Manipulative Auction Design*

Philippe Jehiel†

21st September 2007

Abstract

This paper considers a manipulative auction design framework in which the designer in addition to choosing the auction format(s) is free to choose how much information feedback about the distribution of bids previously observed in similar auctions to report to bidders. A feedback equilibrium is proposed to model the long run interactions of bidders in such environments with partial feedback. In one-object private values auctions with independent distributions of types across bidders, it is shown that a suitable choice of manipulative auction design always allows the seller to extract more revenues than in the classic optimal nonmanipulative auction (Myerson, 1981). It is also shown how the first-price auction in which bidders get only to know the aggregate distribution of bids (across all bidders) generates more revenues than the second-price auction while achieving an efficient outcome in the asymmetric two-bidder case.

1 Introduction

Standard equilibrium approaches of games with incomplete information (à la Harsanyi) assume that players know the distributions of signals held by other players as well as these players' strategies as a function of their signals (see Harsanyi (1995)). Yet, this requires

*I would like to thank I. Esponda, L. Lamy and R. Spiegler for useful comments made at various stages of this project.

†PSE, 48 boulevard Jourdan, 75014 Paris, France and University College London ; jehiel@enpc.fr

a lot of knowledge that need not be easily accessible to players. Modern approaches to equilibrium rely on learning to justify this knowledge (see Fudenberg and Levine (1998)). But, a question arises as to whether enough information feedback is available to players at the learning stage for convergence to equilibrium to be reasonably expected.

For example, consider a series of first-price auctions in an asymmetric private values setting involving each time two new bidders of observable characteristics $i = 1, 2$. That is, each bidder when choosing his bid knows his valuation but not that of the other bidder, and the distributions of valuations which are *a priori* unknown to the bidders are assumed to depend on the characteristic i.[1] Assume further that along the process, bidders are only informed of the aggregate distribution of bids in past auctions without being informed of the characteristics of the bidders of the corresponding bids.[2] It seems then dubious that bidders will be able to play a best-response to the actual distribution of bids of the other bidder, even in the long run because there is no way a bidder can assess the distribution of bids conditional on the characteristic based on the feedback he receives.[3] Instead, bidders $i = 1, 2$ are more likely to play a best-response to the aggregate distribution of bids that mix the distribution of bids of all bidders whether of characteristic $i = 1, 2$.

In this paper, we consider one-object private values auction environments in which the valuations are independently distributed across bidders. We first propose an equilibrium concept to describe the long run interaction of bidders in situations in which (at the learning stage) bidders receive coarse feedback about the distribution of previous bids, and every single bidder participates in just one auction. This allows us to consider situations in which as described above only the aggregate distribution of bids with no reference to the characteristics of bidders is observed at the learning stage, and also situations in which different auction formats are being used and only the aggregate distribution of bids across the various auction formats is being revealed at the learning stage. The equilibrium obtained (which stands for the limiting outcome of the corresponding learning process) is called a feedback equilibrium (it is parameterized by the form of the feedback described as a partition of the set of profile of format and bidders' characteristics received

[1] The asymmetry should be interpreted to mean that the observable characteristic i is informative of the distribution from which the valuation is drawn.

[2] This looks like the norm rather than the exception in many auction setups.

[3] In the language of econometrics, the model is not identifiable.

by the players at the learning stage, see Section 2), and it requires that bidders play a best-response to the aggregate distribution of bids, as given by their feedback.

The feedback equilibrium is closely related to the analogy-based expectation equilibrium (ABEE) introduced in Jehiel (2005) and further developed in Jehiel and Koessler (2007) and Ettinger and Jehiel (2007).[4] We note that the experiment reported in Huck et al. (2007) gives support to the ABEE as a good description of long run behaviors in situations in which past behaviors are not cleanly separated between games. In that experiment, players repeatedly played one of two games with very different best-response structures - this is the analog of bidders being involved in different auction formats in our setting. In one treatment, these players received only feedback about the aggregate distribution of actions of the subjects assigned to the role of their opponent over the two games in the last five rounds.[5] Convergence to the corresponding ABEE (which differed from Nash equilibrium) was mostly observed in this case, thereby suggesting that subjects did best-respond to the aggregate distribution of actions over the two games despite the fact that the games under study had very different best-response structures.

The novel perspective adopted in this paper as compared with the analogy-based expectation equilibrium is that the feedback given to bidders is viewed as a design issue. If the same seller repeatedly sells similar objects to bidders with similar characteristics, this seller is not only free to choose the auction format she likes best, but she should also be able to choose the form of feedback she wishes to provide to bidders. We impose the mild constraint that the feedback should be correct (even if coarse) and that some outside authority can control and enforce it.[6] Other than that, the designer is free to choose the feedback (partition) she likes best.

The main question we ask is as follows. Can the designer generate more revenues than in the optimal auction in which bidders are assumed to play a Nash Bayes equilibrium? We show that this is always so no matter what the distributions of valuations are. That is,

[4] It can also be viewed as a selection of self-confirming equilibrium (the conjecture adopted by the players is the simplest theory consistent with their feedback).

[5] They did not learn about their performance until the end of the experiment, which should be related to our implicit assumption that individual bidders participate in just one auction.

[6] If the feedback were not constrained to be correct, it sounds dubious that bidders would pay attention to it. It may then be in the designer's interest to have an outside authority that guarantees the correctness of the feedback.

by providing partial feedback about the distribution of bids and/or by proposing several auction formats, the designer can always (i.e. whatever the distributions of valuations) achieve an expected revenue that is strictly larger than the revenues generated in the classic optimal auction (Myerson (1981) or Riley and Samuelson (1981)). Interpreting the provision of partial feedback as design manipulation, our result shows that there is always scope for design manipulation when the designer seeks to maximize revenues.[7]

We also address another question. Assume the primary goal of the designer is welfare maximization whereas the auxiliary goal is revenue. In the classic setup, the so called revenue equivalence theorem asserts that the designer can do not better than using a second-price auction (the second-price auction induces an efficient outcome and any efficient mechanism that respects the participation constraints of bidders must achieve a revenue no greater than that of the second-price auction). In our manipulative auction design setup, we show that the designer can sometimes do better. Specifically, in the case of two bidders with asymmetric distributions of valuations, we show that the first-price auction in which the designer provides as feedback the aggregate distribution of bids with no reference to the identity of the bidders always induces an efficient outcome and always generates an expected revenue that is strictly larger than that of the second-price auction no matter what the distributions of valuations are. Our clear-cut revenue comparison should be contrasted with the ambiguous revenue ranking between the first-price auction and the second-price auction obtained in the standard rationality case (see Maskin and Riley (2000)).[8] The result also illustrates that the so called revenue equivalence result does not hold in our setup with manipulative auction design. Two mechanisms that result in the same allocation rule (in equilibrium) and that give the same zero expected utility

[7]It may be mentioned here that the designer need not know the distribution of valuations of the bidders to start with. Her maximization exercise may be viewed as the outcome of a trial and error process where the designer would be assumed to keep track of the past performance of her previously tried auction designs. It may also be mentioned that unlike the bidders, the designer has access to the entire distribution of bids and as such can much more easily have access to the distribution of valuations. For example, by using second-price auctions and observing the distribution of bids according to the characteristic of the bidder, the designer could in a first stage learn these distributions. Based, on this knowledge, she could then use the feedback she likes best, as analyzed in this paper.

[8]Hafalir and Krishna (2007) also obtain a clear-cut revenue comparison for the two bidder case in the standard rationality paradigm when a resale market operates after the first-price auction. Observe that in our coarse feedback treatment, the outcome of the first-price auction is efficient and there is no scope for resale.

to the bidders with lowest valuations need not result in the same expected revenue.

