Manipulative Auction Design*

Philippe Jehiel[†]

26th January 2009

Abstract

This paper considers an auction design framework in which bidders get partial feedback about the distribution of bids submitted in earlier auctions. A feedback equilibrium is proposed to model the long run interactions of bidders in such environments with partial feedback. It is shown that 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 private values two-bidder case with independent distributions. It is also shown how by using several auction formats with coarse feedback a designer can always extract more revenues than in Myerson's optimal auction, and yet less revenues than in the full information case whenever bidders enjoy ex-post quitting rights and the assignment and payment rules are monotonic in bids. These results suggest an important role of feedback disclosure as a novel instrument in mechanism design.

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

*I would like to thank I. Esponda, L. Lamy, R. Spiegler, B. Wilson, J. Wooders, A. Wolinsky, and seminar participants at Roy seminar (Paris), Northwestern University, the conference in honor of Ehud Kalai (Jerusalem, 2007), the Cemmap conference on auctions (UCL, 2008), the workshop on biased beliefs (UCL, 2008), SWET (Madrid, 2008) and the third world congress of game theory for useful comments.

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

players' strategies as a function of their signals (see Harsanyi (1995)). Yet, this requires 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 of similar objects (say, iphones) involving each time new bidders of observable characteristics $i = 1, 2, ...n$ (say, bidders of different age). Assume that every bidder knows his own valuation for the object but not those of other bidders, and assume that the distribution of valuation depends on the observable characteristic i. In many contexts, it seems unlikely that bidders would *a priori* know these distributions and how they depend on i. In such cases, bidders would look at the bids submitted in earlier similar auctions so as to form a judgement as to what the distribution of others' bids is likely to be in the current auction of interest.[1] In many practical auction designs such as those used on ebay or treasury auctions, bidders have access only to the aggregate distribution of bids in past auctions without being informed of the characteristics of the bidders of the corresponding bids. It is then dubious that bidders would be able to play a best-response to the actual distribution of bids of the other bidders because there is no way a bidder can assess the distribution of bids conditional on the observable characteristic based on the feedback he receives. Instead, bidders are more likely to play a best-response to the conjecture that all bidders - no matter what their characteristic is - bid according to the aggregate distribution of bids that mix the distribution of bids of all bidders. In the long run, assuming convergence of the overall process, bidders are not playing a Nash equilibrium but what I call a feedback equilibrium.[2]

As an alternative example, consider promotions in an organization. Each promotion takes the form of a contest between several employees who each make a proposal for the job task specification in case of promotion. The private interest for the promotion may vary across employees, and the criteria for promotion as well as the specification of the

[1]Typically, on ebay, one has access to the history of bids made in auctions of similar objects run in the last month.

[2]In the language of econometrics, the model is not identifiable. Most applied econometricans would particularize the model making an extra symmetry assumption (see Athey and Haile (2006)), which is in line with the feedback equilibrium.

job task (based on the proposals) may change from one promotion to the next. Yet, the criterion for the current promotion is typically known to the contestants. Nash equilibrium would require that contestants know the distribution of others' effort (interpreted here as the contestant's specification of the job task) given the criterion applying to the current promotion. Assuming that contestants form their expectations by looking at past promotions, this would require that employees have access to the joint distribution of effort and corresponding criterion. Yet, if employees only have access to the distribution of effort (and not of the corresponding criterion), contestants will not be able to play a Nash equilibrium. Instead, I propose that contestants play a best-response (given the current criterion) to the conjecture that the distribution of effort is the same irrespective of the criterion (and that it corresponds to the aggregate distribution of effort they have access to).

In this paper, I generalize the above two examples. I consider one-object private values auction environments in which the valuations are independently distributed across bidders (the promotion example can easily be cast in terms of an auction by identifying the effort with the bid and the private interest for the promotion with the valuation). I first propose an equilibrium concept to describe the long run interaction of bidders in situations in which bidders receive coarse feedback about the distribution of previous bids, and every single bidder participates in just one auction. Specifically, my framework allows me to consider situations in which, as in the first-price auction example, only the aggregate distribution of bids with no reference to the characteristics of bidders is disclosed, and also situations in which, as in the promotion example, different auction formats are being used (this is the analog of different criteria being used) and only the aggregate distribution of bids across the various auction formats is being disclosed. The equilibrium obtained - which stands for the limiting outcome of a learning process with corresponding feedback disclosure - is called a feedback equilibrium,[3] and it requires that bidders play a best-response to the aggregate distribution of bids, as given by the feedback they receive.

