UNOBSERVED MECHANISMS

LI, HAO AND MICHAEL PETERS
VANCOUVER SCHOOL OF ECONOMICS

Abstract. A fundamental assumption of mechanism design is that agents are fully aware of commitments that are made by principals, and that they believe that these commitments will be carried out. There is a considerable literature devoted to analyzing what happens when principals can't commit. However, a more relevant consideration for the kinds of mechanism that are used in on line markets is that sellers' can commit, but that not all buyers can observe and understand what these commitments are. In this sense, mechanism design becomes a game with imperfect information in which some agents cannot distinguish different mechanisms offered by principals.

Mechanism designers then face a trade off between the surplus they can extract from informed agents by using commitments, and the surplus they can extract from uninformed buyers from exploiting their ignorance. For instance, a second price auction that is revenue maximizing with complete commitment may not constitute and *equilibrium* for a game of mechanism design. In an auction, the seller commits to a mechanism that makes a price offer to one of the buyers that is contingent on bids he receives. If some buyers do not understand how the bids are being used, the seller might want to deviate from a second price auction to something that looks more like a first price auction in order to extract more surplus from the uninformed bidders.

In this preliminary version of the paper, we consider a two value auction and show that if the probability with which each bidder is uninformed about the mechanism the seller is using is very low, then the standard optimal auction constitutes an *equilibrium* in the sense that them mechanism designer has no incentive to deviate to any alternative mechanism. We also characterize the equilibrium mechanism in the case where buyers are more likely to be uninformed. In this case, the sellers' best mechanism has the property that high value uninformed bidders send randomized messages. These messages are designed to make deviations by the seller unprofitable. They have the unfortunate side effect that the resulting equilibrium cannot be ex post efficient.

We also explore the case with a continuum of types. We show that sellers can design mechanisms with extended messages that allow informed bidders to reveal that they are informed to the seller. The seller can then offer them standard mechanisms with an added incentive constraint - all informed bidders prefer to reveal that they are informed (uninformed bidders cannot pretend to be informed in this mechanism). We give a partial characterization of the equilibrium mechanism for this case, and show that it cannot be ex post efficient.

Airline tickets are sold using dynamic pricing algorithms that use the time that buyers purchase tickets as signals of their type (high value business class vs lower value leisure travelers). It is widely speculated that these price offers also depend on buyers on line search behavior. For example, it is trivial for the seller's website to record whether or not a particular buyer has visited previously. Though buyers might not understand exactly what their messages are, they know they are being used to determine prices. Even so, most buyers are unsure exactly how their messages are being used.

Of course, experienced flyers are more likely to understand the pricing rules. The computer engineers and economists who designed these rules, are also airline travelers, so they understand how these mechanisms work.

What is unusual about this situation is not that there is no commitment - the price offers are made by computer programs that cannot easily be modified once they are running. The unusual part is that many, but not all buyers, cannot see exactly what commitments are built into those programs. The purpose of this paper is to explore the implications of this for the most standard mechanism design problems. Here we mostly confine our attention to independent private value auctions.

One variant of our story concerns the situation in which the message space is just a set of bids. All buyers understand this, but some buyers cannot see directly how the bids are being used to determine allocations. For example, if these bidders believe that the seller is using a second price auction to allocate, they will bid their values. A seller who understands this will be tempted to write a program offers the good to the highest bidder, but offer a price equal to the bid this bidder submitted. This would allow the seller to extract more surplus from buyers who incorrectly believed the seller was using a second price auction.

A first price auction actually works no better. If all buyers believe the seller is using a first price auction, the seller can work out the value that each bidder must have by observing his or her bid. Then the seller would offer the high bidder a trade with price equal to that value instead of the actual bid. Informed buyers would understand this an reduce their bids. It is this trade off that we want to understand.

It isn't hard to show that if no bidder directly observe the mechanism the seller is using, then the only equilibrium mechanism is one in which the seller randomly chooses a buyer and offers the buyer a fixed price.[1] This is equivalent to the case where sellers can't commit. If all buyers can observe the sellers' mechanism then we are back to standard mechanism design. It is the intermediate case that interests us.