While considering the information feedback as a design issue is new (see however Esponda (2007) which is discussed in Section 3), the empirical literature on auctions has been concerned with the related issue of when the available information on bids and valuations allows the researcher to identify the model assuming bidders play a Nash equilibrium of the corresponding auction model (see Athey and Haile (2006)). Somehow the approach developed here assumes that bidders themselves need not have a complete understanding of the game they play, and as a result bidders play a feedback equilibrium as opposed to a Nash equilibrium.

The rest of the paper is organized as follows. In Section 2 we describe the basic definitions of our manipulative auction design framework. In Section 3 we expand on the learning interpretation of the solution concept. In Section 4 we provide some preliminary analysis. The core results of the paper are contained in Section 5. Section 6 contains further results in particular about the complete information case and the incentive for the designer to use shill bidders. Section 7 suggests some further applications of the approach and some avenues for future research. Missing proofs can be found in the Appendix.

2 Basic definitions

There is one object for sale and n bidders with characteristics $i \in I = \{1, ...n\}$. Each bidder i knows his own valuation v_i for the object, but not that of the other bidders $j \neq i$. The distribution of valuations are independent across bidders. The valuation v_i is drawn from a distribution with support $[L, L+1]$ and (continuous) density $f_i(\cdot)$ where we assume that $f_i(v) > 0$ for all $v \in [L, L+1]$ and $L \geq 0$. Bidders do not *a priori* know the densities $f_i(\cdot)$ (not even their own $f_i(\cdot)$).

The auctioneer may use multiple auction formats M_k, $k \in K = \{1, ...r\}$ to sell his object where auction format M_k, $k = 1, ...r$, is selected with probability λ_k. The rules of the auction format M_k are announced before bidders take actions in M_k, and each auction format M_k takes the following form:

- Bidders $i = 1, ...n$ simultaneously submit a bid $b_i \in [0, B]$.

- Based on the profile of bids $b = (b_i)_{i=1}^{i=n}$ bidder i wins the object with probability $\varphi_i^k(b)$ and pays a transfer $\tau_i^k(b)$ to the auctioneer.

- Any bidder i who bids $b_i = 0$ makes no payment, i.e. $\tau_i^k(b) = 0$ whenever $b_i = 0$.

We assume that bidders have quasi-linear preferences and that they are risk neutral. That is, if a bidder with valuation v expects to win the object with probability p and expects to make (an expected) transfer t to the designer, his expected utility is $pv - t$. Accordingly, in format M_k, if bidder i with valuation v_i bids b_i and expects the bid profile $b_{-i} = (b_j)_{j \neq i}$ to be distributed according to the random variable $\widetilde{b}_{-i}$, his perceived expected utility is:

$$u_i^k(v_i, b_i; \widetilde{b}_{-i}) = \underset{\widetilde{b}_{-i}}{E}\left[\varphi_i^k(b_i, b_{-i})v_i - \tau_i^k(b_i, b_{-i})\right]$$

A strategy of bidder i is a a family of bid functions $\beta_i = (\beta_i^k)_k$, one for each auction format M_k where $\beta_i^k(v_i)$ denotes player i's bid in auction M_k when i's valuation is v_i.[9]

Nash equilibrium requires that for each k and v_i, player i plays a best-response to the *actual* distribution of bids of bidders $j \neq i$ in M_k. That is,

$$\beta_i^k(v_i) \in \arg\max_{b_i} u_i^k(v_i, b_i; \beta_{-i}^k)$$

where (with some slight abuse of notation) β_{-i}^k stands for the random variable of bids $\left(\beta_j^k(v_j)\right)_{j \neq i}$ as generated by the densities $(f_j(\cdot))_{j \neq i}$.

In this paper, we do not assume that bidders (necessarily) have access to β_i^k for every i and k, as the designer is assumed to be able to choose how much information feedback to provide to bidders about the distribution of bids.

The class of partial feedback that we consider is described as follows. Each player i is endowed with a partition $\boldsymbol{P}_i$ of the set

$$\{(j, k),\ j \in I \text{ and } k \in K\}$$

where $\boldsymbol{P}_i$ is called the feedback partition of player i (see the learning interpretation below).

[9]Stricctly speaking, allowing for mixed strategies $\beta_i^k(v_i)$ should be a distribution over bids. Yet, for our purpose, considering pure strategies is enough.

A typical element of $\boldsymbol{P}_i$ is denoted by α_i and referred to as a feedback class of player i. The element of $\boldsymbol{P}_i$ containing (j,k) is denoted by $\alpha_i(j,k)$. When making his choice of strategy in auction format M_k, player i is assumed to know only (in addition to the rules of M_k as described by $\varphi_i^k(b)$ and $\tau_i^k(b)$) the aggregate distribution of bids in every α_i. He is further assumed to play a best response to the conjecture that j in M_k bids according to the aggregate distribution of bids in $\alpha_i(j,k)$. Formally, we let $\boldsymbol{A} = (M_k, \lambda_k, \boldsymbol{P}_i)_{i\in I, k\in K}$ denote an auction design. A feedback equilibrium of $\boldsymbol{A}$ is defined as follows:

Definition 1 *A feedback equilibrium of* $\boldsymbol{A} = (M_k, \lambda_k, \boldsymbol{P}_i)_{i\in I, k\in K}$ *is a strategy profile* $\beta = (\beta_i)_{i\in I}$ *such that for every* k *and* v_i*,*

$$\beta_i^k(v_i) \in \arg\max_{b_i} u_i^k(v_i, b_i; \overline{\beta}_{-i}^k)$$

where $\overline{\beta}_{-i}^k = (\overline{\beta}_j^k)_{j\neq i}$ *, the distributions* $\overline{\beta}_j^k$*,* $j \neq i$ *are perceived by bidder* i *to be independent of each other, and* $\overline{\beta}_j^k$ *is the aggregate distribution of bids in* $\alpha_i(j,k)$*. That is,* $\overline{\beta}_j^k$ *is the distribution of bids that assigns weight* $\lambda_{k'} / \sum_{(j",k")\in\alpha_i(j,k)} \lambda_{k"}$ *to the distribution* $\beta_{j'}^{k'}(v_{j'})$ *as generated by the density* $f_{j'}(\cdot)$ *for every* $(j',k') \in \alpha_i(j,k)$*.*

Except otherwise mentioned, the auctioneer selects the auction design $\boldsymbol{A} = (M_k, \lambda_k, \boldsymbol{P}_i)_{i\in I, k\in K}$ and a feedback bidding equilibrium β of $\boldsymbol{A}$ so as to maximize her expected revenue:

$$\underset{v_1,\dots v_n}{E} \sum_{k\in K, i\in I} \lambda_k \tau_i^k(\beta^k(v))$$

That is, we assume that the seller has no intrinsic value for the object (her valuation is 0) and she is risk neutral.[10]

Several comments are in order. First, observe that we do not explicitly require that bidders voluntary accept to participate in $\boldsymbol{A}$. However the requirement that bidder i makes no payment when he bids $b_i = 0$ ensures that player i feels he is weakly better off participating in the auction than staying out (since he can always participate and bid 0). It should be noted though that since bidder i does not have an accurate perception of

[10] We will also discuss the alternative objective that the designer maximizes her revenue arising in the worst feedback equilibrium (strong implementation requirement), and the lexicographic objective that gives priority to efficiency relative to revenues.

other bidders' strategies, it may be that in a feedback equilibrium of $\boldsymbol{A}$, player i's expected payoff is negative. We will sometimes comment on the more stringent participation constraint that bidders whatever their types should actually make no loss in (a feedback) equilibrium.[11] Second, in the current formulation, we do not allow the feedback received by bidders to be about the distribution of profile of bids.[12] While the definition of a feedback equilibrium could be extended to cover that case, the class of partial feedback considered above is enough to prove our main insights (see Theorem 1 and Proposition 1).

