mechanism are determined in the following way:

- When the highest reported valuation is less than r: no offer is made to the buyer; for each $m \geq 1$, with probability $1/m$ an offer t is made to each unobservant buyer.
- When the highest reported valuation is between r and $\underline{w}$: if $m = 0$, the buyer is made an offer equal to the maximum of the second highest reported valuation and r; if $m \geq 1$, no offer is made to the buyer, and instead with probability $1/m$, an offer t is made to each unobservant buyer.

[9] 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 observant buyers is to give it to an unobservant 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^*.

- When the highest reported valuation is between $\underline{w}$ and $\overline{w}$: if $m+k=1$, the buyer is made an offer equal to the maximum of the second highest reported valuation and r; if $m+k \geq 2$, with probability $1/(m+k)$ an offer $\underline{w}$ is made to each observant buyer with reported valuation in the interval $[\underline{w}, \overline{w}]$ and an offer t to each unobservant buyer.

- When the highest reported valuation is above $\overline{w}$: the buyer is made an offer equal to the second highest reported valuation if it is above $\overline{w}$; if the second highest reported valuation is in $[\underline{w}, \overline{w}]$ or $m \geq 1$, an offer $(\underline{w} + (m+k)\overline{w})/(m+k+1)$ is made to the buyer; if it is below $\underline{w}$ and $m=0$, the buyer is made an offer equal to the maximum of the second highest reported valuation and r.

In terms of the offer rule, an equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ is a second-price auction with a reserve price r for observant buyers, combined with a take-it-or-leave-it offer t to unobservant buyers. However, we have an equal priority pool consisting of observant buyers with valuations between $\underline{w}$ and $\overline{w}$ and unobservant buyers. As a result, the second price, or the offer made to the buyer with the highest reported valuation, is the maximum of r and the second highest reported valuation, only if the second highest reported valuation is outside $[\underline{w}, \overline{w}]$, and only if there are no unobservant buyers when the second highest reported valuation is lower than $\underline{w}$.

Formally, using the notation of direct mechanisms introduced in section 2.1, we can represent an equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ as follows. Suppose that v_1 is the highest reported valuation, and v_2 be the second highest reported valuation. The collection of functions $\left\{(q^\epsilon_m(v), p^\epsilon_m(v))_{m=0}^{n-1}, (q^\mu_m(v), p^\mu_m(v))_{m=1}^{n}\right\}$ given by an equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ is

$$
\begin{cases}
q^\epsilon_m(v) = 0 & \text{if} \quad v_1 < r, \text{ or } v_1 \in [r, \underline{w}) \text{ and } m \geq 1 \\
q^\epsilon_m(v) = 1/(m+k), \; p^\epsilon_m(v) = \underline{w} & \text{if} \quad v_1 \in [\underline{w}, \overline{w}] \text{ and } m+k \geq 2 \\
q^\epsilon_m(v) = 1, \; p^\epsilon_m(v) = (\underline{w} + (m+k)\overline{w})/(m+k+1) & \text{if} \quad v_1 > \overline{w}, \text{ and } v_2 \in [\underline{w}, \overline{w}] \text{ or } m \geq 1 \\
q^\epsilon_m(v) = 1, \; p^\epsilon_m(v) = \max\{v_2, r\} & \text{if} \quad \text{otherwise,}
\end{cases}
$$

and

$$\begin{cases} q^{\mu}_{m}(v) = 0 & \text{if} \quad v_1 > \overline{w} \\ q^{\mu}_{m}(v) = 1/(m+k),\ p^{\mu}_{m}(v) = t & \text{if} \quad \text{otherwise.} \end{cases}$$

Suppose that observant buyers truthfully report their valuations in an equal priority mechanism. Then using the allocation rule, we can calculate the probability with which each type of observant buyer trades. This probability of trade function Q^{ϵ} for an observant buyer is

$$\begin{cases} 0 & \text{if} \quad w < r \\ (1-\alpha)^{n-1}F^{n-1}(w) & \text{if} \quad w \in [r, \underline{w}) \\ \sum_{m=0}^{n-1} B(m; n-1, \alpha) \sum_{k=0}^{n-1-m} B_k^{n-1-m}(\underline{w}, \overline{w})/(m+k+1) & \text{if} \quad w \in [\underline{w}, \overline{w}] \\ \sum_{m=0}^{n-1} B(m; n-1, \alpha) F^{n-1-m}(w) & \text{if} \quad w > \overline{w}, \end{cases} \tag{5}$$

where

$$B_k^{n-1-m}(\underline{w}, \overline{w}) = \binom{n-1-m}{k} (F(\overline{w}) - F(\underline{w}))^k F^{n-1-m-k}(\underline{w}).$$

We now provide more convenient formulas for Q^{ϵ}. For $w > \overline{w}$, we have

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

For $w \in [\underline{w}, \overline{w}]$, The trading probability $Q^{\epsilon}(w)$ for $w \in [\underline{w}, \overline{w}]$ plays a critical role in the analysis below, and for convenience we denote it as $\chi(\underline{w}, \overline{w})$. We re-do the double summations over m and k by first summing over k for fixed $l = m + k$ then summing over l, and rewrite $\chi(\underline{w}, \overline{w})$ as

$$\begin{aligned} &\sum_{l=0}^{n-1} \binom{n-1}{l} ((1-\alpha)F(\underline{w}))^{n-1-l} \frac{1}{l+1} \sum_{k=0}^{l} \binom{l}{k} ((1-\alpha)(F(\overline{w}) - F(\underline{w})))^k \alpha^{l-k} \\ =&\sum_{l=0}^{n-1} \binom{n-1}{l} ((1-\alpha)F(\underline{w}))^{n-1-l} \frac{1}{l+1} ((1-\alpha)(F(\overline{w}) - F(\underline{w})) + \alpha)^l. \end{aligned}$$

It follows that

$$\chi(\underline{w},\overline{w})=\frac{((1-\alpha)F(\overline{w})+\alpha)^n-((1-\alpha)F(\underline{w}))^n}{n((1-\alpha)(F(\overline{w})-F(\underline{w}))+\alpha)}. \tag{6}$$

The function χ gives the probability that a buyer whose valuation is in the pooling interval $[\underline{w},\overline{w}]$ receives an offer. The logic in $\chi(\underline{w},\overline{w})$ is that the buyer has the same chance of receiving an offer as any of the unobservant buyers and any other observant buyer whose reported valuation is in the interval $[\underline{w},\overline{w}]$, as long as there are no observant buyers with valuation above $\overline{w}$. This explains why in the formula (6) the denominator is the expected number of buyers in the equal priority pool, and the numerator is the total probability that there is at least one buyer, observant or unobservant, with that priority.

The following result gives the necessary and sufficient condition for the equal priority mechanism $\{r,\underline{w},\overline{w};t\}$ to be incentive compatible.

Lemma 1 *The equal priority mechanism $\{r,\underline{w},\overline{w};t\}$ is incentive compatible if and only if*

$$\int_r^{\underline{w}}(1-\alpha)^{n-1}F^{n-1}(w)dw\geq\chi(\underline{w},\overline{w})(\underline{w}-t) \tag{7}$$

Two arguments are needed to establish Lemma 1. The first is to show that the rules of allocation and offers are the ones that make truthful reporting incentive compatible by observant buyers. Note that when observant buyers report their valuations truthfully, they accept their offers with probability one. Since the allocation rule is monotone, we just need to show that the payoff of observant buyers $U^{\epsilon}(w)$ from truthful reporting satisfies (3) for each w.[10]

The second is to show that when t satisfies condition (7) no observant buyer can improve his payoff by pretending to be unobservant. Since $Q^{\epsilon}(w)=\chi(\underline{w},\overline{w})$ for all $w\in[\underline{w},\overline{w}]$, it follows from (3) that $U^{\epsilon}(w)$ is linear with slope $\chi(\underline{w},\overline{w})$. By construction, this is the same slope as the increasing part of the payoff function $U^{\mu}(w)$ for unobservant buyers, which is given by

$$U^{\mu}(w)=\chi(\underline{w},\overline{w})\max\{w-t,0\},$$

[10]Indeed, the offer rule is constructed from the allocation rule of the equal priority mechanism to ensure that it is incentive compatible with respect to valuations.

because unobservant buyers have the same allocation priority as observant buyers whose valuations are in $[\underline{w}, \overline{w}]$. Moreover, since by construction $Q^{\epsilon}(w)$ is strictly increasing for $w \in [r, \underline{w})$ and $w > \overline{w}$, it follows from (3) that the payoff function $U^{\epsilon}(w)$ is strictly convex for $w \geq r$ outside $[\underline{w}, \overline{w}]$. The equal priority auction $\{r, \underline{w}, \overline{w}; t\}$ is therefore incentive compatible if and only if $U^{\epsilon}(\underline{w}) \geq U^{\mu}(\underline{w})$. This is precisely (7).