I explore whether and when it is the case that the designer is better off when bidders play a feedback equilibrium rather than a Nash equilibrium and how the answer is affected

[3]It is parameterized by the form of the feedback described as a partition of the set of profile of format and bidders' characteristics received by the players at the learning stage, see Section 2.

by the specific forms of feedback disclosure and auction rules. Addressing such questions opens new avenues in mechanism design. As it turns out, providing coarse feedback about past behaviors - as considered in this paper - can enhance the designer's objective, which suggests the relevance of an instrument not previously considered in mechanism design.[4]

To highlight the potential role of feedback disclosure in mechanism design, I assume that (in addition to the format(s)) the designer is free to choose which kind of feedback (within the class specified above) to disclose to bidders. In the promotion example, choosing which feedback to disclose sounds natural for the designer given that the feedback is in principle under the control of the organization. In the auction example, this is an idealization that would fit well if a single seller had many similar objects to auction off over time. My framework in the auction case can more usefully be interpreted as providing insights as to what kind of feedback policy an auction house such as ebay should adopt so as to increase the revenues of sellers (which in turn determine the revenues of the auction house through the fees).

The first question I address is as follows. Suppose the main objective of the designer is welfare maximization whereas the auxiliary goal is revenue, as is the case in many government auctions. Can the designer do better than using a second-price auction (or equivalently an ascending-price auction) by using a coarse feedback device?

In the classic setup (relying on Nash equilibrium), the so called revenue equivalence theorem implies that the designer can do no better. This is so because 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 (see, for example, Milgrom (2004) for an exposition of the revenue equivalence theorem).

In my setup, I show that the designer can sometimes do better, thereby illustrating a failure of the revenue equivalence theorem when the solution concept is the feedback equilibrium. Specifically, in the case of two bidders with asymmetric distributions of valuations, I show that the first-price auction in which the designer provides as feedback the aggregate distribution of bids with no reference to the characteristic of the bidders (as considered above) always induces an efficient outcome and always generates an expected

[4]For a related investigation in the context of symmetric first-price auctions with affiliated signals, see the independent work of Esponda (2008a), which is discussed in subsection 5.1.

revenue that is strictly larger than that of the second-price auction no matter what the distributions of valuations are. Such a 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 with asymmetric bidders (see Maskin and Riley (2000)).[5]

The second question I address is as follows. Suppose bidders can always veto the transaction ex-post and suppose that the designer is solely interested in revenue maximization. Can the designer generate more revenues than in the classic optimal auction analyzed in Myerson (1981) and Riley and Samuelson (1981)?

I show that this is always so and that the best revenue so obtained lies strictly below the full information maximal revenue. In other words, a clever use of coarse feedback coupled with a judicious choice of format(s) may reduce the informational rent left to bidders even though some rent must be left to bidders (given that the full information maximal revenues cannot be reached).

The rest of the paper is organized as follows. In Section 2 the model is introduced together with the feedback equilibrium and the mechanism design problem. In Section 3, the first-price auction with bidder-anonymous feedback partition is analyzed. In Section 4, it is shown how one can generate more revenues than in Myerson's optimal auction. Section 5 offers a discussion including how the paper relates to the literature. Section 6 concludes.

2 Basic definitions

An object is to be auctioned off and there are n bidders $i = 1, ...n$. 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 $[c, d]$ and (continuous) density $f_i(\cdot)$ where $f_i(v) > 0$ for all $v \in [c, d]$ and $d > 0$. Bidders have quasi-linear preferences and they are risk neutral. That is, if a bidder

[5] Hafalir and Krishna (2008) 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 thus there is no room for resale.

with valuation v expects to win the object with probability p and expects to make (an expected) transfer t to the seller, his expected utility is $pv - t$. The seller's valuation is set at 0.

The object is auctioned off through possibly different auction formats M_k, $k \in K = \{1, ...r\}$ where format M_k is used with probability λ_k. Each auction format M_k (which together with λ_k is chosen by the designer) is restricted to take the following form:

- Bidders $i = 1, ...n$ simultaneously submit a bid $b_i \in [c, d]$.
- Based on the profile of bids $b = (b_i)_{i=1}^{i=n}$, with probability $\varphi_i^k(b)$, bidder i is offered to get the object in exchange for a payment $\tau_i^k(b)$ to the seller. The transaction takes place if bidder i approves the terms of the contract.
- For every i, k, $\varphi_i^k(b)$ is a non-decreasing function of b_i and a non-increasing function of b_j, $j \neq i$; $\tau_i^k(b)$ is a non-decreasing functions of b_i and of b_j, $j \neq i$.