In the two value case, we show that the equilibrium mechanism depends on the probability with which bidders are uninformed. If this probability is low enough, the standard optimal auction is an equilibrium mechanism. As this probability rises the optimal auction breaks down. Uniformed bidders begin to offer messages that no longer fully reveal their type, though informed bidders continue to bid truthfully. As the probability that bidders are uninformed increases, the messages uninformed bidders send become less and less informative. In the limit the mechanism collapses to a fixed price.

In this preliminary version of the paper we provide a result like this for the two value case.

We then consider the case in which the seller can use a randomized mechanism that makes it possible for the informed buyer to reveal that he is informed to the seller. Certain messages are treated differently by the seller, but only informed buyers are able to observe what these messages are. The messages are then partially

[1]This conclusion must be modified if the seller can make offers to other bidders if this initial offer is refused.

verifiable in that a buyer can prove to the seller that he is informed, but he can't prove to the seller that he is uninformed. The seller wants to offer the informed buyer and optimal auction in this case, but faces a new constraint in that each informed buyer type who is expected to reveal that he is informed must prefer to do this instead of pretending to be uninformed. To suppose this, uninformed buyers must send randomized messages. As a consequence, two new inefficiencies are introduced. First, in cases where all buyers are informed, the seller incorrectly uses the randomized messages to choose an allocation. Second, some low value informed buyers will never trade in the revenue maximizing auction. As a consequence, these buyers will prefer to pretend to be uninformed since they will have a chance to trade if other traders are uninformed.

1. A two-value model

There is a single seller with one indivisible good for sale. The seller's reservation value for the good is zero. There are two ex ante identical bidders with independent private values for the good. Each bidder has the high value v_H with probability p_H, and the low value $v_L \in (0, v_H)$ with probability $p_L = 1 - p_H$. To make the problem interesting, we assume that

$$v_L > p_H v_H.$$

Each bidder is independently uninformed about the seller's mechanism with probability μ. The value of μ is common knowledge between the two bidders and the seller. Only the bidder himself knows whether or not he is informed about the seller's mechanism.

The timing of the "unobservable mechanism design game" is as follows. The seller first commits to a mechanism. With probability $1 - \mu$ each bidder learns the mechanism, and with the remaining probability μ he remains uninformed. The mechanism is then played by the two bidders. The set of available pure actions for a bidder consists of two reports, v_H or v_L, regardless of whether his value is v_H or v_L, or whether he is informed or uninformed, and if he is informed, what the seller's actual mechanism is. We allow the bidder to randomize between the two reports. Given a realized profile of two reports, the seller's mechanism is implemented and the game ends.

We restrict the seller's feasible commitments to the following direct mechanisms. Each bidder's message space contains two elements: a report of v_H and a report of v_L, and the seller's mechanism maps any two reports to a selection between the two bidders and a single take-it-or-leave-it offer to the selected bidder. Denote the offers as y_{HH} when both bidders report v_H, y_{HL} when the reports are v_H and v_L, and y_{LL} when both reports are v_L. Since the two bidders are ex ante identical, we require the seller's mechanism to treat them symmetrically. The selection is allowed to be stochastic and so are the offers.

We have imposed two restrictions on feasible mechanisms that are potentially with loss. First, the bidder's message space is restricted to the standard type space. However, in our problem of unobserved commitment, at the interim stage bidders can differ in their knowledge about the seller's mechanism as well as their value. Second, the seller is restricted to a single take-it-or-leave-it offer. However, in standard problems with observed commitment but without bidding, the seller generally benefits from making a sequence of take-it-or-leave-it offers instead of a single one.

If $\mu = 1$, then the seller cannot do better than randomly selecting a bidder with a take-it-or-leave-it offer of v_L. This gives the revenue of v_L.