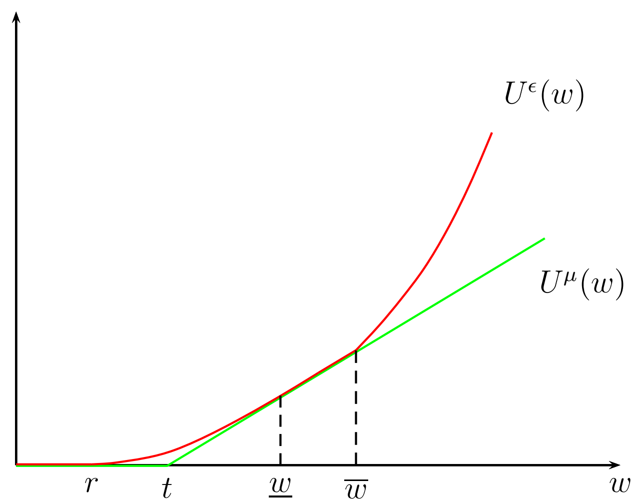

Figure 1. An equal priority mechanism with a binding incentive constraint.

Figure 1 above shows the payoffs to observant and unobservant buyers in an equal priority mechanism with a binding incentive compatible constraint (7). The green line represents the payoff function $U^{\mu}(\cdot)$ of an unobservant buyer or an observant buyer that acts as one. (It is zero for valuations below t.) The slope of the green line is $\chi(\underline{w}, \overline{w})$. The red curve represents the payoff function $U^{\epsilon}(w)$ to an observant buyer. It coincides with the green line for valuations in the pooling interval $[\underline{w}, \overline{w}]$ because the incentive condition (7) is binding, and is strictly convex for valuations between r and $\underline{w}$, and above $\overline{w}$. (It is equal to zero for valuations below r.)

In any equal priority mechanism, observant buyers with low valuations, between r and $\underline{w}$, and those with high valuations, above $\overline{w}$, are strictly worse off by pretending to be unobservant. If the incentive compatibility constraint (7) is binding, it is a matter of indifference for observant buyers with valuations in $[\underline{w}, \overline{w}]$ whether they truthfully report their valuations or wait for the take-it-or-leave-it offer t just like an unobservant buyer. Indeed, the same truth telling equilibrium among observant buyers is implemented if we change the offer rule, so that the offer received by an observant buyer with valuations in the pooling interval $[\underline{w}, \overline{w}]$ is always t, instead of the maximum of the second highest bid and reserve price r when there are no other buyers in the equal priority pool, and $\underline{w}$ when there is at least one unobservant buyer in the pool. Furthermore, by revenue equivalence, the seller's revenue from observant buyers is the same if all observant buyers with valuations in the pooling interval $[\underline{w}, \overline{w}]$ behave in the same way as unobservant buyers. Since the allocation probability $q^\mu(v)$ and the offer $p^\mu(v)$ for unobservant buyers depend only on the size of the equality priority pool, i.e., $m + k$, and not on its composition, the seller's revenue from unobservant buyers is also unaffected by whether or not observant buyers with valuations in $[\underline{w}, \overline{w}]$ pretend to be unobservant.

Any equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ with a binding incentive condition (7) is therefore payoff-equivalent for all buyers and the seller to the following *indirect* mechanism. All buyers, observant or unobservant, are asked to place their bids; the seller reveals a random password; unobservant buyers do not know the password and their bids are treated as meaningless babbles; observant buyers can match the password and have their bids accepted as valid, except that those in the pooling interval $[\underline{w}, \overline{w}]$ are treated as babbles; the allocation and offer rules otherwise mimic those in the equal priority mechanism, with l representing the total number of buyers who babble:

- When the highest bid is less than r: the seller keeps the object if $l = 0$; otherwise, with probability $1/l$ the seller makes an offer t to each babbling buyer.
- When the highest bid is between r and $\underline{w}$: if $l = 0$, the bidder wins and pays the maximum of the second highest reported valuation and r; if $l \geq 1$, with probability $1/l$, the seller makes an offer t to each babbling buyer.
- When the highest bid is above $\overline{w}$: the bidder wins; he pays the second highest bid if it is above $\overline{w}$, the maximum of r and the second highest bid if it is below $\underline{w}$ and $l = 0$, and otherwise $(\underline{w} + l\overline{w})/(l + 1)$.

The above indirect mechanism is what we refer to as the equal priority "auction" in the introduction. As in a standard auction, the winner is the one with the highest valid bid, and the auction commits all bidders to paying the price charged by the seller when they win. To bidders who submit valid bids, it looks like a second price auction: the winner pays the maximum of the second highest bid and a reserve price. A non-standard part is that bids in the interval $[\underline{w}, \overline{w}]$ are treated as uninformative. The other non-standard part is that the seller's reserve price depends on the number l of babbling buyers: it is r when $l = 0$, and otherwise it is $(\underline{w} + l\overline{w})/(l + 1)$. To bidders in the auction, the reserve price is therefore "secret." This is because the seller has the outside option of offering the object to a babbling buyer.

3.1 Optimal equal priority mechanism

Under an equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$, the seller's expected revenue from observant buyers is given by the first term in (4),

$$R(\delta) = n(1 - \alpha) \int_0^1 Q^\epsilon (w) \, \phi(w) f(w) dw, \tag{8}$$

where $Q^\epsilon(w)$ is specified in (5), and the revenue from unobservant buyers is given by the second term in (4), which is equal to

$$\sum_{m=1}^{n} B(m; n, \alpha) \sum_{k=0}^{n-m} B_k^{n-m}(\underline{w}, \overline{w}) \frac{m}{m+k} \pi(t) = n\alpha\chi(\underline{w}, \overline{w})\pi(t). \tag{9}$$

The optimal equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ maximizes the sum of (8) and (9) subject to $r \leq \underline{w} \leq \overline{w}$ and (7).

The following lemma characterizes optimal equal priority mechanisms. We assume that $\pi(\cdot)$ is strictly concave. This 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 $w \geq r^*$.[11] The valuation r^* represents the optimal reserve price in a standard auction,

[11] 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.

regardless of the number of buyers.[12]

Lemma 2 *Suppose that $\pi(\cdot)$ is strictly concave. If $\{r, \underline{w}, \overline{w}; t\}$ is an optimal equal priority mechanism, then*

$$0 < r < r^* < t < \underline{w} < \overline{w} < 1.$$

Further, (7) holds with equality, and

$$\alpha(\pi(t) - \phi(\overline{w})) = (1-\alpha)\left((\underline{w} - t)(\phi(\overline{w}) - \phi(\underline{w}))f(\underline{w}) + \int_{\underline{w}}^{\overline{w}} (\phi(\overline{w}) - \phi(w))f(w)dw \right); \quad (10)$$

$$-\alpha\pi'(t) = (1-\alpha)(\phi(\overline{w}) - \phi(\underline{w}))f(\underline{w}); \quad (11)$$

$$-\phi(r)f(r) = (\phi(\overline{w}) - \phi(\underline{w}))f(\underline{w}). \quad (12)$$

Our proof (in the appendix) first uses variational arguments to establish that the optimal mechanism is interior, satisfying $r < t < \underline{w} < \overline{w}$. In particular, $\underline{w} < \overline{w}$, so that the pooling interval $[\underline{w}, \overline{w}]$ is non-degenerate as long as unobservant buyers are present in the model, i.e., $\alpha > 0$. If the interval were degenerate, then the seller could cut the offer t to unobservant buyers and pool observant buyers with them by decreasing $\underline{w}$ and increasing $\overline{w}$. We show that cutting the price offer t has a first order revenue gain from unobservant buyers, and the corresponding pooling has only a second order revenue loss from observant buyers.

In any optimal equal priority mechanism, the incentive condition (7) for observant buyers with valuations in the pooling interval $[\underline{w}, \overline{w}]$ is binding. Otherwise, in Figure 1 we would have a line segment in the payoff function $U^\epsilon(\cdot)$ for observant buyers parallel to, and above, the linear part of the payoff function $U^\mu(\cdot)$ for unobservant buyers. The seller would then want to either shrink the pooling interval, by increasing $\underline{w}$ and decreasing $\overline{w}$, or raise the take-it-or-leave-it offer t to unobservant buyers.