The auction formats I consider require that the bidders approve the terms of the contract ex-post. This ensures that a bidder no matter what bidding strategy he employs is better off participating in the auction rather than staying outside (see Compte and Jehiel (2007) for further elaborations on ex-post quitting rights in mechanism design). The monotonicity assumptions made on $\varphi_i^k(b)$ and $\tau_i^k(b)$ are in line with the auction interpretation: As a bidder increases his bid, his probability of winning (weakly) increases as well as the price this bidder must pay. As a competing bidder increases his bid, the probability of winning (weakly) decreases and the price paid in case of transaction (weakly) increases. It should be mentioned that the first-price auction, and the second-price auction both belong to this class of auctions, and that is always possible to pick an auction in this class that achieves Myerson's optimal revenue.

When auction format M_k prevails, bidder i is informed of the functions $\varphi_i^k(b)$ and $\tau_i^k(b)$ that apply to him in this format. 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 in M_k is:

$$u_i^k(v_i, b_i; \widetilde{b}_{-i}) = \underset{\widetilde{b}_{-i}}{E} [\varphi_i^k(b_i, b_{-i}) \max(v_i - \tau_i^k(b_i, b_{-i}), 0)]$$

A strategy of bidder i is 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 bidder i's bid in format M_k when i's valuation is v_i.[6]

Nash equilibrium requires that for each k and v_i, bidder 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, bidders are not assumed to know (or have access to) β_i^k for every i and k. Instead, each bidder i receives partial feedback about the distribution of bids observed in past auctions. They play a best-response to this feedback (in a sense to be defined next) and a steady state is assumed to have been reached.

Specifically, I consider the following class of partial feedback. Each bidder 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 bidder i. A typical element of $\boldsymbol{P}_i$ is denoted by α_i and referred to as a feedback class of bidder i. The element of $\boldsymbol{P}_i$ containing (j,k) is denoted by $\alpha_i(j,k)$. The interpretation of $\boldsymbol{P}_i$ is that bidder i gets only informed of the empirical distribution of bids of bidders engaged in past auctions where all bids b_j submitted in M_k with $(j,k) \in \alpha_i$ are treated alike (i.e., they are not distinguished).

I further assume that a steady state has been reached so that the empirical distributions of previous bids correspond to the actual distributions (applying now and in the future). When making his choice of strategy in auction format M_k, bidder i is thus assumed to know only (in addition to $\varphi_i^k(b)$ 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 bidder j in format M_k bids according to the aggregate distribution of bids in $\alpha_i(j,k)$, the feedback class to which (j,k) belongs.

Formally, let $\boldsymbol{A} = (M_k, \lambda_k, \boldsymbol{P}_i)_{i \in I, k \in K}$ denote an auction design. The solution concept is defined as:

[6]Strictly 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.

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 $\overline{\beta}_{-i}^k = (\overline{\beta}_j^k)_{j\neq i}$, *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* $(j',k') \in \alpha_i(j,k)$, *and the distributions* $\overline{\beta}_j^k$, $j \neq i$ *are perceived by bidder* i *to be independent of each other.*[7]

Remarks. 1) It should be mentioned that the feedback received by bidders is about the distribution of individual bids and not about the distribution of bid profiles.[8] 2) The feedback equilibrium will be further interpreted in subsection 2.2. It will be related to other approaches in the literature in particular the analogy-based expectation equilibrium and the self-confirming equilibria in subsection 5.1.

Various objectives for the designer will be considered. The first objective will be a lexicographic criterion with welfare ranked first and revenues ranked second. The second objective will be revenues. In all cases, the designer is assumed to be risk neutral (and as already mentioned to have no intrinsic value for the object). She chooses the auction design $\boldsymbol{A} = (M_k, \lambda_k, \boldsymbol{P}_i)_{i\in I,k\in K}$ so as to maximize the expected value of her objective assuming that bidders behave according to a feedback equilibrium of $\boldsymbol{A}$.[9]

2.1 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.

[7] The independence assumption is in line with the non-cooperative interpretation of the solution concept.

[8] Accordingly, every bidder i treats every bidder j's distribution of bids, $j \neq i$, as being independent of each other. This fits in well when bidders do not have access (or pay attention to) whether the past bids they observe were submitted at the same or different times.

[9] As in Myerson (1981), one also implicitly assumes that the designer can choose the feedback equilibrium she likes best. Yet, which feedback equilibrium is played turns out to be inessential provided weakly dominated strategies are never played.

1) *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. This is the situation studied in Section 3.