If $\mu = 0$, we have a standard mechanism design problem with observable commitments. The following direct mechanism is an optimal "auction." Each bidder is asked to report either v_H or v_L. If both report v_H, the seller randomly selects one and makes a take-it-or-leave-it offer of v_H. Similarly, if both report v_L, the seller randomly selects one with a take-or-leave-it offer of v_L. If one bidder reports v_H and the other reports v_L, the seller makes an offer of $\frac{1}{2}(v_H + v_L)$ to the former. The seller's revenue is given by

$$p_H^2 v_H + 2p_H p_L \frac{1}{2}(v_H + v_L) + p_L^2 v_L = p_H v_H + p_L v_L.$$

A low-value bidder's payoff is zero. A high-value bidder's payoff is positive, and is determined by the indifference between reporting v_H and reporting v_L:

$$p_H(v_H - v_H) + p_L\left(v_H - \frac{1}{2}(v_H + v_L)\right) = p_L \frac{1}{2}(v_H - v_L).$$

Revenue equivalence in discrete-value auctions implies that any two offers $y_{HH} \in [v_L, v_H]$ and $y_{HL} \in [v_L, v_H]$, together with $y_{LL} = v_L$, are optimal if they satisfy the above indifference condition of the high-value bidder.

2. Optimal mechanism with unobservability

For any $\mu > 0$, we say that a mechanism is "optimal" if it is an equilibrium mechanism, and if there is no other equilibrium mechanism that gives the seller a higher revenue. We use perfect Bayesian equilibrium as the solution concept. This implies that an uninformed bidder on the equilibrium path "knows" the mechanism that the seller has committed to just as an informed bidder does. However, off the equilibrium path, the seller's deviations are unobserved by an uninformed bidder.

First, we show that when the probability of unobservability μ is sufficiently low, the optimal auction given above for $\mu = 0$ remains optimal. In this case, the seller has no incentive to deviate from the optimal auction because the loss in revenue from informed bidders overweighs the gain from uninformed bidders through changing the optimal auction.

Proposition 1. *If $\mu \leq 1/(1+\sqrt{p_L})$, then the optimal auction for $\mu = 0$ is optimal.*

Proof. [Sketch of proof] The proposition follows from two claims.

First, we claim that if $\mu \leq 1/(1 + \sqrt{p_L})$, the optimal auction for $\mu = 0$ is an equilibrium mechanism. That is, there is an equilibrium where the seller's strategy is the optimal auction for $\mu = 0$. We establish this claim by constructing the rest of the equilibrium. On the equilibrium path, each bidder reports his value truthfully with probability one, regardless of whether he is uninformed or informed. If he is uninformed, truthful reporting is his equilibrium strategy. If he is informed, then after any deviation by the seller, the reporting strategy of the bidder forms a continuation equilibrium given the deviating mechanism and given truthful reporting by any uninformed bidder. The revenue-maximizing deviation by the seller is to set $\tilde{y}_{HH} = \tilde{y}_{HL} = v_H$, together with $\tilde{y}_{LL} = v_L$, so as to extract all surplus from the high-value uninformed bidder. (The seller could also exclude low value bidder in deviation by making a take-it-or-leave-it offer of v_H, but by the assumption of $v_L > p_H v_H$, this revenue from this deviation is less than the revenue from the optimal auction for $\mu = 0$.) In the continuation equilibrium, the high-value informed bidder reports v_L with probability one. The maximum deviation revenue is

$$(1 - (1 - p_H\mu)^2)v_H + (1 - p_H\mu)^2 v_L.$$

So long as $\mu \le 1/(1+\sqrt{p_L})$, the above does not exceed the revenue from the optimal auction for $\mu = 0$. This implies that the seller has no incentive to deviate from the optimal auction for $\mu = 0$, and thus we have constructed an equilibrium.