The three conditions (10), (11) and (12) are the first order conditions for an interior op-

[12] In much of the auction literature, the seller has the fixed outside option of keeping the object. 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 observant buyers is to give it to an unobservant buyer with a take-it-or-leave-it offer, and is endogenous because it is chosen by the seller. We do not need to assume that $\phi(w)$ is strictly increasing for valuations below r^*.

timum.[13] In an optimal equal priority mechanism, the reserve price r for selling to observant buyers when there are no unobservant buyers is set below the standard optimal reserve price r^* in the absence of unobservant buyers, as can be seen from (12). This sacrifices revenue when all observant buyers have low valuations and there are no unobservant buyers, but provides incentives for observant buyers to truthfully report their valuations instead of pretending to be unobservant. Correspondingly, (11) implies that the take-it-or-leave-it price t to unobservant buyers is raised above the optimal monopoly price r^* in the absence of observant buyers. This reduces the revenue when all buyers are unobservant, but provides disincentive for observant buyers to pretend to be unobservant.

If the seller does not give the object to an observant buyer, she can always make a take-it-or-leave-it offer to an unobservant buyer if there is one. Absent incentives, the seller would set the reserve price $\overline{r}(t)$ for observant buyers so that the virtual valuation is equal to the expected profit $\pi(t)$ of making the offer t to an unobservant buyer:

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

By condition (10), the optimal equal priority mechanism has $\phi(\overline{w}) < \pi(t)$. This means that the seller gives the object to observant buyers even though their virtual valuations are lower than the value of the seller's "outside option" $\pi(t)$. This reason for doing this is to provide incentives for truthful reporting by observant buyers with valuations just above $\overline{w}$ rather than wait for the take-it-or-leave-it offer by pretending to be unobservant.

When all buyers are surely observant the revenue from the optimal equal priority mechanism converges to the revenue from the standard auction with reserve price r^*, as it becomes optimal for the seller not to distort the reserve price r at all to provide incentives (equation (12)). The pooling interval shrinks to a single valuation v_0 as α goes to 0,[14] satisfying

[13] They are all derived with variational arguments without explicitly using a multiplier for (7). For example, condition (10) is obtained by marginally changing $\underline{w}$ and $\overline{w}$ such that (7) is satisfied and then considering the effects on the seller's revenue. From the proof in the appendix, it can be seen that the value of the multiplier associated with (7) is the right hand side of (11) multiplied by n. This turns out to be the integral of the multiplier function $\lambda(\cdot)$ in the proof of Theorem 2 over the valuation support $[0, 1]$.

[14] The limit of $\chi(\underline{w}, \overline{w})$ as α goes to 0 and $\underline{w}$ and $\overline{w}$ shrink to the same point of v_0 is $F^{n-1}(v_0)$. That is, when all other buyers are almost surely observant, a deviating observant buyer will be the only buyer in the equal priority pool and will win the object with probability one if all other buyers (who are observant) have valuation below v_0.

the binding constraint (7) that an observant buyer with valuation v_0 is indifferent between truthfully reporting it and receiving a take-it-or-leave-it offer t_0 when all other buyers have valuations below v_0,

$$\int_{r^*}^{v_0} F^{n-1}(w)dw = F^{n-1}(v_0)(v_0 - t).$$

The limit values of v_0 and t_0 satisfy the above indifference condition and the limit version of first order conditions (10) and (11), given by

$$\pi'(t_0)(v_0 - t) + \pi(t_0) - \phi(v_0) = 0.$$

We have $t_0 > r^*$ and $\pi(t_0) > \phi(v_0)$. When α is arbitrarily close to 0, the incentives for observant buyers not to pretend to be unobservant are provided by raising the take-it-leave-it offer to an unlikely unobservant buyer above r^*, and not selling to unobservant buyers even when the profit from doing so exceeds virtual valuations of observant buyers.

In the opposite limit of $\alpha = 1$, buyers are surely unobservant, and the revenue from the optimal equal priority mechanism converges to the revenue from a take-it-or-leave-it offer r^*. By (11), the seller no longer distorts t to provide incentives for observant buyers. From (10), the upper-bound of the pooling interval converges to $\overline{r}(r^*)$, satisfying

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

as the need for the seller to provide incentives for observant buyers with valuations just above the upper-bound becomes second order. From the binding constraint (7), the lower-bound of the pooling interval becomes r^*.[15] This is to prevent an unlikely observant buyer with a valuation equal to the lower bound from pretending to be unobservant, as the buyer has close to zero chance of making the winning bid with the limit reserve price r_1 satisfying (12)

$$-\phi(r_1)f(r_1) = \pi(r^*)f(r^*).$$

As long as α is strictly less than 1, however, the mechanism is what provides incentives for

[15] The limit of $\chi(\underline{w}, \overline{w})$ as α goes to 1 is $1/n$, as an unlikely observant buyer will surely face $n-1$ unobservant buyers in the equal priority pool after pretending to be unobservant.

observant buyers with valuations just below the lower bound of the interval not to pretend to be unobservant.

3.2 Optimal direct mechanisms

We want to show that an optimal equal priority mechanism provides the seller the highest expected revenue among all direct mechanisms. Optimizing over all incentive compatible direct mechanisms is difficult, due to the continuum of incentive constraints for observant buyers with any valuation w not to pretend to be unobservant. Instead we adopt an indirect approach, by incorporating the continuum of constraints through a multiplier function. This is known as the Lagrangian relaxation method.

Recall that a direct mechanism δ consists of a series of functions $(q^{\epsilon}_m(v), p^{\epsilon}_m(v))_{m=0}^{n-1}$ and $(q^{\mu}_m(v), p^{\mu}_m(v))_{m=1}^{n}$. We first use the assumption that $\pi(\cdot)$ is strictly concave to simplify the optimal design problem. Replacing all these offers with the expected offer reduces the deviation payoff to observant buyers from pretending to be unobservant. Concavity then implies a greater revenue from unobservant buyers.

Lemma 3 *If $\pi(\cdot)$ is strictly concave, then in any optimal direct mechanism, $p^{\mu}_m(v)$ is independent of m and v.*

Using Lemma 3, we denote the constant price offered to the unobservant as p^{μ}. Define

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

to be the total probability of an offer expected by an unobservant buyer (or a deviating observant buyer).

Next, we drop the transfers $(p^{\epsilon}_m(v))_{m=0}^{n-1}$ to observant buyers, and construct the relaxed Lagrangian using only allocations $(q^{\epsilon}_m(v))_{m=0}^{n-1}$. Once we show that an optimal equal priority $\{r, \underline{w}, \overline{w}; t\}$ solves the relaxed Lagrangian, we can then use the offer rule in section 3 to construct the transfers $(p^{\epsilon}_m(v))_{m=0}^{n-1}$ and the resulting payoff function $U^{\epsilon}(\cdot)$, and apply Lemma 1 to conclude that the solution is incentive compatible.

We are thus led to the following maximization problem: Choose $(q_m^{\epsilon}(v))_{m=0}^{n-1}$, $(q_m^{\mu}(v))_{m=1}^{n}$, and p^{μ} to maximize

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

subject to the feasibility constraint (2), $Q^{\epsilon}(\cdot)$ is non-decreasing, and for every w

$$\int_0^w Q^{\epsilon}(x)\,dx \geq Q^{\mu}\max\{w-p^{\mu},0\}. \tag{13}$$

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

$$\begin{aligned} n(1-\alpha)&\int_0^1 Q^{\epsilon}(w)\,\phi(w)f(w)dw + n\alpha Q^{\mu}\pi(p^{\mu}) \\ &+\int_0^1 \lambda(w)\left(\int_0^w Q^{\epsilon}(x)\,dx - Q^{\mu}\max\{w-p^{\mu},0\}\right)dw, \end{aligned}$$

with the same choice variables and constraints except (13). That is, by introducing the Lagrangian function, we incorporate a continuum of constraints (13) into the objective function of the relaxed problem as an extra term.

The above relaxed problem has different solutions depending on the choice of $\lambda(\cdot)$. Regardless of the choice of $\lambda(\cdot)$, however, the value of the relaxed problem is an upper bound on the value of the full problem, because the solution to the full problem is feasible for the relaxed problem and because the extra term in the objective function of the relaxed problem is non-negative by construction. We will try to construct a function $\lambda(\cdot)$ such that the solution to the relaxed problem is an optimal equal priority mechanism. Since the equal priority mechanism yields an upper bound on the seller's revenue in the full problem, and since it satisfies all the constraints in the full problem, it solves the full problem.