3 Interpretation

The interpretation of our auction design framework is as follows. A designer faces the problem of repeatedly selling similar objects (a new one each period) to one of n potential buyers with observable characteristics $i \in I$. These potential buyers are replaced every time a new object is for sale and the observable characteristic i may affect the distribution $f_i(\cdot)$ of the valuation of the bidder with characteristic i.[13] Bidders are assumed not to be aware of $f_i(\cdot)$ (even though they can observe the characteristic of other bidders). To simplify the exposition of the interpretation, we assume that the designer has got sufficient experience to know $f_i(\cdot)$.[14] Now the designer can change her auction format from one period to another while ensuring that the frequency of format M_k is λ_k. In addition, the designer who collects the bids in all periods can decide how much information she wants to pass to new bidders about past bids. If she can target the feedback to each bidder according to his characteristic, she can decide to tell bidder i only about the aggregate distribution of past bids (the empirical frequencies) in every α_i.[15] If she does so every period, and if behaviors stabilize, it must be to a feedback equilibrium provided bidders

[11] Even though this need not be immediately perceived by a player, it could be argued that making losses is salient and likely to be transmitted to other bidders.

[12] Every bidder i treats every bidder j's distribution of bids as independent of each other.

[13] The replacement scenario corresponds to an assumption made in recurrent games (Jackson and Kalai, 1997).

[14] Observe that this knowledge may be obtained by the designer in an initial phase in which she would use second price auctions.

[15] That is, bidder i would be informed of the aggregate distribution of bids $\left\{b_j^k, (j,k) \in \alpha_i\right\}$ with no reference to which (j,k) generated the bid.

consider the simplest theory that is consistent with their feedback.[16] Our auction design framework adopts the viewpoint that the designer can optimize on the auction formats M_k, their frequencies λ_k and the feedback partitions $\boldsymbol{P}_i$ provided to bidders, and that behaviors have stabilized to a corresponding feedback equilibrium of $\boldsymbol{A}$.[17]

It may be argued that in a number of applications, the designer would have a hard time providing a different feedback to the various bidders (as the information transmitted may be shared among bidders). In such cases, the designer can always restrict herself to public feedback in which case the feedback partitions of all bidders must coincide, $\boldsymbol{P}_i = \boldsymbol{P}_j$ for all i, j. As will be clear, our main results still hold if the designer is constrained to use public feedback partitions.

It should be mentioned that a feedback equilibrium is very closely related to the analogy-based expectation equilibrium introduced in Jehiel (2005), further developed in Jehiel and Koessler (2007) and Ettinger and Jehiel (2007). The feedback partition $\boldsymbol{P}_i$ of player i is very similar to the analogy partition considered in Jehiel (2005) with the mild difference that here we allow the feedback partition to include decision nodes of player i himself. Except for this mild difference, a feedback equilibrium can be viewed as a special case of an analogy-based expectation equilibrium. The main novelty of the approach taken here is that the feedback partitions are viewed as a choice made by the designer. That is, they are not exogenously given as in Jehiel (2005).

The only other paper we are aware of that considers information feedback in auctions as a design issue is by Esponda (2007). He considers first-price auctions in which the *same* bidders get involved over sequence of auctions, and get information about the joint distribution of highest bids (and possibly second-highest bids) and their own valuation and bid. In a symmetric first-price auction with private and affiliated values he shows that symmetric self-confirming equilibria (of the static auction) generate at least as much revenues as the Nash equilibrium.

[16]It is in this sense that a feedback equilibrium can be viewed as a selection of self-confirming equilibrium for the signal structure corresponding to the chosen feedback partition.

[17]Strictly speaking, the system should stabilize to the feedback equilibrium the designer likes best. This is similar to the requirement of weak implementation generally made in mechanism design. One interpretation is that the designer could suggest a default belief that would fit the equilibrium she likes best. Alternatively, one may wish to reinforce the notion of implementation to require that all feedback equilibria deliver good outcomes. Our main results would still hold under this more stringent notion of implementation.

Apart from the obvious difference that Esponda considers first-price auctions with symmetric bidders and affiliated signals whereas we consider general auction formats with arbitrary yet independent distributions of valuations, Esponda's result shares some similarities with our insight that partial feedback may help achieve greater revenues in first-price auctions (see Proposition 1 below). Yet, there are notable differences between our framework and his that we now discuss. First, Esponda considers a setting in which the same bidders keep participating in the auctions whereas we have in mind situations in which new bidders arrive each time. This difference in turn explains why in our setting the feedback of bidders is not conditional on their own valuation whereas in Esponda's setting it is.[18] Second, Esponda's solution concept is the self-confirming equilibrium whereas we rely on the feedback equilibrium (which is a selection of self-confirming equilibrium, see above). As such Esponda's analysis can never rule out that providing partial feedback does no better than providing full feedback (since the Nash equilibrium is always a self-confirming equilibrium whatever the feedback). More problematic: his result is for symmetric setups and symmetric self-confirming equilibria, as there is no guarantee that an asymmetric self-confirming equilibrium generates more revenues than the Nash equilibrium even in a symmetric setup.[19] By contrast, our insights about first-price auctions concern the case of asymmetric bidders with independent distributions, and the selection imposed by the feedback equilibrium (based on complexity considerations) ensures a strict superiority of providing partial feedback (see Proposition 1).

Examples of feedback partitions and auction designs:

The following classes of auction designs with public feedback (all $\boldsymbol{P}_i$ are the same) will play a central role in the analysis.

1) *Format-anonymous feedback partition*: In this case, bidders know the aggregate distribution of bids across the different auction formats M_k, $k \in K$, but they differentiate

[18] One should note that in Esponda's paper, bidders simply ignore the distribution of highest bids conditional on other realizations of the valuation. If bidders somehow mistakenly mixed these distributions for various realizations of their valuations (because say there are not enough data for each specific realization of the valuation), then providing partial feedback about the highest bid might result in a revenue fall in the affiliated environment he considers.

[19] I view this as problematic as I fail to see what mechanism would lead bidders to have symmetric behaviors given that there are many possible conjectures under the partial feedback considered in Esponda and many different best-responses associated to these conjectures.

the distribution of bids for the various bidders $i \in I$. That is, for all $i \in I$, $\boldsymbol{P}_i = \left\{ \bigcup_{k \in K} \{(j,k)\} \right\}_{j \in I}$. For example, the object could be sold either through a sealed bid first price auction or through a sealed bid second price auction (in equal proportion, say), and players would receive feedback about the aggregate distribution of bids across the two auction formats.

2) *Bidder-anonymous feedback partition*: In this case, there is only one auction format, and the feedback is about the aggregate distribution of bids across all bidders. That is, $K = \{1\}$, and for all $i \in I$, $\boldsymbol{P}_i = \left\{ \bigcup_{j \in I} \{(j,1)\} \right\}$. For example, the object could be sold through a first-price auction, and players would receive feedback about the aggregate distribution of bids with no mention of the characteristics of the bidders who generated the various bids.

4 Preliminaries

We make a few preliminary observations (the claims are proved in the appendix). First, by picking a single auction format M and the finest feedback partition, the designer can always replicate the revenue generated in the standard rationality case in M. It follows that the designer can always achieve at least as much as Myerson (1981)'s optimal revenue. The question is whether the designer can achieve more revenues.

Second, consider an auction format M in which player i has a dominant strategy. Then no matter what the auction design is, a feedback equilibrium requires that player i plays his dominant strategy in M. This is an obvious statement, since player i will find his strategy best no matter what his expectation about the distribution of others' bids is. In particular, if one of the auction formats is a second-price auction SPA it will always be part of an equilibrium for them to bid their true valuation in SPA.