2) *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 the distribution of bids for the various bidders $i \in I$. That is, for all $i \in I$, $\boldsymbol{P}_i = \left\{ \bigcup_{k \in K} \{(j,k)\} \right\}_{j \in I}$. This situation corresponds to the contest application mentioned in introduction, and it will be considered in Section 4.

2.2 Interpretation

The interpretation of a feedback equilibrium is that it stands for the limiting outcome of a learning process in which 1) at every stage there is a new auction and new bidders of observable characteristics $i \in I$;[10] 2) auction format M_k is used with frequency λ_k and 3) a bidder with characteristic i receives as feedback the aggregate empirical distribution of past bids in every α_i.[11] If behaviors stabilize in such a learning process, it must be to a feedback equilibrium provided bidders consider the simplest theory that is consistent with the feedback they receive (see more below on the interpretation of the simplest theory).

The mechanism design perspective considered in the paper corresponds to the idealization that a single 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 $\boldsymbol{A}$. As mentioned in introduction, such a view is appropriate in situations in which a single seller repeatedly sells similar objects or in situations in which a single organization repeatedly organizes contests for promotion. It may also be useful to understand the incentives of an auction house such as ebay which is

[10]The profile of bidders' charateristics is assumed to stay the same throughout the process (see the discussion section for some elaboration on the case in which the number of bidders may vary stochastically).

[11]That is, a bidder with characteristic i is informed of the aggregate empirical distribution of past bids $\{b_j^k, (j,k) \in \alpha_i\}$ with no reference to which (j,k) generated the bid.

obviously interested in increasing the sellers' revenues (through the fees these generate) and which can control both the formats and the feedback about past auctions that is disclosed to bidders.

The approach developed in this paper has a non-Bayesian element in the sense that upon learning the coarse feedback the designer reports to them, bidders do not update their belief about the distribution of others' bids based on some (possibly subjective) prior. Instead, bidders are assumed to consider the simplest theory consistent with the feedback they receive: They play a best-response to the conjecture that the distribution of bids is uniformly the same over the various (j,k) that are bundled in the same feedback class.[12] I believe this is a natural assumption in many practical situations of interest in which 1) subjects would have no other data than the feedback they receive to form their prior and[13] 2) the functions φ_j^k and τ_j^k that govern bidder j's incentive in format M_k are either the same across all (j,k) that belong to the same feedback class or they are not known to subjects with characteristic i, $i \neq j$.[14] In such situations, it would seem rather hard (in fact impossible) for subjects to understand how the distribution of bids vary across different (j,k) that belong to the same feedback class simply based on the aggregate empirical distribution they are informed about: assuming that the distribution is the same across these various (j,k) seems focal, and I propose it gives good account for bidders' mode of thinking in such situations.[15]

[12] Technically speaking, the feedback equilibrium can be viewed as a refinement of the self-confirming equilibrium in which bidders adopt the "simplest" conjecture as just described.

[13] This in particular requires that bidders have no prior knowledge about the densities f_i not even about the frequencies λ_k.

[14] In the first-price auction with bidder-anonymous feedback partition the former property applies (ϕ_j^k and τ_j^k are the same across all bidders). And it may be argued that in the contest for promotion example mentioned in introduction the latter property is often met given that the criteria used for different promotions are often not very transparent to outsiders.

[15] Such an assumption is very much in line with how applied econometricans operate (see footnote 2 in introduction). From another perspective, the finding in Huck et al. (2008) gives support to such an assumption. There, when subjects did not know their opponents' payoffs and received as feedback the aggregate empirical distribution over two possible games, they played as if this aggregate distribution applied to every single game.

2.3 Preliminaries

A few preliminary observations follow. First, by picking a single auction format M and by adopting the finest feedback partition, the designer can always replicate the revenue generated in the standard rationality case in M. Thus, if the designer seeks to maximize revenues, she can always achieve a revenue at least as large as Myerson (1981)'s optimal revenue. The question is whether she can achieve larger revenues.

Second, consider an auction format M in which player i has a dominant strategy. Then in any auction design including format M, a feedback equilibrium requires that bidder i plays his dominant strategy in M. This is a straightforward observation, since bidder i will find his strategy best no matter what his expectation about the distribution of others' bids is, and thus no matter how the auction design is further specified.

Third, one of the auction designs that will be studied falls in the following class. 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 $\Gamma^{ba}(\boldsymbol{A})$ defined by the auction format M in which the distribution of bidder i has the average density $\overline{f}(v_i) = \sum_{i \in j} f_j(v_i)/n$ instead of $f_i(v_i)$.