The multiplier function $\lambda(\cdot)$ is the shadow cost (benefit) of violating (relaxing) the constraints (13). The second term in the relaxed Lagrangian is the total shadow value. The relaxed problem is then choosing feasible allocations $(q_m^{\epsilon}(v))_{m=0}^{n-1}$ and $(q_m^{\mu}(v))_{m=1}^{n}$, together with p^{μ}, to maximize the sum of the resulting revenues from observant and unobservant buyers and the shadow values. The key to our construction of the desired $\lambda(\cdot)$ is that, first, it satisfies complementary slackness so that the extra term in the relaxed Lagrangian is zero; and second, the allocations of an optimal equal priority auction characterized by Lemma 2 maximize the sum of the revenues and the shadow values. More precisely, we use integration by parts and rewrite the Lagrangian as

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

We want to choose $\lambda(\cdot)$ to have the following properties: (i) It is equal to 0 outside of $[\underline{w}, \overline{w}]$ where the constraint (13) is slack. (ii) It is non-negative on $[\underline{w}, \overline{w}]$ and makes the expression in the first bracket in the above Lagrangian constant, so that it is point wise maximizing to have constant $Q_m^{\epsilon}(w)$ for all $w \in [\underline{w}, \overline{w}]$. (iv) The constant value of the expression in the first bracket in the above Lagrangian matches the constant value of the expression in the second bracket, so that it is point wise maximizing to give the same allocation priority to observant buyers with valuations in the pooling interval and unobservant buyers. (iv) The value of the expression in the first bracket is strictly increasing and greater than that in the second bracket for $w > \overline{w}$, and increasing and smaller for $w < \underline{w}$, so that observant buyers have increasing and higher priorities than unobservant buyers if their valuations are higher than $\overline{w}$, and increasing and lower priorities if their valuations are lower than $\underline{w}$.

Theorem 2 *Suppose that $\pi(\cdot)$ is strictly concave. Then, there is no incentive compatible direct mechanism that yields a strictly greater revenue than an optimal equal priority mechanism.*

Putting together Theorems 2 and 1, we have shown that when $\pi(\cdot)$ is concave, the outcome of a U-equilibrium of the game $\mathcal{G}$ where unobservant buyers babble corresponds to an optimal equal priority mechanism. Conversely, once we solve for the optimal equal priority mechanism, we can construct a U-equilibrium with the same outcome. Since equal priority mechanisms are relatively straightforward to describe and optimize over, our result provides a simple characterization of equilibrium outcomes of the unobserved mechanism design game in the important class where unobservant buyers babble.[16]

The relative simplicity of optimal equal priority mechanisms also allows us to understand welfare implications of unobserved mechanism design. The seller is of course worse off compared to when all buyers are observant, as unobservability reduces the power of commitment necessary for standard optimal auctions. This means that the seller has incentives to "educate" buyers about her mechanism. But such attempt would be thwarted so long as the commitments in the mechanism remain unverifiable.

When all n buyers are surely observant, they face the standard optimal reserve price of r^*. In a U-equilibrium of the unobserved mechanism design game $\mathcal{G}$ with $\alpha > 0$, the seller sets $r < r^*$, so an observant buyer with a valuation between r and r^* is better off than when there are no unobservant buyers around. Observant buyers with higher valuations are affected by the presence of unobservant buyers in two opposing ways: they can win even though some unobservant buyer has a higher valuation, but they may also lose to an unobservant with a lower valuation. The net effect is generally ambiguous, but we can show that observant buyers with sufficiently high valuations benefit from having unobservant buyers around if the number of buyers is sufficiently large.[17]

For unobservant buyers, the relevant welfare comparison question is how they are affected by the presence of observant buyers. If there are no observant buyers, unobservant buyers have an equal chance of receiving a take-it-or-leave-it offer equal to r^*. Since in a babbling equilibrium of $\mathcal{G}$ the seller sets the take-it-or-leave-it offer t strictly above r^*, an unobservant

[16] Indeed, the first order conditions (10), (12) and (11), together with the binding constraint (7), are sufficient as well as necessary for an optimal equal priority mechanism. The sufficiency comes from the fact that the proof of Theorem 2 uses only the first order conditions. That is, Theorem 2 actually shows that an equal priority mechanism that satisfies the first order conditions are optimal among all direct mechanisms, and a fortiori, optimal among all equal priority mechanisms.

[17] To see this, note that

$$U^\epsilon(1) = \int_r^1 Q^\epsilon(w)dw > \int_{\overline{w}}^1 ((1-\alpha)F(w)+\alpha)^{n-1}dw.$$

The above is greater than $\int_{r^*}^1 F^{n-1}(w)dw$ when n is sufficiently large, because by integration by parts, it is implied by

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

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

buyer with a valuation w just above r^* is worse off in equilibrium than when there are no observant buyers around. For unobservant buyers with higher valuations, they have a higher priority than observant buyers with valuations below $\underline{w}$, which makes them better off in equilibrium, but lose out to observant buyers with valuations above $\overline{w}$. The net effect is again ambiguous, but we can show that unobservant buyers are all worse off in equilibrium than when there are no observant buyers if the number of buyers is large.[18]

4 Discussion

The main insight of the present paper has to do with how unobservability affects mechanism design. Although they do not observe the seller's mechanism, unobservant buyers in our model are meant to be fully rational rather than behavioral. In any equilibrium of the imperfect information game $\mathcal{G}$, as rational players unobservant buyers know the strategies of the seller and the observant buyers. To focus on the main insight of unobserved mechanism design, we have taken a short cut by assuming that there is a separate message space $\mathcal{M}^\epsilon$ for observant buyers that is not accessible to unobservant buyers. This prevents unobservant buyers from participating in the seller's mechanism in the same way as observant buyers do, even though they know the seller's mechanism. At the same time, there is a message space $\mathcal{M}^\mu$ for unobservant buyers that is also accessible to observant buyers. This is where the assumption that unobservant buyers are fully rational leads us to conclude that communication of their valuations by unobservant buyers is severely restricted - in fact in this paper there is no communication at all - because the seller can't refrain from exploit such communication through deviations that unobservant buyers don't see. Since observant buyers can pretend to be unobservant by choosing their message from $\mathcal{M}^\mu$, and since the seller gets part of the revenue from unobservant buyers, the presence of unobservant buyers affects the seller's auction design with observant buyers, even though unobservant buyers don't know how to participate in the same way as observant buyers.

[18]We have

$$U^\mu(1) = \chi(\underline{w}, \overline{w})(1-t) < ((1-\alpha)F(\overline{w}) + \alpha)^{n-1}(1-r^*).$$

The above is less than $(1-r^*)/n$ when n is sufficiently large. Since the payoff functions are piece wise linear, an unobservant buyer with any valuation is worse off in equilibrium.

A more general approach to unobserved mechanism design would be an imperfect information game $\overline{\mathcal{G}}$ where the message space $\overline{\mathcal{M}}$ in the seller's mechanism used by all buyers, observant or unobservant, is common knowledge. An intuitive conjecture is that, any U-equilibrium outcome in $\mathcal{G}$ with some $\mathcal{M}^\epsilon$ and $\mathcal{M}^\mu$ can be replicated by a suitably defined U-equilibrium in $\overline{\mathcal{G}}$ with a sufficiently rich $\overline{\mathcal{M}}$, and vice versa.[19] The general approach is beyond the scope of this paper as it raises two challenging technical issues. First, we need to formally define randomizations over mechanisms in $\overline{\mathcal{G}}$ in order to give the seller the option of preventing unobservant buyers from participating in the mechanism in the same way as observant buyers. The idea is that even though unobservant buyers know the seller's mixed strategy and can use messages in $\overline{\mathcal{M}}$ as observant buyers do, the realized mechanism is seen by observant buyers only and is never correctly "guessed" by unobservant buyers.[20] Second, we need to be able to compute the expected payoff of the seller and the expected payoffs of observant and unobservant buyers for any valuation and over any mixture of mechanisms, given strategies of the observant and the unobservant. For any fixed realization of a mixture, for each valuation, the revenue of the seller, the expected payoff of an observant buyer, and the expected payoff of an unobservant buyer, are all well-defined from the strategies of the observant and the unobservant. The expectations of these payoffs over a given mixture are examples of functional integration. The domain of the integration is a function space, which represents the space of mechanisms; each integrand is a functional, which represents the revenue of the seller, the payoff of an observant buyer or the payoff of an unobservant buyer with a given valuation, for each mechanisms in the domain.