Second, we claim that there is no equilibrium in which the seller obtains a higher revenue than the revenue from the optimal auction for $\mu = 0$. We establish this claim by arguing that any equilibrium revenue in an unobservable mechanism design game with some $\mu > 0$ is attainable by some mechanism with $\mu = 0$. This is because in the continuation game after the seller commits to the equilibrium mechanism, the uninformed bidder "knows" the seller's mechanism. The revelation principle does not immediately apply because the uninformed bidder and the informed bidder can have different reporting strategies, but uninformed and informed bidders have the same incentive constraints and their reporting strategies have the same impact on the seller's revenue. This means that the revenue in the continuation equilibrium is attainable by an incentive compatible mechanism in the standard observable commitment environment. □

To extend the above result to beyond the two-type model, we can try to establish the two claims in the proof of Proposition 1 more generally. The first claim, which establishes the maximal deviation revenue given the optimal auction with observable commitments, involves an interesting mixed "behavioral" mechanism design problem: with some probability μ each bidder is behavorial and has a reporting strategy fixed at the equilibrium of the optimal auction with observable commitments, while with the remaining probability the bidder is rational and requires the deviation mechanism to satisfy the standard incentive constraints. The restriction to direct mechanisms implies that in deviation the seller is not allowed to use self-selection to induce the bidders to reveal whether or not they are informed. The second claim, which establishes that unobservability cannot benefit the seller, is a variant of the usual revelation-principle type of argument. Although the seller may not be able to replicate a given continuation equilibrium on the path of an unobservable commitment game with a truth-telling equilibrium of a standard direct mechanism, it is always possible to replicate the revenue.

Next, for each value of μ that exceeds the bound given in Proposition 1, we modify the optimal auction for $\mu = 0$ and show that it is optimal. As the optimal auction for $\mu = 0$, the mechanism we construct selects each bidder with probability one half when both report v_H and when both report v_L, with offers y_{HH} and y_{LL} respectively, and selects the bidder who alone reports v_H with probability one, with offer y_{HL}. Also, $y_{LL} = v_L$. In the continuation game on the path, a low value bidder reports v_L with probability one, regardless whether he is uninformed or informed. Unlike in the optimal auction for $\mu = 0$, the modified auction induces an uninformed high value bidder to randomize between reporting v_H and reporting v_L, while informed high value bidder reports v_H with probability one.

Proposition 2. *For any $\mu > 1/(1 + \sqrt{p_L})$, there exists a modified auction that is optimal.*

Proof. [Sketch of proof] There are two steps.

First, we show that with appropriate choices for y_{HH} and y_{HL}, the modified auction is an equilibrium mechanism. Let r^u be the probability of reporting v_H by an uninformed high value bidder on the equilibrium path. Given that low value bidder reports v_L with probability one and an informed high value bidder reports v_H with probability one, we can define

$$\overline{r} \equiv p_H(\mu r^u + 1 - \mu)$$

as the ex ante probability that a bidder reports v_H. Then, for any given $r^u \in [0, 1]$, there exist $y_{HH}, y_{HL} \in [v_L, v_H]$ such that the incentive constraint of a high value bidder binds:

$$\overline{r}\frac{1}{2}(v_H - y_{HH}) + (1 - \overline{r})(v_H - y_{HL}) = (1 - \overline{r})\frac{1}{2}(v_H - v_L).$$

Using this constraint, we can write the seller's revenue as

$$\overline{r}^2 y_{HH} + 2\overline{r}(1 - \overline{r})y_{HL} + (1 - \overline{r})^2 v_L = \overline{r}v_H + (1 - \overline{r})v_L.$$

For any given probability r^u that an uninformed high value bidder reports v_H, since an informed high value bidder already reports v_H with probability one on the equilibrium path, the revenue-maximizing deviation by the seller is to charge $\tilde{y}_{HH} = \tilde{y}_{HL} = v_H$. (The seller could also exclude low value bidder in deviation by making a take-it-or-leave-it offer of v_H, but the revenue from this deviation is less than the equilibrium revenue

$$(1 - (1 - \overline{r})^2)v_H < \overline{r}v_H + (1 - \overline{r})v_L,$$

due to the assumption of $v_L > p_H v_H$.) An informed high value bidder will report v_L with probability one. The maximum deviation revenue is thus