Third, some of the auction designs $\boldsymbol{A}$ considered below have the feature that the various auction formats M_k are such that $\varphi_i^k(b_i, b_{-i}) = \varphi_i(b_i, b_{-i})$ for all $k \in K$ (for example in all formats the object is allocated to the player who submitted the highest bid). When the *anonymous format feedback partition* prevails, one can relate the feedback equilibria of $\boldsymbol{A}$ to the Nash Bayes equilibria of the following game referred to as $\Gamma(\boldsymbol{A})$:

Game $\Gamma(\boldsymbol{A})$: Each bidder i (simultaneously) submits a bid b_i; the object is assigned to bidder i with probability $\varphi_i(b_i, b_{-i})$; prior to bidding, bidder i is privately informed of his valuation v_i drawn from $f_i(\cdot)$ and of his method of payment k defined by $\tau_i^k(b_i, b_{-i})$; the methods of payment k are identically and independently drawn across bidders and every bidder i is subject to the method of payment k with probability λ_k.

Claim 1: Suppose that the format anonymous feedback partitions prevail and that in all auction formats M_k of $\boldsymbol{A}$, we have that $\varphi_i^k(b_i, b_{-i}) = \varphi_i(b_i, b_{-i})$ for all $k \in K$ and $i \in I$. Then a strategy profile β is a feedback equilibrium of $\boldsymbol{A}$ if and only if it is a Nash Bayes equilibrium of $\Gamma(\boldsymbol{A})$.

Fourth, another auction design that we will study has the following form. There is one auction format M, which respects the anonymity of bidders. That is, consider two bid profiles b and b' obtained by permuting the bids of players i and j, then $\varphi_i(b) = \varphi_j(b')$ and $\tau_i(b) = \tau_j(b')$ and for all $m \neq i, j$, $\varphi_m(b) = \varphi_m(b')$, $\tau_m(b) = \tau_m(b')$. Consider the *anonymous bidder feedback partition* defined above. One can relate the feedback equilibria of $\boldsymbol{A}$ to the Nash Bayes equilibria of game $\Gamma'(\boldsymbol{A})$ defined by the auction format M in which the distribution of bidder i has density $\overline{f}(v_i) = \sum_{i \in j} f_j(v_i)/n$ instead of $f_i(v_i)$.

Claim 2: A symmetric strategy profile is a feedback equilibrium of $\boldsymbol{A}$ if and only if it is a Bayes Nash equilibrium of $\Gamma'(\boldsymbol{A})$.

5 Main Results

5.1 Revenues

Our first main result establishes that the designer can always achieve strictly larger revenues by a suitable choice of manipulative auction design as compared with the best nonmanipulative auction design of Myerson (1981).

Theorem 1 *The largest revenue that the designer can achieve in a manipulative auction design is strictly larger than the revenue generated in Myerson's optimal auction.*

A possible idea to prove Theorem 1 might go as follows. Consider the regular two-bidder symmetric case ($f_1(\cdot) = f_2(\cdot) = f(\cdot)$ and $v \to v - \frac{1-F(v)}{f(v)}$ is strictly increasing over $[L, L+1]$) with large L ($L - \frac{1}{f(L)} > 0$). We know that the second-price auction (SPA) implements Myerson's optimal auction in this case. Consider now a mix of second price auctions (SPA) and first price auctions (FPA) in which bidders get only to know the aggregate distribution of bids over the two auction formats. Given that bids in SPA are higher than bids in FPA, one might conjecture that the confusion between the two distributions of bids would lead bidders to bid more aggressively in FPA than in the usual rationality case. Yet, this is not always so because the incentive to bid more aggressively in FPA is not driven by the first-order stochastic dominance relation in the distribution of bids.[20] However, we show that by mixing SPA with some other well chosen format M^*, the confusion between SPA and M^* can increase revenues.

Proof of Theorem 1- The symmetric two-bidder regular case.

Consider the following auction design in which we assume that the maximum possible bid coincides with the maximum valuation, i.e., $B = L+1$. There are two auction formats $M_1 = SPA$ (the second price auction) and $M_2 = M^*$ which will be defined shortly. SPA is chosen with probability $\lambda_1 = 1 - \varepsilon$ and $M_2 = M^*$ is chosen with probability $\lambda_2 = \varepsilon$ where ε is assumed to be small enough. The feedback partitions are the format-anonymous feedback partitions for both players.[21] That is, for $i = 1, 2$, $\boldsymbol{P}_i = \{\{(j,1),(j,2)\}\}_{j \in I}$.[22]

Format M^* is parameterized by $z \in (L, L+1)$ and it is defined as follows. The allocation rule $\phi^*(b)$ is the same as that of SPA: the bidder who submits the highest bid wins the object; if both bidders submit the same bid, each bidder receives the object with probability 1/2. The payment scheme is the same as that of the SPA when a bidder sends a bid less than z. That is, when $b_i \leq z$, bidder i makes a payment only when he

[20] In the case of uniform distributions, the bids arising in the corresponding feedback equilibrium would coincide with the bids arising in NE.

[21] As we argued above the equilibrium analysis in our manipulative auction design is the same as the equilibrium analysis in the standard rationality case of an auction setup in which each bidder's payment rule is drawn at random independently across bidders and it corresponds either to the rule of the second price auction with probability $1-\varepsilon$ or the payment rule of format M_2 with probability ε. The revenue generated by the format is however different from that generated in our manipulative auction design, and by Myerson's optimal auction theorem we know it is no greater than that in Myerson's optimal auction.

[22] Given the symmetry of bidders, we could also consider the coarsest feedback partition, e.g. $\boldsymbol{P}_i = \{\{(1,1),(1,2),(2,1),(2,2)\}\}$. That is, the format and bidder anonymous feedback partition.

wins and this payment is the second highest bid, i.e. $\min(b_1, b_2)$. When $b_i > z$, $b_i \leq B$, bidder i makes a payment equal to $p(z,\varepsilon)$ whether or not he wins where

$$z - p(z,\varepsilon) - \varepsilon\frac{1-F(z)}{2}z = \int_L^z (z-v)f(v)dv \qquad (1)$$

We now check that the following strategy profile is a feedback equilibrium in this auction design:

- In format $M_1 = SPA$, each bidder i submits a bid b_i equal to his true valuation v_i : $\beta_i^1(v_i) = v_i$.
- In format $M_2 = M^*$,

$$\beta_i^2(v_i) = \begin{cases} v_i \text{ if } v_i < z \\ B \text{ if } v_i \geq z \end{cases}$$

Observe that if bidders follow the above strategies, the aggregate distribution of bids over the two auction formats assign a weight $1-\varepsilon$ to the distribution $f(\cdot)$ of bids and a weight ε to the distribution which has a density $f(v_i)$ for $v_i \in [L, z]$ and a mass $1 - F(z)$ on $B = L + 1$. In format SPA, bidding his own valuation is a weakly dominant strategy. In format M_2, given the form of the aggregate distribution of bids, the payment scheme of M_2, and condition (1), a bidder with valuation z is indifferent between bidding $b_i = z$ (in which case he perceives to get $\int_L^z (z-v)f(v)dv$) and bidding B (in which case he perceives to get $z - p(z,\varepsilon) - \varepsilon\frac{1-F(z)}{2}z$)[23]. It is now readily verified that the above strategies are a feedback equilibrium (because bidding $b_i \in (z, B)$ is always dominated by bidding B whenever $b_i = B$ is perceived to be profitable, and within the range $b_i \in [0, z]$ submitting the bid closest to his own valuation is always perceived to be the best alternative).

The rest of the argument consists in showing that for ε small enough and z close enough to $L+1$, the revenue generated by this auction design is strictly larger than the revenues generated by the second-price auction.

A preliminary observation is that the revenue generated in our auction design writes

$$(1-\varepsilon)R^{SPA} + \varepsilon R^*$$

[23] He always pays $p(z,\varepsilon)$ and wins the object with probability $1 - \varepsilon\frac{1-F(z)}{2}$.

where R^{SPA} is the revenue of SPA and R^* is the revenue generated in our auction design when the format is M^*. We thus have to show that R^* is larger than R^{SPA} for suitable choices of z ad ε. It can be shown that it is always so when z is close enough to $L+1$ and ε is set sufficiently small (see details in the appendix).