[19]The following is a verbal argument that the outcome of any U-equilibrium $\{\gamma, \sigma^\epsilon, \sigma^\mu\}$ in $\mathcal{G}$, where $\gamma = \{P, q\}$ is defined over $\mathcal{M}^\epsilon$ and $\mathcal{M}^\mu$, can be replicated by a mixed U-equilibrium $\{\psi, \overline{\sigma}^\epsilon, \overline{\sigma}^\mu\}$ in $\overline{\mathcal{G}}$ through "password" mechanisms. The message space $\overline{\mathcal{M}}$ in $\overline{\mathcal{G}}$ is the product of $\mathcal{M}$ and $[0,1]$, the latter meant to represent the space of a random password, with a typical element denoted as x. A message by buyer i is $\overline{b}_i = (b_i, x_i)$, where $b_i \in \mathcal{M}^\epsilon \cup \mathcal{M}^\mu$ and $x_i \in [0,1]$. With x uniformly distributed on $[0,1]$, each realized mechanism $\overline{\gamma}(x) = \{\overline{P}, \overline{q}\}$ in the mixture ψ has the identical selection and offer probabilities $\overline{q}$ and $\overline{P}$ as q and P, if the first part of each message in a profile either comes from $\mathcal{M}^\epsilon$ when the second part matches the password x or comes from $\mathcal{M}^\mu$ when the second part doesn't match x, and otherwise keeps the good. The rest of the argument for replication is to construct the strategies $\overline{\sigma}_i^\mu$ and $\overline{\sigma}_i^\epsilon$ to mimic σ_i^μ and σ_i^ϵ through password-matching.

[20]We may formally think of a mixture of mechanisms as a stochastic process, where a mechanism - mappings from profiles of bids in $\overline{\mathcal{M}}$ to profiles of probabilities of selection and offers - becomes a sample function. Then we can apply the Kolmogorov extension theorem to construct a probability measure associated with the stochastic process. This requires us to specify the distributions of selection probabilities and offers for each profile of bids induced by the mixture.

We have assumed that the output of the seller's mechanism is a single take-it-or-leave-it offer in the unobserved mechanism design game. If this offer is rejected, which it will sometimes be if it is made to an unobservant bidder, the game ends without trade. We view this game form as a "canonical" one, because under the standard mechanism design problem with observant agents, it is never optimal for the seller to make offers that might be rejected. As we have mentioned in the introduction, a separate motivation for this particular game form is that we can obtain empirically relevant results through equal priority auctions.

The assumption of a single take-it-or-leave-it offer is without loss for observant buyers, since, as in a standard auction, they will always want to accept the offer when they are made one. For the unobservant this assumption is perhaps unrealistic. Once the seller learns who the unobservant buyers are, the seller is likely to approach them in sequence with offers. One question is how this might change if the seller could follow up a rejection by making a possibly lower offer to one of the other unobservant bidders.

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 model, the seller first makes a commitment in terms of how a particular sequence of offers 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, and unobservant buyers babble. The main insights are intact - an unobservant buyer receives an expected offer independent of the buyer's valuation, while observant buyers face an outside option of waiting for their turn to receive an offer if they decide to pretend to be unobservant. We conjecture that the equilibrium outcome with babbling by unobservant buyers can be characterized by a similar equal priority mechanism as in the present model, with the single offer to unobservant buyers replaced with a decreasing sequence of offers. The seller's equilibrium revenue should be higher than the present single-offer model, because being able to make a sequence of offers improves the seller's revenue from unobservant buyers, without necessarily increasing the value of outside option to observant buyers who pretend to be unobservant.

A more challenging question with multiple offers arises if the seller's algorithm is not restricted to at most one take-it-or-leave-it offer for each buyer. Since an unobservant buyer does not observe the seller's deviations to other algorithms, rejecting an offer from the seller could reveal information about his valuation that could be exploited later by the seller. Yet we can make one observation. The equilibrium when the seller's algorithm is restricted to at most one take-it-or-leave-it offer for each buyer can be supported as an equilibrium when the algorithms are unrestricted. Imagine that an unobservant buyer disappears after rejecting an offer, believing the seller's algorithm makes at most one offer to each buyer. Given this belief by unobservant buyers, committing to an algorithm that potentially makes multiple offers to a given buyer would only affect the behavior of observant buyers. This then becomes unprofitable because observant buyers observe the seller's commitment.

There may be other equilibria when the seller's algorithm is not restricted to at most one take-it-or-leave-it offer for each buyer. It would be interesting to find out if any of these equilibria makes the seller better off compared to the equilibrium when the seller can make at most one offer to a buyer. We defer these questions to future research since it not clear at this point what is the best way to generalize to multiple offers to each buyer.

We have assumed that buyers are either fully observant or fully unobservant. 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.[21] 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.

[21] This corresponds to the main result of Akbarpour and Li [2020], who frame the issue of partial observability in terms of limited commitment by the seller.

5 Related Literature

As mentioned above, the idea that consumers might not notice prices is an old one in the marketing literature. 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.[22] These papers considered the same problem that we do, which is how this unobservant buyers would affect the prices that firms offer. The difference is that we are interested in mechanisms, not prices.

The presence of unobservant buyers provides type dependent outside options to observant buyers. This is the basic problem 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 observant. This makes it resemble later papers on competing mechanisms 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 unobservant 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. [2019], or Skreta [2015]). To our knowledge, our model is the first to study commitment 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

[22] See also Varian [1980], or Stahl [1994]. In Varian [1980], unobservant buyers are loyal to a specific seller, observant buyers are just interested in the lowest price.

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 equilibrium mechanisms.

Our observant buyers can "prove" they are observant in the same sense as Porath et al. [2014]. The main difference is that they assume that the social choice function is known by all the players, while in our model the driving force is the presence of buyers who are unobservant of the seller's mechanism. They also assume players have complete information about the state, but in our model only buyers know their own valuations.

Finally, our observant buyers can pretend they are unobservant but not the other way around. The one-sidedness of this incentive condition is similar to Denekere and Severinov [2006], who study an optimal non linear pricing problem with a fraction of consumers constrained to reporting their valuations truthfully.[23] As in our paper, their mechanism separates "honest" consumers from "strategic" consumers who can misrepresent their valuations costlessly. The main difference is that we start with a standard independent private value auction problem rather than a non linear pricing problem. More importantly, our unobservant buyers are uncommunicative in the class of equilibria we focus on, but they are rational rather than behavioral.

6 Appendix: Omitted Proofs

Proof of Lemma 1

We verify that the expected payoff of an observant buyer with valuation w matches $U^{\epsilon}(w)$ given by (3) and (5). There are four cases.

(i) By truthfully reporting his valuation, an observant buyer with $w < r$ never wins the object, and thus the expected payoff is 0, matching $U^{\epsilon}(w)$ in (5) and (3) for $w < r$.

(ii) By truthful reporting, an observant buyer with $w \in [r, \underline{w})$ wins the object only when $m = 0$ and all $n-1$ other observant buyers have valuation at most w, pays the maximum

[23] See also Sher and Vohra [2015]. They use graph theory to study a more general non linear pricing problem with voluntary provision of hard evidence.

of r and the second highest valuation. Thus, the expected payoff is

$$w(1-\alpha)^{n-1}F^{n-1}(w) - \left(r(1-\alpha)^{n-1}F^{n-1}(r) + \int_r^w x \, d\left((1-\alpha)^{n-1}F^{n-1}(x)\right)\right).$$

By integration by parts, the above matches $U^\epsilon(v)$ in (3) and (5) for $v \in [r, \underline{w})$.

(iii) By truthful reporting, an observant buyer with $w \in [\underline{w}, \overline{w}]$ wins the object with probability one when $m = 0$ and all $n-1$ other observant buyers have valuation at most $\underline{w}$, and pays the maximum of r and the second highest valuation. The contribution of this event to the buyer's expected payoff is

$$\begin{aligned} & w(1-\alpha)^{n-1}F^{n-1}(\underline{w}) - \left(r(1-\alpha)^{n-1}F^{n-1}(r) + \int_r^{\underline{w}} x \, d\left((1-\alpha)^{n-1}F^{n-1}(x)\right)\right) \\ =& U^\epsilon(\underline{w}) + (w - \underline{w})(1-\alpha)^{n-1}F^{n-1}(\underline{w}). \end{aligned}$$

The buyer also wins the object with probability $1/(m+k+1)$ when there are m unobservant buyers, all $n-m-1$ other observant buyers have valuation at most $\overline{w}$, and $m+k$ is at least 1 (where k is the number of observant buyers with valuation on $[\underline{w}, \overline{w}]$), and pays $\underline{w}$. The contribution of this event to the buyer's expected payoff is

$$(w - \underline{w})\left(\chi(\underline{w}, \overline{w}) - (1-\alpha)^{n-1}F^{n-1}(\underline{w})\right).$$

The sum of the above two expressions matches $U^\epsilon(w)$ in (3) and (5) for $w \in [\underline{w}, \overline{w}]$.