$$(1 - (1 - p_H\mu r^u)^2)v_H + (1 - p_H\mu r^u)^2 v_L.$$

Thus, any r^u such that

$$1 - p_H(\mu r^u + 1 - \mu) \le (1 - p_H\mu r^u)^2$$

makes the corresponding modified auction an equilibrium mechanism. The left-hand side is decreasing and linear in r^u, while the right-hand side is decreasing and convex in r^u. At $r^u = 0$, the inequality always holds; at $r^u = 1$, it is violated because $\mu > 1/(1+\sqrt{p_L})$. Thus, there is a unique r^u_* such that the seller is indifferent between the equilibrium mechanism and the revenue-maximizing deviation. Any $r^u \in [0, r^u_*]$ induces an equilibrium mechanism, and the seller's revenue is the highest with r^u_*.

Second, we argue that there is no other equilibrium mechanism that gives the seller a higher revenue than that corresponding to r^u_*. In any equilibrium, high value bidder's incentive constraint must bind; otherwise, the seller would want deviate by raising the price offers. Given this, the seller's equilibrium revenue has a one-to-one relation with the ex ante probability $\overline{r}$ that a bidder reports v_H, given by

$$\overline{r} = p_H(\mu r^u + (1 - \mu)r^i),$$

where r^i is the probability that an informed high value bidder reports v_H. We claim that under any optimal mechanism, on the equilibrium path either $r^u = 0$ or $r^i = 1$, or both. For if $r^u > 0$ and $r^i < 1$, the seller can construct another equilibrium mechanism, with r^u slightly lower and r^i slightly higher, such that $\overline{r}$ remains unchanged. The incentive constraint of high value bidder remains unchanged for the same offers y_{HH} and y_{HL}, and so is the seller's equilibrium revenue, but the maximum deviation revenue is reduced. As a result, the seller can further modify the mechanism to increase $\overline{r}$ and thus to increase the revenue, contradicting the assumption of optimality. Since we have already shown that any modified auction with $r^u \in [0, r^u_*]$ is an equilibrium, the optimal mechanism has $r^u = r^u_*$ and $r^i = 1$. □

The proofs for Proposition 1 and Proposition 3 can be unified by considering the mixed behavioral mechanism design problem mentioned above. Fix any μ and any $r^u \in [0, 1]$, imagine that the seller chooses $r^i \in [0, 1]$ to maximize the revenue, subject to the indifference condition of high value bidder if $r^i > 0$. Then, the argument in the proof of Proposition 3 implies that the optimal r^i is either 1, in order to maximize the ex ante probability $\overline{r}$, or 0, in order to dispense with the indifference condition of high value bidder. Comparing these two options then leads to the conclusions of Proposition 1 and Proposition 3. If for some μ and r^u the solution is $r^i = 1$, then there is a modified auction with these values of μ and r^u that is an equilibrium mechanism; otherwise, this value of r^u cannot be supported in an equilibrium mechanism for the given value μ. For $\mu \leq 1/(1 + \sqrt{p_L})$, we can support $r^u = 1$ in an equilibrium mechanism, which is the result of Proposition 1. As we increase the value μ, there is a maximum value of r^u, which is r^u_* given in the proof of Proposition 3, that can be supported in an equilibrium mechanism.

In the optimal modified auction, the offers y_{HH} and y_{HL} are not unique. Any $y_{HH}, y_{HL} \in [v_L, v_H]$ that bind the incentive constraint of high value bidder corresponding to r^u_* are optimal. We have seen from the proof of Proposition 3 that r^u_* is uniquely determined by

$$1 - p_H(\mu r^u_* + 1 - \mu) = (1 - p_H \mu r^u_*)^2.$$

It is straightforward to show that as μ increases, r^u_* increases but both μr^u_* and $p_H(\mu r^u_* + 1 - \mu)$ decrease. Thus, high value bidder's payoff increases and the seller's optimal revenue decreases as μ increases. As μ converges to 1, the seller's optimal revenue goes to v_L.