Roughly, $R^* > R^{SPA}$ for ε small and z large because as ε goes to 0, the price $p(z, \varepsilon)$ paid in M^* by high valuation types $v > z$ does not vary much as z moves locally from $z = L + 1$ ($\frac{\partial p}{\partial z}(z = L+1, \varepsilon = 0) = 0$) whereas the corresponding expected payment $b^{FPA}(z) = \frac{\int_L^z vf(v)dv}{F(z)}$ made in the second price auction decreases more significantly for the same local change ($\frac{db^{FPA}}{dz}(z = L+1) > 0$). The result then follows by rearranging the expression of the expected revenues (and by observing that $p(z = L+1, \varepsilon = 0) = b^{FPA}(z = L+1)$). **Q. E. D.**

The intuition behind the proof of Theorem 1 is as follows. The manipulative auction design considered in the proof consists of a mix between a main auction format in which bidders have a dominant strategy (the second price auction in the simple case) and an auxiliary auction format, which is very similar to the main auction format except for very high valuation types who are proposed to make a lump-sum payment to the designer and win the object if they happen to have the highest bid. Because the designer employs the format anonymous feedback partition, high valuation bidders in the auxiliary format have the feeling they have a very high chance of winning the auction (in fact higher than in reality because in the auxiliary format unlike in the main format all high valuation bidders would submit the highest bid). As a result the high valuation bidders are ready to pay more than they would in the standard rationality case. The manipulation employed by the designer has the effect of reducing the informational rent left to the high valuation bidders. As such it can be designed so that all bidders whatever their type get a positive expected payoff in equilibrium.

Remark. By considering the coarsest feedback partition instead of the format-anonymous feedback partition, the feedback equilibrium obtained in the auction environment in the proof of Theorem 1 is easily shown to be the only equilibrium not employing weakly dominated strategies. It follows that the conclusion of Theorem 1 holds even if bidders can coordinate on the feedback equilibrium (not employing weakly dominated strategies) that is least favorable to the designer.

5.2 Efficiency and revenues

In a number of applications, the designer may be interested in efficiency as well as revenues. For example, suppose that the primary objective of the designer is efficiency while revenue is only the secondary objective. In the standard rationality paradigm, the so called revenue equivalence result holds. That is, if two mechanisms result in the same allocation rule and the expected payment made by any bidder i with minimal valuation $v_i = L$ is 0 then the two mechanisms must yield the same revenues. Since an efficient outcome can be achieved by a second-price auction SPA, the standard approach concludes that the designer can do no better than using a SPA.

We now observe that the designer can sometimes achieve strictly larger revenues (than that obtained through the SPA) while still preserving efficiency, thereby illustrating a failure of the allocation equivalence in our manipulative auction design setup. Besides, this gain in revenues is achieved by using a fairly standard auction format (with, of course, a non-standard, i.e. coarse, feedback partition).

Proposition 1 *Let there be two bidders $i = 1, 2$ with asymmetric distributions ($F_1(\cdot) \neq F_2(\cdot)$ on a set of strictly positive measure). Consider a symmetric feedback equilibrium of the first price auction with anonymous bidder feedback partition. It induces an efficient outcome and generates a strictly higher revenue than the second-price auction.*

Proof of Proposition 1:

Step 1: Consider the first-price auction with anonymous bidder feedback partition. The feedback equilibrium is for $i = 1, 2$, $\beta_i(v) = \beta(v) = \frac{\int_L^v x\overline{f}(x)dx}{\overline{F}(v)}$ where $\overline{f}(x) = \frac{f_1(x)+f_2(x)}{2}$ and $\overline{F}(v) = \frac{F_1(v)+F_2(v)}{2}$. It follows that the outcome is efficient in our auction design.

Proof of step 1. Any symmetric feedback equilibrium must be a Nash Bayes equilibrium of the FPA with symmetric bidders and density $\overline{f}(v)$ and vice versa (see Claim 2 in Section 4). Given the analysis of the FPA with symmetric bidders, we may conclude as desired. **Q. E. D.**

Call R the revenue generated in the first price auction with bidder anonymous feedback partition. Call R^{SPA} the revenue generated in the second-price auction. Finally, call $\overline{R}$

the expected revenue generated in the second-price auction with symmetric bidders and density of valuations $\overline{f}(v) = \frac{f_1(x)+f_2(x)}{2}$. These revenues write:

$$
\begin{aligned}
R &= \int_L^{L+1} \beta(v) \left[f_1(v) F_2(v) + f_2(v) F_1(v) \right] dv \\
R^{SPA} &= \int_L^{L+1} v f_1(v) \left[1 - F_2(v) \right] dv + \int_L^{L+1} v f_2(v) \left[1 - F_1(v) \right] dv
\end{aligned}
$$

and we also have (the identity between the two can be obtained as a consequence of the allocation equivalence):

$$
\begin{aligned}
\overline{R} &= 2 \int_L^{L+1} v \overline{f}(v) \left[1 - \overline{F}(v) \right] dv \\
\overline{R} &= 2 \int_L^{L+1} \beta(v) \overline{f}(v) \overline{F}(v) dv
\end{aligned}
$$

Step 2: $\overline{R} > R^{SPA}$

Proof of step 2. Using the first expression of $\overline{R}$, we have that $\overline{R} - R^{SPA}$ can be written as

$$
\begin{aligned}
& \int_L^{L+1} v \left[-\frac{1}{2} \left(F_1(v) + F_2(v) \right) \left(f_1(v) + f_2(v) \right) + f_1(v) F_2(v) + f_2(v) F_1(v) \right] dv \\
= & \int_L^{L+1} -\frac{v}{2} \left(f_1(v) - f_2(v) \right) \left(F_1(v) - F_2(v) \right) dv \\
= & \int_L^{L+1} \frac{1}{4} \left(F_1(v) - F_2(v) \right)^2 dv
\end{aligned}
$$

where the last equality is obtained by integration by parts (noting that $F_1(v) - F_2(v) = 0$ for $v = L$ and $L+1$). Since $\int_L^{L+1} \frac{1}{4} \left(F_1(v) - F_2(v) \right)^2 dv > 0$ whenever $f_1 \neq f_2$. Step 2 follows. **Q. E. D.**

Step 3: $R > \overline{R}$

Proof of step 3. Using the second expression of $\overline{R}$, we have that $R - \overline{R}$ can be written as

$$
\begin{aligned}
& \int_{L}^{L+1} \beta(v) \left[f_1(v)F_2(v) + f_2(v)F_1(v) - 2\frac{f_1(v)+f_2(v)}{2} \cdot \frac{F_1(v)+F_2(v)}{2} \right] dv \\
= & \int_{L}^{L+1} -\frac{1}{2}\beta(v) \left(f_1(v) - f_2(v)\right) \left(F_1(v) - F_2(v)\right) dv \\
= & \int_{L}^{L+1} \frac{1}{4}\frac{d\beta(v)}{dv} \left(F_1(v) - F_2(v)\right)^2 dv
\end{aligned}
$$

where the last equality is obtained by integration by parts (noting that $F_1(v) - F_2(v) = 0$ for $v = L$ and $L+1$). Since $\frac{d\beta(v)}{dv} > 0$ for all v, we have that $\int_{L}^{L+1} \frac{1}{4}\frac{d\beta(v)}{dv} \left(F_1(v) - F_2(v)\right)^2 dv > 0$ whenever $F_1 \neq F_2$ on a positive measure set. Step 3 follows. **Q. E. D.**

Proposition 1 clearly follows from steps 1, 2, 3. **Q. E. D.**

What is the intuition for the above result? Providing feedback about the aggregate distribution of bids among the two bidders leads the bidders to feel that they are in competition with a fictitious bidder who has a distribution of valuations that is the average distribution between the distributions of the two bidders (this is essentially step 1 in the proof). In the two bidder case, the price level in the second-price auction is determined by the lowest valuation, hence by the weak bidder. The manipulation generated by the bidder-anonymous feedback partition is good for revenues because it makes the strong bidder feels the weak bidder is less weak than he really is (this is what steps 2 and 3 formalize).