(iv) By truthful reporting, an observant buyer with $w > \overline{w}$ wins the object with probability one when $m = 0$ and the second highest bid is below $\underline{w}$, and he pays the maximum of the second highest bid and the reserve price r. The contribution to the expected payoff is

$$U^\epsilon(\underline{w}) + (w - \underline{w})(1-\alpha)^{n-1}F^{n-1}(\underline{w}).$$

He also wins with probability one when the second highest bid is below $\overline{w}$ and $m + k \geq 1$, and pays $(\underline{w} + \overline{w}(m+k))/(m+k+1)$. The contribution to the expected payoff is

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

Finally, the observant buyer with $w > \overline{w}$ wins with probability one and pays the second highest bid x when it is above $\overline{w}$, which occurs with probability

$$\sum_{m=0}^{n-1} B(m; n-1, \alpha)(F^{n-1-m}(x) - F^{n-1-m}(\overline{w})).$$

By integration by parts, the contribution to the expected payoff is

$$\int_{\overline{w}}^{w} \sum_{m=0}^{n-1} B(m; n-1, \alpha)(F^{n-1-m}(x) - F^{n-1-m}(\overline{w}))dx$$
$$= \int_{\overline{w}}^{w} \sum_{m=0}^{n-1} B(m; n-1, \alpha) F^{n-1-m}(x) dx - (w - \overline{w})((1-\alpha)F(\overline{w}) + \alpha)^{n-1}.$$

The sum of the three expressions for the contributions to the expected payoff matches $U^{\epsilon}(w)$ in (3) and (5) for $w > \overline{w}$.

Proof of Lemma 2

Fix an incentive compatible, optimal equal priority mechanism $\{r, \underline{w}, \overline{w}; t\}$ with $r \leq \underline{w} \leq \overline{w}$. When $r \leq t \leq \underline{w}$, define

$$D = U^{\epsilon}(\underline{w}) - U^{\mu}(\underline{w}) = \int_{r}^{\underline{w}} (1-\alpha)^{n-1} F^{n-1}(w) dw - \chi(\underline{w}, \overline{w})(\underline{w} - t),$$

and let R be the revenue, which is the sum of (8) and (9). If $0 < r < \underline{w}$, or if $0 = r < \underline{w}$ and $dr > 0$, or if $0 < r = \underline{w}$ and $dr < 0$, we have

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

If $0 < t < \underline{w}$, or $0 = t < \underline{w}$ and $dt > 0$, or $0 < t = \underline{w}$ and $dt < 0$, we have

$$\frac{\partial D}{\partial t} = \chi(\underline{w}, \overline{w}); \; \frac{\partial R}{\partial t} = n\alpha\chi(\underline{w}, \overline{w})\pi'(t).$$

If $t < \underline{w} < \overline{w}$, or if $t = \underline{w} < \overline{w}$ and $d\underline{w} > 0$, or $t < \underline{w} = \overline{w}$ and $d\underline{w} < 0$, we have

$$\begin{aligned}
\frac{\partial \chi(\underline{w}, \overline{w})}{\partial \underline{w}} &= \frac{(1-\alpha)f(\underline{w})}{(1-\alpha)(F(\overline{w}) - F(\underline{w})) + \alpha} \left(\chi(\underline{w}, \overline{w}) - ((1-\alpha)F(\underline{w}))^{n-1}\right); \\
\frac{\partial D}{\partial \underline{w}} &= (1-\alpha)^{n-1}F^{n-1}(\underline{w}) - \chi(\underline{w}, \overline{w}) - \frac{\partial \chi(\underline{w}, \overline{w})}{\partial \underline{w}}(\underline{w} - t); \\
\frac{\partial R}{\partial \underline{w}} &= n(1-\alpha)((1-\alpha)^{n-1}F^{n-1}(\underline{w}) - \chi(\underline{w}, \overline{w}))\phi(\underline{w})f(\underline{w}) \\
&\quad + n((1-\alpha)(\pi(\underline{w}) - \pi(\overline{w})) + \alpha\pi(t))\frac{\partial \chi(\underline{w}, \overline{w})}{\partial \underline{w}}.
\end{aligned}$$

If $\underline{w} < \overline{w} < 1$, or if $\underline{w} = \overline{w} < 1$ and $d\overline{w} > 0$, or if $\underline{w} < \overline{w} = 1$ and $d\overline{w} < 0$, we have

$$\begin{aligned}
\frac{\partial \chi(\underline{w}, \overline{w})}{\partial \overline{w}} &= \frac{(1-\alpha)f(\overline{w})}{(1-\alpha)(F(\overline{w}) - F(\underline{w})) + \alpha} \left(((1-\alpha)F(\overline{w}) + \alpha)^{n-1} - \chi(\underline{w}, \overline{w})\right); \\
\frac{\partial D}{\partial \overline{w}} &= -\frac{\partial \chi(\underline{w}, \overline{w})}{\partial \overline{w}}(\underline{w} - t); \\
\frac{\partial R}{\partial \overline{w}} &= n(1-\alpha)\left(\chi(\underline{w}, \overline{w}) - ((1-\alpha)F(\overline{w}) + \alpha)^{n-1}\right)\phi(\overline{w})f(\overline{w}) \\
&\quad + n((1-\alpha)(\pi(\underline{w}) - \pi(\overline{w})) + \alpha\pi(t))\frac{\partial \chi(\underline{w}, \overline{w})}{\partial \overline{w}}.
\end{aligned}$$

The proof of the lemma is divided into seven steps.

(i) We claim that $r \leq t \leq \underline{w}$. We can rule out $t < r$ right away, because it violates (7). To rule out $t > \underline{w}$, note that in this case (7) is slack. From the expression of $\partial R/\partial t$, concavity of $\pi(\cdot)$ and the optimality of $\{r, \underline{w}, \overline{w}; t\}$ together imply that $t = r^*$. If $r < \underline{w}$, then since $\underline{w} < t = r^*$, we have $r < r^*$. From the expression of $\partial R/\partial r$, a marginal increase in r would increase (8), contradicting the optimality of $\{r, \underline{w}, \overline{w}; t\}$. Thus, $r = \underline{w}$. If $\underline{w} < \overline{w}$, then from the expression of $\partial R/\partial \underline{w}$, a marginal increase in $\underline{w}$ would increase the revenue, contradicting the assumption of optimality. Thus, $r = \underline{w} = \overline{w} < t = r^*$. From the expressions of $\partial R/\partial \underline{w}$ and $\partial R/\partial \overline{w}$, a increase in $\underline{w}$ and $\overline{w}$ by the same marginal amount would increase the revenue, a contradiction. Thus, $t \leq \underline{w}$.

(ii) We claim that $r < t < \underline{w}$. We can rule out $r = t < \underline{w}$ right away, because it violates (7). To rule out $r < t = \underline{w}$, note that in this case (7) is slack. Since $r < t$, either $r < r^*$ or $t > r^*$, or both. If $r < r^*$, then by raising r marginally, the seller could increase the revenue because $\partial R/\partial r > 0$. If $t > r^*$, then by lowering t marginally, the seller could increase the revenue because $\partial R/\partial t < 0$. Either way, we have a contradiction to the assumption of optimality. Finally, we rule out $r = t = \underline{w}$. If $r = t = \underline{w} < r^*$, then by raising t marginally, the seller relaxes (7), and increases the revenue because $\partial R/\partial t > 0$. If $r = t = \underline{w} > r^*$, then by lowering r marginally, the seller relaxes (7), and increases the revenue because $\partial R/\partial r < 0$. If $r = t = \underline{w} = r^*$, then by lowering r marginally, the seller relaxes (7) because $\partial D/\partial r < 0$, without changing the revenue because $\partial R/\partial r = 0$. With (7) slack, the seller could then increase the revenue by either further raising $\underline{w}$ marginally if $\underline{w} = r^* < \overline{w}$, because $\phi(\underline{w}) = 0$ implies $\partial R/\partial \underline{w} > 0$, or by raising both $\underline{w}$ and $\overline{w}$ by the same infinitesimal amount if $\underline{w} = \overline{w} = r^*$, because $\partial R/\partial \underline{w} + \partial R/\partial \overline{w} > 0$. In each case, we have a contradiction to the assumption of optimality.