We have restricted the seller to direct mechanisms both on and off the equilibrium path. An implication of this restriction is that the seller cannot separate informed bidders from uninformed ones through additional messages. If this restriction is lifted off the path, the seller can improve the deviation payoff for a given incentive compatible mechanism on the path. This would change the characterization of the equilibrium mechanisms given in Proposition 1 and Proposition 3.

Indeed, it is easy to see that if the seller is allowed to other messages for self-selection of informed bidders, the optimal auction can never be an equilibrium mechanism with observability, in contrast to Proposition 1. For example, in addition to the two messages understood by an uninformed bidder, the seller can create one more message for an informed high value bidder to reveal himself. If no bidder reveals himself to be informed with v_H, the seller extracts all surplus from the high-value uninformed bidder when there is one, as in the revenue-maximizing deviation in the proof of Proposition 1; otherwise, the price offers are such that the informed high value bidder is indifferent between revealing himself and pretending to have v_L. An informed high value bidder can of course pretend to be uninformed, but since the surplus of an uninformed high value bidder is fully extracted, there is no such incentive. The seller's maximal revenue with self-selection off the path can be shown to be

$$(\mu p_H p_L + p_H)v_H + ((1-\mu)p_H p_L + p_L^2)v_L,$$

which is always higher than the revenue $p_H v_H + p_L v_L$ from the optimal auction with observability.

On the equilibrium path of any mechanism, since uninformed bidders correctly understands the equilibrium mechanism, it is without loss to assume that there are two messages, v_H and v_L that are used with positive probability. Thus, any equilibrium mechanism is outcome-equivalent to a modified auction even when the seller is allowed to use additional messages off the path. As argued in the proof of Proposition 3, in any equilibrium mechanism, informed high value bidder reports v_H with probability one, that is, $r^i = 1$, while uninformed bidders may randomize, with probability r^u of reporting v_H. The equilibrium revenue is $\overline{r}v_H + (1-\overline{r})v_L$, where $\overline{r} = p_H(\mu r^u + 1 - \mu)$ is the ex ante probability that a bidder reports v_H on the equilibrium path, so it is increasing in r^u. Without additional messages off the path, the seller's optimal mechanism is given by r^u_*. The following result establishes that if the seller can use additional messages for informed bidders to signal themselves off the path, the only equilibrium mechanism is the modified auction with $r^u = 0$ regardless of μ.

Proposition 3. *If the seller can use additional messages off the path, then for any μ, any equilibrium mechanisms is outcome-equivalent to the modified auction with $r^u = 0$.*

Proof. [Sketch of proof] Fix any $\mu \in (0,1)$. Given that the seller can use additional messages off the path, for any $r^u \in [0,1]$ on the equilibrium path, the revelation principle implies that it is without loss to consider truth-telling equilibria in the mixed behavioral mechanism design problem off the path. As in the standard optimal auction problem, the good is assigned to the bidder who alone reports v_H, regardless of whether this bidder and the other v_L-reporting bidder reveal themselves to be informed or uninformed. Denote the price offer as $y_{H^k L^{k'}}$, where $k, k' \in \{u, i\}$ denote respectively whether the bidder that reports v_H and the bidder that reports v_L are uninformed or informed. Clearly, in any solution to the problem, we have $y_{H^u L^k} = v_H$ for $k \in \{u, i\}$. Also, we can assume that the seller uses an even lottery to assign the good when both bidders report v_L, regardless of whether they reveal themselves to be informed or uninformed, and we can set the price offer $y_{L^k L^{k'}} = v_L$ for any $k, k' \in \{u, i\}$. Finally, revenue equivalence can be shown to apply to this problem, so without loss we assume that the bidder revealing himself to be informed and having v_H gets the good with probability one against a bidder revealed to be uninformed and having v_H; denote the price offer as $y_{H^i H^u}$. (Revenue equivalence holds because whenever the seller assigns the good to the uninformed high value bidder the price offer will be v_H, implying that the same revenue is obtained as if the seller instead assigns the good the informed high value bidder and charges $y_{H^i H^u} = v_H$, with appropriate adjustments to other price offers so as to keep the incentive compatibility constraint of the high value informed bidder binding.)