It should be noted that the result of the above Proposition need not hold when there are three or more bidders. This is because averaging the distributions across all bidders need not anymore strengthen the distribution of the second highest bid. Technically, while steps 1 and 3 of the proof still hold in the three or more bidder case, step 2 need not hold in general. For the sake of illustration, if there are two bidders with a distribution of valuations concentrated around $L + 1$ and a third bidder whose valuation is concentrated around L, it is readily verified that the first-price auction with bidder anonymous feedback partition generates less revenues than the second-price auction (which achieves a revenue approximately equal to $L + 1$). Of course, if there is only one bidder whose distribution of valuations is concentrated around $L + 1$ while the other bidders have a distribution of valuations concentrated around L, the first-price auction with bidder anonymous feedback partition generates more revenues than the second-price auction (which generates a revenue very close to L). These extreme cases illustrate that the revenue comparison between the second-price-auction and the first-price auction with bidder anonymous feedback partition can go either way in the case of more than two bidders.

It should also be mentioned that in the first-price auction with anonymous bidder feedback partition no bidder whatever his valuation makes a loss in equilibrium. This is because bidders bid below their valuation (bidding above his valuation is dominated) and thus they cannot make losses. The manipulation exploited by the designer has here the effect of reducing the informational rent left to the bidders. The most extreme revenue gain provided by the exploitation is when the distribution of valuations of one bidder is concentrated around $L+1$ whereas the distribution of valuations of the other bidder is concentrated around L in which case the first-price auction with anonymous feedback partition provides a revenue gain of $\frac{1}{2}$ compared to the revenue L of the second-price auction.[24]

Comment. In some applications, the distribution of winning bids as opposed to the aggregate distribution of all bids is available to bidders. From this information, bidders can compute an optimal strategy based on the assumption that all bidders bid according to the same distribution (this might be argued to be the simplest conjecture in this case). In the two-asymmetric-bidder scenario considered in Proposition 1, it is not difficult to show that bidders would then bid according to $\beta^*(v) = \frac{\int_L^v x f^*(x)dx}{F^*(v)}$ where $F^*(v) = (F_1(v)F_2(v))^{\frac{1}{2}}$ and $f^*(v) = \frac{dF^*(v)}{dv}$. Such a behavior would always generate higher revenues than in the second price auction, as in Proposition 1.[25]

[24] In some cases, the first-price auction with bidder anonymous feedback partition may provide a revenue that is larger than the revenue in the optimal auction of Myerson, even though this is not always true.

[25] The analog of steps 2 and 3 in the proof of Proposition 1 would still hold. Letting R^* denote the revenue in the manipulative auction setup (as just defined) and $\overline{R}^*$ the revenue in the second-price auction with symmetric bidders and density of valuations $f^*(\cdot)$, one can establish that

$$\begin{aligned} \overline{R}^* - R^{SPA} &= \int_L^{L+1} [\sqrt{F_1(v)} - \sqrt{F_2(v)}]^2 \, dv > 0 \\ R^* - \overline{R}^* &= 0 \end{aligned}$$

6 Further results

6.1 Complete information

In the above analysis, we have assumed that there was some uncertainty about bidders' valuations. When each bidder i's valuation can take a single value v_i, the classic revenue-maximizing mechanism allows the designer to extract a revenue equal to $\max_{i \in I} v_i$ (by making a take-it-or-leave-it offer of v_{i^*} to bidder $i^* \in \arg\max_{i \in I} v_i$).

We now ask ourselves whether the designer can extract more that v_{i^*} in our manipulative auction design setup? We find that this is so, but (unsurprisingly) it comes at the cost of having at least one bidder get a strictly negative payoff. As such, it seems that manipulations in the complete information case are less likely to remain undetected than in the incomplete information case analyzed earlier.[26]

To start with, observe that if the designer is able to extract strictly more than v_{i^*}, it is necessary that at least one bidder $i \in I$ gets a negative payoff. Indeed, the maximum surplus among all agents is v_{i^*} and if the designer gets strictly more than v_{i^*}, it must be that at least one bidder $i \in I$ gets a strictly negative payoff.

To illustrate the possibility of manipulation in the complete information case, consider the following setup.[27] There are three bidders $i = 1, 2, 3$ with $v_1 > v_2 > v_3 > 0$, and a single auction format described as follows. The bidder with highest bid wins the object. Bidder 1 pays v_1 if he wins and $b_1 \neq 0$ (and he pays 0 otherwise). Bidder 3's payment rule is that of the second-price auction. Finally, bidder 2 makes a lump-sum payment of l with $0 < l < \frac{v_2 - v_3}{4}$ if $b_2 \neq 0$ in addition to the transfer of the second-price auction made upon winning the auction. He pays 0 if $b_2 = 0$. Assume that bidder 2 gets only to know the aggregate distribution of bids of bidders 1 and 3 while bidders 1 and 3 receive the finest feedback.

[26] In the complete information case to be considered now, manipulation takes the form of extracting positive revenue also from bidders who do not acquire the object, hence the negative payoff. This is very different from manipulation in the incomplete information case highlighted above, which allowed the designer to reduce the informational rent left to the high valuations types.

[27] Our example assumes that there are three bidders. It obviously extends to the case of more than three bidders. If the designer is allowed to use shill bidders as in the next subsection, it can also be extended to the case of two bidders (in which case the third bidder of the example can be thought of as a shill bidder).

Claim 3: Bidding $\beta_i(v_i) = v_i$ for $i = 1, 2, 3$ in the above auction design is a feedback equilibrium, and it generates a revenue $v_1 + l > v_1$ to the designer.

Proof of Claim 3: This is straightforward for bidders 1 and 3. Bidder 2 if he makes a non-zero bid finds it optimal to bid $b_2 = v_2$ (this follows from the usual analysis of the second-price auction). The only remaining issue is whether bidder 2 chooses to make a non-zero bid given the extra lump-sump payment $l > 0$ required in this case (he would not in the standard rationality case). Given his feedback partition, bidders 2 thinks each of bidders 1 and 3 bid v_1 or v_3 each with probability $\frac{1}{2}$. So 2 thinks that with probability $\frac{1}{4}$ both 1 and 3 will bid v_3 in which case apart from the lump-sump payment 2 would make a net gain of $v_2 - v_3$. Given that $l < \frac{v_2 - v_3}{4}$ it follows that bidder 2 thinks he is better off making a non-zero bid. **Q. E. D.**

6.2 Shill bidding

So far, we have assumed that the only players in the auction were the bidders $i \in I$. It might be argued that the designer could also employ shill bidders in addition to the real bidders $i \in I$. In the standard case, this does not help the designer obtain a better outcome, but in our manipulative mechanism design setup it might.

To illustrate the potential benefit of employing shill bidders, consider the following framework. There are two real bidders $i = 1, 2$ whose valuations are drawn from the same distribution with density $f(\cdot)$. If a first-price auction with fine feedback is considered, bidders bid according to the bid function $\beta(v) = \frac{\int_L^v xf(x)dx}{F(v)}$. Now consider the alternative format M^* in which there are many bidders $i = 3, 4....N$ in addition to bidders $i = 1, 2$ and all bidders $i = 3, 4, ...N$ are shill bidders instructed (or incentivized) to bid according to a distribution of bids with density $g(\cdot)$. In format M^*, only the bids b_1 and b_2 of bidders $i = 1, 2$ matter (for the determination of who wins and what the payments are). Restricted to these two bids, the rules of M^* are those of the first-price auction. The bidder $i = 1, 2$ with highest bid wins the auction and pays b_i to the designer. In the fine feedback case, bidders $i = 1, 2$ would bid $\beta(v) = \frac{\int_L^v xf(x)dx}{F(v)}$, but in the bidder-anonymous feedback partition case, bidders would now bid $\frac{\int_L^v x\overline{f}(x)dx}{\overline{F}(v)}$ where $\overline{f}(x) = \frac{2f(x)+(N-2)g(x)}{N}$. In the limit as N grows to infinity, their bidding strategy would converge to $\beta_g(v) = \frac{\int_L^v xg(x)dx}{G(v)}$

and the corresponding revenue obtained by the seller would get close to

$$R_g = 2\int_L^v \beta_g(v)F(v)f(v)dv$$

Restricting attention to this class of manipulations, the designer's problem is to find a density $g(\cdot)$ which maximizes $\int_L^v \beta_g(v)F(v)f(v)dv$ where $\beta_g(v) = \frac{\int_L^v xg(x)dx}{G(v)}$.