(iii) We claim that $r < t < \underline{w} < \overline{w}$. Suppose instead $\underline{w} = \overline{w} = \hat{w}$, and consider decreasing both $\underline{w}$ and $\overline{w}$ by the same marginal amount. We have $\partial D/\partial \underline{w} + \partial D/\partial \overline{w} < 0$, and $\partial R/\partial \underline{w} + \partial R/\partial \overline{w}$ has the same sign as $\pi(t) - \phi(\hat{w})$. Thus, we must have $\pi(t) > \phi(\hat{w})$: otherwise, the seller relaxes (7) without decreasing the revenue, which would then allow the seller to increase the revenue by either raising r or lowering t, as $r < t$ implies $r < r^*$ or $t > r^*$, or both. Since $\phi(1) = 1$, it follows from $\pi(t) > \phi(\hat{w})$ that $\hat{w} < 1$. Consider perturbing the equal priority mechanism by reducing $\underline{w}$ from $\hat{w}$ and raising $\overline{w}$ from $\hat{w}$ such that

$$-(\chi(\hat{w}, \hat{w}) - (1-\alpha)^{n-1} F^{n-1}(\hat{w})) d\underline{w} = (((1-\alpha)F(\hat{w}) + \alpha)^{n-1} - \chi(\hat{w}, \hat{w})) d\overline{w}.$$

By construction,

$$-\frac{\partial \chi(\hat{w}, \hat{w})}{\partial \underline{w}} d\underline{w} = \frac{\partial \chi(\hat{w}, \hat{w})}{\partial \overline{w}} d\overline{w}.$$

This implies that (7) is relaxed, because

$$\frac{\partial D}{\partial \underline{w}} d\underline{w} + \frac{\partial D}{\partial \overline{w}} d\overline{w} = ((1-\alpha)^{n-1} F^{n-1}(\hat{w}) - \chi(\hat{w}, \hat{w})) d\underline{w},$$

which is strictly positive. The seller's revenue is unchanged, because

$$\begin{aligned}\frac{\partial R}{\partial \underline{w}} d\underline{w} + \frac{\partial R}{\partial \overline{w}} d\overline{w} =& n(1-\alpha) f(\hat{w}) \left(\chi(\hat{w}, \hat{w}) - (1-\alpha)^{n-1} F^{n-1}(\hat{w})\right) (\pi(t) - \phi(\hat{w})) d\underline{w} \\ &+ n(1-\alpha) f(\hat{w}) \left(((1-\alpha)F(\hat{w}) + \alpha)^{n-1} - \chi(\hat{w}, \hat{w})\right) (\pi(t) - \phi(\hat{w})) d\overline{w},\end{aligned}$$

which is equal to 0 by construction. The seller could now increase the revenue by either raising r or lowering t, as $r < t$ implies $r < r^*$ or $t > r^*$, or both. This contradicts the assumption of optimality.

(iv) We claim that (7) binds, $r < r^* < t$, and $\pi(t) > \phi(\overline{w})$. If (7) is slack, then since $r < t$ implies that $r < r^*$ or $t > r^*$, or both, the seller could increase the revenue by either raising r or lowering t, a contradiction to the assumed optimality. If $r^* \leq r < t$, the seller could relax (7) by lowering r marginally without decreasing the revenue, which then would allow the seller to increase the revenue by lowering t. Similarly, if $r < t \leq r^*$, the seller could relax (7) by raising t marginally without decreasing the revenue, which then would allow the seller to increase the revenue by raising r. Finally, we show that $\pi(t) > \phi(\overline{w})$. Otherwise, by lowering $\overline{w}$ marginally, the seller relaxes (7) because $\partial D/\partial \overline{w} < 0$, and increases the revenue, as $\partial R/\partial \overline{w}$ has the same sign as

$$\begin{aligned}&\alpha(\pi(t) - \phi(\overline{w})) + (1-\alpha)(\pi(\underline{w}) - \pi(\overline{w})) - \phi(\overline{w})(F(\overline{w}) - F(\underline{w})) \\ =& \alpha(\pi(t) - \phi(\overline{w})) - \int_{\underline{w}}^{\overline{w}} (\phi(\overline{w}) - \phi(w)) f(w) dw \\ <& \alpha(\pi(t) - \phi(\overline{w})),\end{aligned}$$

contradicting the assumed optimality. Note that $\pi(t) > \phi(\overline{w})$ implies $\overline{w} < 1$.

(v) To obtain (10), consider perturbations $d\underline{w}$ and $d\overline{w}$, while keeping r and t unchanged. An optimality condition is that

$$\frac{\partial R}{\partial \underline{w}} d\underline{w} + \frac{\partial R}{\partial \overline{w}} d\overline{w} = 0,$$

for all perturbations $d\underline{w}$ and $d\overline{w}$ satisfying

$$\frac{\partial D}{\partial \underline{w}} d\underline{w} + \frac{\partial D}{\partial \overline{w}} d\overline{w} = 0.$$

Thus we have

$$\frac{\partial R / \partial \underline{w}}{\partial D / \partial \underline{w}} = \frac{\partial R / \partial \overline{w}}{\partial D / \partial \overline{w}}.$$

Using the expressions for $\chi(\underline{w}, \overline{w})$, $\partial \chi(\underline{w}, \overline{w}) / \partial \underline{w}$ and $\partial \chi(\underline{w}, \overline{w}) / \partial \overline{w}$, straightforward algebra lead us to the first-order condition (10) for an optimal equal priority mechanism with respect to $\underline{w}$ and $\overline{w}$. Also, (10) implies that

$$\frac{\partial R / \partial \overline{w}}{\partial D / \partial \overline{w}} = -n(1-\alpha)(\phi(\overline{w}) - \phi(\underline{w})) f(\underline{w}).$$

(vi) To obtain (11), consider perturbations dt and $d\overline{w}$. The optimality condition is

$$\frac{\partial R / \partial t}{\partial D / \partial t} = \frac{\partial R / \partial \overline{w}}{\partial D / \partial \overline{w}}.$$

This gives the first order condition (11) with respect to t and $\overline{w}$.

(vii) Lastly, to obtain (12), consider perturbations dr and $d\overline{w}$, while keeping t and $\underline{w}$ unchanged. The resulting optimality condition is

$$\frac{\partial R / \partial r}{\partial D / \partial r} \geq \frac{\partial R / \partial \overline{w}}{\partial D / \partial \overline{w}},$$

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

$$-\phi(r) f(r) \leq (\phi(\overline{w}) - \phi(\underline{w})) f(\underline{w}).$$

Note that $-\phi(0) f(0) = 1$. Since $\phi(\overline{w}) < \pi(t) < \pi(r^*) < r^*$, and $\underline{w} > t > r^*$,

$$(\phi(\overline{w}) - \phi(\underline{w})) f(\underline{w}) = (\phi(\overline{w}) - \underline{w}) f(\underline{w}) + 1 - F(\underline{w}) < 1.$$

It follows that the optimal r is interior and so (12) holds.

Proof of Lemma 3

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

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

Since $\max\{w-p, 0\}$ is convex in p for any w,

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

Thus, replacing all functions $\{p^\mu_m(\cdot)\}_{m=1}^{n}$ with a single offer p^μ reduces the deviation payoff of an observant buyer. The seller's revenue from unobservant buyers is

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

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

Proof of Theorem 2

Suppose that $\{r, \underline{w}, \overline{w}; t\}$ is an optimal equal priority mechanism. By Lemma 2, the first order conditions (10)-(12) are satisfied. We construct a non-negatively valued multiplier function $\lambda(w)$ for all $w \in [0,1]$ such that the allocative rule $(q^\epsilon_m(v))_{m=0}^{n-1}$ and $(q^\mu_m(v))_{m=1}^{n}$ defined by $\{r, \underline{w}, \overline{w}; t\}$, together with $p^\mu = t$, solves the Lagrangian relaxation. By Lemma 1, the offer rule $(p^\epsilon_m(v))_{m=0}^{n-1}$ we have specified for an equal priority mechanism supports a truthful reporting equilibrium among observant buyers. The conclusion then follows immediately. The proof is divided into four steps.

(i) Construction of the multiplier fucntion. Let $\lambda(w) = 0$ for all $w \notin [\underline{w}, \overline{w}]$, and let

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

for all $w \in (\underline{w}, \overline{w})$, with $\lambda(\underline{w})$ and $\lambda(\overline{w})$ 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 $w \in [\underline{w}, \overline{w}]$ such that $f'(w) \leq 0$. By (10) we have $\phi(\overline{w}) < \pi(t) < \pi(r^*) < r^*$. Since $w \geq \underline{w} > t > r^*$, we have $\lambda(w) > 0$ at any $w \in [\underline{w}, \overline{w}]$ such that $f'(w) > 0$. Thus, $\lambda(w)$ as constructed is non-negative for any w.

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

$$K^{\epsilon}(w) = n(1-\alpha)\phi(w) + \int_w^1 \lambda(x)dx/f(w);$$
$$K^{\mu} = n\alpha\pi(p^{\mu}) - \int_0^1 \lambda(x)\max\{x - p^{\mu}, 0\}dx.$$

We can then rewrite the Lagrangian as

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

where $Q_0^{\epsilon}(w)$ is the probability that an observant buyer with valuation w wins the object when all buyers are observant, and q_n^{μ} is the probability that each unobservant buyer wins the object when all buyers are unobservant.