The incentive compatibility constraint for the high value informed bidder is

$$\begin{aligned}
&(1-\mu)\left(p_H\frac{1}{2}(v_H - y_{H^iH^i}) + p_L(v_H - y_{H^iL^i})\right)\\
&\quad + \mu\Big(p_H r^u(v_H - y_{H^iH^u}) + (p_L + p_H(1-r^u))(v_H - y_{H^iL^u})\Big)\\
=&(1-\mu)p_L\frac{1}{2}(v_H - v_L) + \mu(p_L + p_H(1-r^u))\frac{1}{2}(v_H - v_L).
\end{aligned}$$

This can always be satisfied by appropriately choosing the price offers $y_{H^iH^i}, y_{H^iL^i}, y_{H^iH^u}, y_{H^iL^u} \in [v_L, v_H]$ appropriately. If the high value informed bidder reports truthfully, his total probability of getting the good is

$$1 - (1-\mu)p_H\frac{1}{2},$$

while if he deviates and reports v_L, he gets the good at the price of v_L with total probability

$$p_L\frac{1}{2} + \mu p_H(1-r^u)\frac{1}{2}.$$

The revenue is given by

$$\begin{aligned}
&(1-\mu)^2\Big(p_H^2 y_{H^iH^i} + 2p_Hp_L y_{H^iL^i} + p_L^2 v_L\Big)\\
&+2(1-\mu)\mu\Big(p_H r^u(p_H y_{H^iH^u} + p_L v_H) + (p_L + p_H(1-r^u))(p_H y_{H^iL^u} + p_L v_L)\Big)\\
&+\mu^2\Big((1-(p_L + p_H(1-r^u))^2)v_H + (p_L + p_H(1-r^u))^2 v_L\Big).
\end{aligned}$$

Substituting in the binding incentive compatibility constraint for the high value informed bidder, we rewrite the revenue as

$$\begin{aligned}
&2(1-\mu)p_H\left(\left(1-(1-\mu)p_H\frac{1}{2}\right)v_H - \left(p_L\frac{1}{2} + \mu p_H(1-r^u)\frac{1}{2}\right)(v_H - v_L)\right)\\
&+v_H\Big(2(1-\mu)\mu p_L p_H r^u + \mu^2(1-(p_L + p_H(1-r^u))^2)\Big)\\
&+v_L\Big((1-\mu)^2 p_L^2 + 2(1-\mu)\mu p_L(p_L + p_H(1-r^u)) + \mu^2(p_L + p_H(1-r^u))^2\Big).
\end{aligned}$$

The above is strictly greater than the equilibrium revenue

$$(p_H(\mu r^u + 1 - \mu))v_H + (1 - p_H(\mu r^u + 1 - \mu))v_L$$

for any $r^u > 0$, and they are equal when $r^u = 0$. □

The idea of the above result is simple. For any $r^u > 0$, in the revenue-maximizing deviation mechanism the payoff to an informed high value bidder is pinned down by the probability of getting the good if the bidder lies and reports v_L through the binding incentive compatibility constraint. The probability is the same on the equilibrium path, and thus the payoff is the same. However, unless $r^u = 0$, uninformed high value bidders lose part of their payoff to the seller.

The equilibrium revenue is equal to $p_H(1-\mu)v_H + (1 - p_H(1-\mu))v_L$. This converges to the take-it-or-leave-it offer when μ goes to 1, and converges to the revenue from the unconstrained optimal auction as μ goes to 0.

For any $\mu \in (0, 1)$, since the revenue from the modified auction is increasing in $\overline{r}$, and the payoff of high value bidders is decreasing in $\overline{r}$, compared to the case where the seller is not allowed to use additional messages off the path, the seller is worse off and high value bidders are better off. The equilibrium behavior of uninformed high value bidders is simpler, which may indicate the model with additional new messages off the path is more natural.