This is not a trivial optimization problem. Yet, we make the following simple observation.

Claim 4: Consider a smooth density $g(\cdot)$ where for some $x \in (L, L+1)$, $g(x) > 0$ and

$$\int_x^{L+1} \frac{F(v)f(v)}{G(v)}dv \neq \frac{F(x)f(x)\int_L^x G(v)dv}{(G(x))^2} \tag{2}$$

Then, the designer's revenue R_g can be increased by a local change of g in the neighborhood of x.

As a simple corollary of Claim 4, we find that it is never optimal to have $g(\cdot) = f(\cdot)$. That is, there is always scope for manipulation of this sort.

Corollary 1 *There is always scope for using shill bidders in first-price auctions.*

Proof of Corollary: Let $g(\cdot) = f(\cdot)$. (2) writes: $F(x)[1 - F(x)] = f(x)\int_L^x F(v)dv$. But this cannot be satisfied for x close to $L+1$ as $f(x) > 0$ for all $x \in [L, L+1]$. **Q. E. D.**

7 Extensions and future work

In this paper, we have shown that there is always scope for manipulation in private values auctions. That is, a designer interested in revenues can profitably design an auction setup in which she provides partial rather than total feedback about other bidders' strategies.

There are several interesting directions for future research that we now wish to mention. First, the analysis of this paper does not tell us about the optimal manipulative auction design. We note that finding the optimal mechanism design does not seem to be an easy exercise as there is no analog here of the revelation principle, and the space of auction designs is *a priori* quite large. Second, we have stressed here the manipulative nature of the feedback device that a designer interested in revenue maximization can profitably use. Even though we are not aware (in the real world) of deliberate use of such manipulative strategies to maximize revenues, it seems that the idea of providing partial rather than total feedback about the distribution of bids observed in similar auctions is quite common in practical auction design where it is often regarded as a way to preserve the anonymity/privacy of bidders and/or as a way to combat collusion (that would otherwise be easier to sustain). From the latter perspective, it would be interesting in future research to formalize collusion in our auction setup and to analyze how coarse feedback partitions can be used to combat collusion. Third, our analysis of manipulative auction design could be extended to study the case of stochastic number of bidders and the case of correlated distributions of valuations. Fourth, while in our setting only the designer can provide feedback about the distribution of past bids, one could imagine that bidders themselves try to provide additional feedback to better combat manipulation.[28] Fifth, our analysis and main results even though formulated for the auction setup can be exported to other applications. For example, a result similar to that of Proposition 1 can be exported to contests to suggest the desirability of not providing the characteristics of competitors in asymmetric two-player contests as a way to increase the effort made by the stronger contestant without altering the normative desideratum that a better contestant should win with larger probability no matter what his (observable) characteristics are.

[28] In the context considered here in which each bidder participates in just one auction, the incentive to do so is rather limited, but it would be different if the same bidders participated several times.

Appendix

Proof of Claim 1: Consider an equilibrium β of $\Gamma(\boldsymbol{A})$. In $\Gamma(\boldsymbol{A})$, bidder i whatever his payment method expects every other bidder $j \in I$ to be facing the payment method k' with probability $\lambda_{k'}$, hence to be playing according to strategy $\beta_j^{k'}(\cdot)$ with probability $\lambda_{k'}$. Thus, in $\Gamma(\boldsymbol{A})$, when the payment method is k, bidder i plays a best-response $\beta_i^k(v_i) \in \arg\max_{b_i} u_i^k(v_i, b_i; \overline{\beta}_{-i}^k)$ where $\overline{\beta}_j^k = \sum_{k'} \lambda_{k'} \beta_j^{k'}$ and $\beta_j^{k'}$ is the distribution of bids of bidder j when j has the method of payment k'. But, this corresponds exactly to the definition of a feedback equilibrium of $\boldsymbol{A}$. **Q. E. D.**

Proof of Claim 2: Consider a symmetric feedback equilibrium β of $\boldsymbol{A}$ (where $\beta(v)$ refers to the equilibrium bid of any bidder with valuation v).[29] By definition, bidder i plays a best-response to the distribution of bids of other bidders that has assigns density $\frac{\sum_{j \in I} f_j(v)}{n}$ to the bid $\beta(v)$. But, this is the definition of a Bayes Nash equilibrium of $\Gamma'(\boldsymbol{A})$. The converse part is also immediate. **Q. E. D.**

Proof of Theorem 1 (cont'd):

Lemma 1 *$\exists D > 0$ and $\overline{z} < L+1$ s.t. $\forall \varepsilon$, $z > \overline{z}$*

$$\frac{1}{2} f(L+1)(L+1-z) - D\varepsilon < \frac{\partial p(z,\varepsilon)}{\partial z} < \frac{3}{2} f(L+1)(L+1-z) + D\varepsilon$$

Proof 1 *Differentiating (1) with respect to z yields*

$$\frac{\partial p(z,\varepsilon)}{\partial z} = 1 - F(z) - \varepsilon \left[-\frac{f(z)}{2} z + \frac{1-F(z)}{2} \right]$$

Noting that $-D < -\frac{f(z)}{2} z + \frac{1-F(z)}{2} < D$ for some D and that $F'(L+1) = f(L+1)$ yields the result (by the theorem of intermediate values and the continuity of $f(\cdot)$)[30]*.*

It will be convenient to let $b^{FPA}(v) = \frac{\int_L^v x f(x) dx}{F(v)}$ denote the (standard) equilibrium bid in the first price auction. Observe that $\frac{db^{FPA}}{dv}(L+1) = f(L+1)\left(1 - \int_L^{L+1} x f(x) dx\right) > 0$.

[29] The anonymity properties of M_1 ensure the symmetry (across bidders) of the best-response correspondence.

[30] $\overline{z}$ should be chosen so that for $z > \overline{z}$, $\frac{1}{2} f(L+1) < f(z) < \frac{3}{2} f(L+1)$.

(Remember that $f(\cdot) > 0$ on $[L, L+1]$.) We also have that for all ε, $p(z = L+1, \varepsilon) = b^{FPA}(L+1)$.

In M^*, the revenue from bidder i conditional on $v_i \geq v_j$ times the probability that $v_i \geq v_j$ writes:

$$\int_L^z f(x) \left(\int_L^x y f(y) dy \right) dx + \int_z^{L+1} p(z) f(x) \left(\int_L^x f(y) dy \right) dx$$

or

$$A(z) + C(z, \varepsilon)$$

where

$$\begin{aligned} A(z) &= \int_L^z f(x) \left(\int_L^x y f(y) dy \right) dx + \int_z^{L+1} b^{FPA}(x) f(x) \left(\int_L^x f(y) dy \right) dx \\ C(z, \varepsilon) &= \int_z^{L+1} \left[p(z, \varepsilon) - b^{FPA}(x) \right] f(x) F(x) dx \end{aligned}$$

The total revenue in M^* is no smaller than $2(A+C)$ as it is no smaller than the sum of the above revenues for $i = 1, 2$.[31]

Besides, it is readily verified that for all z, $2A(z) = R^{SPA}$ (this can be seen directly by replacing $b^{FPA}(x)$, and it is nothing else than a consequence of the standard allocation equivalence).