(ii) We claim that $p^{\mu} = t$ maximizes the Lagrangian. For any $w \in [\underline{w}, \overline{w}]$, by construction

$$\int_w^1 \lambda(x)dx = n(1-\alpha)f(w)(\phi(\overline{w}) - \phi(w)).$$

Using integration by parts, we have

$$\begin{aligned}
&\int_0^1 \lambda(w)\max\{w - p^{\mu}, 0\}dw \\
=&-\int_{\underline{w}}^{\overline{w}} (w - p^{\mu})\, d\left(\int_w^1 \lambda(x)dx\right) \\
=&n(1-\alpha)\left((\underline{w} - p^{\mu})f(\underline{w})(\phi(\overline{w}) - \phi(\underline{w})) + \int_{\underline{w}}^{\overline{w}} f(w)(\phi(\overline{w}) - \phi(w))dw\right) \\
=&n(1-\alpha)\left((\underline{w} - p^{\mu})f(\underline{w})(\phi(\overline{w}) - \phi(\underline{w})) + \phi(\overline{w})(F(\overline{w}) - F(\underline{w})) - (\pi(\underline{w}) - \pi(\overline{w}))\right).
\end{aligned}$$

By (10), we have

$$K^{\mu} = n\alpha\phi(\overline{w}) + n\alpha(\pi(p^{\mu}) - \pi(t)) + (p^{\mu} - t)n(1-\alpha)f(\underline{w})(\phi(\overline{w}) - \phi(\underline{w})).$$

The above is strictly concave in p^{μ}. By (11), it is maximized at $p^{\mu} = t$, with the maximum

$$K_t^{\mu} = n\alpha\phi(\overline{w}).$$

(iii) Comparison of $K^{\epsilon}(\cdot)$ and K_t^{μ}. For $w \in [\underline{w}, \overline{w}]$, we have

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

For all $w > \overline{w}$, since $\pi(\cdot)$ is strictly concave,

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

and so

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

For all $w < \underline{w}$,

$$K^{\epsilon}(w) = n(1-\alpha)\phi(w) + \int_{\underline{w}}^{\overline{w}} \lambda(x)dx/f(w) = n(1-\alpha)(\phi(w) + f(\underline{w})(\phi(\overline{w}) - \phi(\underline{w}))/f(w)).$$

We claim that

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

for all $w < \underline{w}$, and thus $K^{\epsilon}(w) < K^{\epsilon}(\overline{w})$ and

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

To establish the claim, recall that in showing that the constructed multiplier function $\lambda(w)$ is positive for $w \in [\underline{w}, \overline{w}]$, we have proved that $f(w)(\phi(w) - \phi(\overline{w}))$ is strictly increasing in w for all $w \geq \phi(\overline{w})$. This immediately implies that the claim holds for any $w \in [\phi(\overline{w}), \underline{w})$. For

$w < \phi(\overline{w})$, we have

$$f(w)(\phi(w) - \phi(\overline{w})) = f(w)(w - \phi(\overline{w})) - (1 - F(w)) < -(1 - F(w)) < -(1 - F(r^*)),$$

where the last inequality follows because $\phi(\overline{w}) < \pi(t) < \pi(r^*) < r^*$, while

$$f(\underline{w})(\phi(\overline{w}) - \phi(\underline{w})) < f(r^*)\phi(\overline{w}) < f(r^*)r^*,$$

where the first equality comes from $f(w)(\phi(w) - \phi(\overline{w}))$ being strictly increasing in w for all $w \geq \phi(\overline{w})$. The claim then follows from the definition of r^*.

(iv) We claim that the allocations $(q_m^\epsilon(v))_{m=0}^{n-1}$ and $(q_m^\mu(v))_{m=1}^{n}$ specified by $\{r, \underline{w}, \overline{w}; t\}$ are a point-wise maximizer of the Lagrangian relaxation. We disaggregate $Q_m^\epsilon(w)$ and write the Lagrangian as

$$(1-\alpha)^{n-1} \int_0^1 K^\epsilon(w) Q_0^\epsilon(w) f(w) dw + \alpha^{n-1} K_t^\mu q_n^\mu$$
$$+ \sum_{m=1}^{n-1} \mathbb{E}_v \left[\frac{B(m; n-1, \alpha)}{n-m} \sum_{i=1}^{n-m} K^\epsilon(v_i) q_m^\epsilon(\rho_m^i(v)) + B(m-1; n-1, \alpha) K_t^\mu q_m^\mu(v) \right].$$

Fix any realized number m of unobservant buyers such that $1 \leq m \leq n-1$, and consider the last term in the above objective function. Suppose that for some realized valuation profile v we have $v_i > \overline{w}$ for some $i = 1, \ldots, n-m$, but $q_m^\mu(v) > 0$. By (2), we can decrease $q_m^\mu(v)$ marginally by $dq_m^\mu(v) > 0$ and increase $q_m^\epsilon(\rho_m^i(v))$ by $m dq_m^\mu(v)$. Since

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

the effect on the seller's revenue is strictly positive. Therefore, $q_m^\mu(v) = 0$ for any v such that $v_i > \overline{w}$ for some $i = 1, \ldots, n-m$. Further, since $K^\epsilon(w)$ is strictly increasing for $w > \overline{w}$, we have $q_m^\epsilon(\rho_m^i(v)) = 1$ for $v_i = \max\{v_1, \ldots, v_{n-m}\}$. Finally, since

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

for all $w \leq \overline{w}$, with equality if $w \in [\underline{w}, \overline{w}]$, if v is such that $\max\{v_1, \ldots, v_{n-m}\} \leq \overline{w}$, there is a maximizer of the Lagrangian such that $q^\epsilon_m(\rho^i_m(v)) = 0$ whenever $v_i < \underline{w}$, and $q^\epsilon_m(\rho^i_m(v)) = q^\mu_m(v)$ if $v_i \in [\underline{w}, \overline{w}]$.

For $m = 0$ and the first term in the Lagrangian, the strict concavity of $\pi(\cdot)$ implies $K^\epsilon(w)$ for $w < \underline{w}$ crosses 0 at most once and only from below. Thus, for r that satisfies (12), it is point-wise maximizing to set $q^\epsilon_0(\rho^i_0(v)) = 1$ if $v_i = \max\{v_1, \ldots, v_n\}$ and $v_i > \overline{w}$, or if $v_i = \max\{v_1, \ldots, v_n\}$ and $v_i \in [r, \underline{w})$; set $q^\epsilon_0(\rho^i_0(v)) = 1/k$ if $v_i \in [\underline{w}, \overline{w}]$, $\max\{v_1, \ldots, v_n\} \in [\underline{w}, \overline{w}]$ and $\#\{j : v_j \in [\underline{w}, \overline{w}]\} = k$; and set $q^\epsilon_0(\rho^i_0(v)) = 0$ otherwise.

For $m = n$ and the second term in the Lagrangian, it is optimal to set $q^\mu_n = 1/n$ because $K^\mu_t > 0$.

References

[1] M. Akbarpour and S. Li. "Credible auctions: A trilemma." *Econometrica*, 88(2): 425–467, 2020.

[2] E. Ben-Porath, E. Dekel, and B.L. Lipman. "Optimal allocation with costly verification." *The American Economic Review*, 104(12): 3779–3813, 2014.

[3] H. Bester and R. Strausz. "Contracting with imperfect commitment and the revelation principle: The single agent case." *Econometrica*, 69(4): 1077–1098, 2001.

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

[5] R. Denekere and S. Severinov. "Screening when some agents are nonstrategic: does a monopoly need to exclude?" *Rand Journal of Economics*, 37(4): 816–841, 2006.

[6] P.R. Dickson and A.G. Sawyer. "The price knowledge and search of supermarket shoppers." *Journal of Marketing*, 54(3): 251–334, 1990.

[7] K. Hendricks and T. Wiseman. "How to sell (or procure) in a sequential auction market." Technical report, University of Texas - Austin, 2020.

[8] A. Kolotilin, H. Li, and W. Li. "Optimal limited authority for principal." *Journal of Economic Theory*, 148(6): 2344–2382, 2013.

[9] Q. Liu, K. Mierendorff, and X. Shi. "Auctions with limited commitment." *American Economic Review*, 109(3): 976–910, 2019.

[10] P. McAfee. "Mechanism design by competing sellers." *Econometrica*, 61(6): 1281–1312, 1993.

[11] R. Myerson. "Optimal auction design." *Mathematics of Operations Research*, 6: 58–73, 1981.

[12] I. Sher and R. Vohra. "Price discrimination through communication." *Theoretical Economics*, 10: 597–648, 2015.

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

[14] D. Stahl. "Oligopolistic pricing and advertising." *Journal of Economic Theory*, 64(1): 162–177, 1994.

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

[16] G. Virag. "Competing auctions: Finite markets and convergence." *Theoretical Economics*, 5: 241–274, 2010.