Our final observation is that one can choose z close enough to $L+1$ and ε small enough so that $C(z, \varepsilon) > 0$, thereby showing that the revenue generated in our auction design is larger than R^{SPA}. This can be seen by differentiating $C(z, \varepsilon)$ with respect to z:

$$\frac{\partial C}{\partial z}(z, \varepsilon) = -\left[p(z, \varepsilon) - b^{FPA}(z) \right] f(z) F(z) + \int_z^{L+1} \frac{\partial p(z, \varepsilon)}{\partial z} f(x) F(x) dx$$

Letting $z(\eta) = L + 1 - \eta$, consider $\overline{\eta} < L + 1 - \overline{z}$ (see Lemma 1) and $\underline{a}$, $\overline{a}$ such that for all $\eta < \overline{\eta}$, $\frac{1}{2} \frac{db^{FPA}(L+1)}{dv} < \frac{db^{FPA}(z(\eta))}{dv} < \frac{3}{2} \frac{db^{FPA}(L+1)}{dv}$ and $\underline{a} < f(z(\eta)) F(z(\eta)) < \overline{a}$.

Noting that $p(L+1, \varepsilon) = b^{FPA}(L+1)$ we get from the intermediate value theorem

[31] It is in fact larger since a bidder with valuation above z will sometimes pay $p(z)$ even though he does not have the larger valuation.

(applied to $p(z,\varepsilon) - b^{FPA}(z)$) and the above lemma that for all $\eta < \overline{\eta}$,

$$\frac{\partial C}{\partial z}(z,\varepsilon) < -\frac{\underline{a}}{2}\frac{db^{FPA}(L+1)}{dv} \cdot \eta + \frac{3}{2}f(L+1)(2\overline{a}) \cdot \eta^2 + D(2\overline{a})\eta \cdot \varepsilon$$

Letting $\varepsilon = \eta^2$ and observing that

$$\begin{aligned} \frac{\partial C}{\partial \varepsilon}(z,\varepsilon) &= -\frac{1-F(z)}{2}z\int_z^{L+1} f(x)F(x)dx \\ &> -\frac{L+1}{2}\overline{a}\eta \end{aligned}$$

we get that along the path $(z(\eta), \varepsilon = \eta^2)$,

$$\frac{d}{d\eta}C(z(\eta), \varepsilon = \eta^2) = -\frac{\partial C}{\partial z} + \frac{\partial C}{\partial \varepsilon}\frac{\partial \varepsilon}{\partial \eta},$$

which is no smaller than

$$\frac{\underline{a}}{2}\frac{db^{FPA}(L+1)}{dv} \cdot \eta + G \cdot \eta^2$$

where $G = -\frac{3}{2}f(L+1)(2\overline{a}) - D(2\overline{a})\eta - (L+1)\overline{a}$. Hence, there must exist $\widehat{\eta}$ such that for all $\eta < \widehat{\eta}$, $\frac{d}{d\eta}C(z(\eta), \varepsilon = \eta^2) > 0$. This implies that for all $0 < \eta < \widehat{\eta}$, $C(z(\eta), \varepsilon = \eta^2) > 0$ since $C(z(0), \varepsilon = 0) = 0$).

1) *The asymmetric regular case.* Assume that $c_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)We}$ is increasing with v_i. The optimal auction in this case requires allocating the object to bidder $i^* \in \arg\max_{i \in I} c_i(v_i)$ whenever $c_{i^*}(v_{i^*}) > 0$ (and otherwise the seller should keep the object). This is achieved in a direct mechanism implementable in dominant strategy in which bidder i^* is then required to pay $\max_{j \neq i}\left[c_i^{-1}\left(c_j(v_j)\right), c_i^{-1}(0)\right)\right]$. The manipulative auction design that mixes this format (probability $1-\varepsilon$) with the format in which if bidder i announces a type $v_i > c_i^{-1}(z)$ he makes a lump sump payment of $p_i(z,\varepsilon)$ and the format anonymous feedback partition is considered achieves a strictly larger revenue that Myerson's optimal revenue for suitable choices of z and ε (close enough to 0 and $L+1$, respectively) and a suitable choice of $p_i(z,\varepsilon)$ (so that type $c_i^{-1}(z)$ is indeed indifferent between bidding z and bidding

B).[32]

2) *The general case.* We note that the function $x \to x - \frac{1-F_i(x)}{f_i(x)}$ is always strictly increasing for x sufficiently close to $L+1$. We can then proceed similarly as in 1) from the upper intervals of types for which there is no bunching. **Q. E. D.**

Proof of Claim 4: This is shown by computing the difference between R_g and R_{g^*} where

$$G^+(v) = \left| \begin{array}{c} G(x) \text{ for } x < v < x + \varepsilon \\ G(v) \text{ for } v < x \text{ or } x + \varepsilon < v \end{array} \right.$$

There is a revenue gain for all $v > x + \varepsilon$, which in aggregate writes (up to order ε^2 terms):

$$Gain \approx \int_x^{L+1} \frac{F(v)f(v)}{G(v)} dv \cdot g(x) \frac{\varepsilon^2}{2}$$

There is a revenue loss for all $v \in (x, x + \varepsilon)$, which in aggregate writes (up to order ε^2 terms):

$$Loss \approx \frac{g(x)F(x)f(x) \int_L^x G(v)dv}{(G(x))^2} \frac{\varepsilon^2}{2}$$

If $g(x) > 0$ and $\int_x^{L+1} \frac{F(v)f(v)}{G(v)} dv > \frac{F(x)f(x) \int_L^x G(v)dv}{(G(x))^2}$, G^+ is the local change that improves expected revenue. If $\int_x^{L+1} \frac{F(v)f(v)}{G(v)} dv < \frac{F(x)f(x) \int_L^x G(v)dv}{(G(x))^2}$, then G^- will do where

$$G^-(v) = \left| \begin{array}{c} G(x) \text{ for } x - \varepsilon < v < x \\ G(v) \text{ for } v < x - \varepsilon \text{ or } x < v \end{array} \right.$$

Q. E. D.

[32] This is shown exactly as in the two-bidder symmetric and regular case noting that $\frac{\partial p_i}{\partial z} = 0$ whenever $z = L + 1, \varepsilon = 0$.

References

[1] Athey, S. and P. Haile (2006): 'Empirical Models of Auctions,' Ninth World Congress of the Econometric Society.

[2] Esponda, I. (2007): 'Information Feedback in First Price Auctions', mimeo.

[3] Ettinger, D. and P. Jehiel (2007): 'A Theory of Deception,' mimeo.

[4] Fudenberg D. and Levine, D. K. (1998): *The Theory of Learning in Games*, MIT Press.

[5] Hafalir I. and V. Krishna (2007): 'Asymmetric Auctions with Resale,' forthcoming *American Economic Review.*

[6] Harsanyi, J. C. (1995): 'Games with Incomplete Information,' *The American Economic Review*, **85**, 291-303.

[7] Huck, S., P. Jehiel, and T. Rutter (2007): 'Learning Spillovers and Analogy-based Expectations: A Multi-Game Experiment,' mimeo.

[8] Jackson, M. and E. Kalai (1997): 'Social Learning in Recurring Games,' *Games and Economic Behavior*, **21**, 102-134.

[9] Jehiel, P. (2005): 'Analogy-based Expectation Equilibrium', *Journal of Economic Theory*, **123**, 81-104.

[10] Jehiel, P. and F. Koessler (2007): 'Revisiting Games of Incomplete Information with Analogy-based Expectations', forthcoming *Games and Economic Behavior.*

[11] Maskin E., and J. Riley (2000): 'Asymmetric Auctions,' *The Review of Economic Studies*, **67**, 413-438.

[12] Myerson, R. (1981): 'Optimal Auction Design,' *Mathematics of Operations Research*, **6**, 58-73.

[13] Riley, J. and W. Samuelson (1981): 'Optimal Auctions,' *The American Economic Review*, **71**, 381-392.