Claim 1: A symmetric strategy profile is a feedback equilibrium of $\boldsymbol{A}$ if and only if it is a Bayes Nash equilibrium of $\Gamma^{ba}(\boldsymbol{A})$.

Fourth, another class of auction designs $\boldsymbol{A}$ considered below is such that the various auction formats M_k in $\boldsymbol{A}$ satisfy: $\varphi_i^k(b_i, b_{-i}) = \varphi_i(b_i, b_{-i})$ 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 such auction designs $\boldsymbol{A}$ to the Nash Bayes equilibria of the following game referred to as $\Gamma^{fa}(\boldsymbol{A})$:

Game $\Gamma^{fa}(\boldsymbol{A})$: 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.[16]

Claim 2: Suppose that the format anonymous feedback partitions prevail and that in all auction formats M_k of $\boldsymbol{A}$, $\varphi_i^k(b_i, b_{-i}) = \varphi_i(b_i, b_{-i})$ 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 Bayes Nash equilibrium of $\Gamma^{fa}(\boldsymbol{A})$.

3 Efficiency and revenues

Assume the designer is interested both in efficiency and revenues, and suppose that the primary objective of the designer is efficiency while revenue is only the secondary objective (this corresponds to the objective in many government auctions). 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 $v_i = c$ is 0 then both 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.

I 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 a manipulative auction design setup. Besides, this gain in revenues is achieved by using a fairly standard auction format (with the bidder-anonymous feedback partition).

Proposition 1 *Assume that all valuations are non-negative, i.e.* $c \geq 0$*, and consider a two bidder* $i = 1, 2$ *auction setup with asymmetric distributions* ($F_1(\cdot) \neq F_2(\cdot)$ *on a set of strictly positive measure). There is a unique feedback equilibrium of the first price auction with anonymous bidder feedback partition. Moreover, this feedback equilibrium induces an efficient outcome and it generates a strictly higher revenue than the second-price auction.*

[16] Compared to the true auction design, the difference is that the methods of payments are independently distributed across bidders in Γ^{fa} whereas they are (perfectly) correlated in Γ.

The revenue gain is

$$\int_c^d \frac{1}{4}\left(F_1(v)-F_2(v)\right)^2 dv+\int_c^d \frac{1}{4}\frac{d\beta(v)}{dv}\left(F_1(v)-F_2(v)\right)^2 dv>0$$

where $\beta(v)=\int_c^v x\overline{f}(x)dx/\overline{F}(v)$, $\overline{f}(x)=\frac{f_1(x)+f_2(x)}{2}$ *and* $\overline{F}(v)=\int_c^v \overline{f}(x)dx$.[17]

What is the intuition for the above result ? First, observe that the use of the bidder-anonymous feedback partition leads the bidders (whatever their characteristic) to best-respond to the same distribution of bids, which given the anonymous character of the first-price auction, ensures efficiency. Second, the use of the bidder-anonymous-feedback partition 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 various bidders (this essentially follows from claim 1 above). 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 enhances revenues because it makes the strong bidder feels the weak bidder is less weak than he really is.[18]

When there are more than two bidders, the first-price auction with anonymous bidder feedback partition remains efficient, but the revenue comparison with the second-price auction can go either way depending on the form of the asymmetry of the distributions.[19] When the distributions of valuations are nearly the same across bidders (say the cumulative functions differ up to ε), then the revenues of the two auction designs differ according to

[17] $\beta(\cdot)$ is the equilibrium bid function in a symmetric two-bidder FPA with density of valuations $\overline{f}$. As such, $\beta(\cdot)$ is an increasing function.

[18] This is essentially what step 2 of the proof formalizes. Step 3 on the other hand formalizes the idea that making the distribution more asymmetric across bidders (moving from $(\overline{f},\overline{f})$ to (f_1,f_2)) makes the distribution of the highest bid among $i=1,2$ more skewed toward larger values.

[19] 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. To see that revenues can go either way, consider first a situation with two bidders whose distribution of valuations is concentrated around d and a third bidder whose distribution of valuation is concentrated around c. 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 d). Consider next a situation with one bidder whose distribution of valuations is concentrated around d while other bidders have a distribution of valuations concentrated around c, the first-price auction with bidder anonymous feedback partition generates more revenues than the second-price auction (which generates a revenue very close to c). Thus the revenue comparison can go either way.

a smaller magnitude (of order ε^2).[20] And when there is a large number of bidders (say the set of bidders is replicated) then the second-price auction generates slightly more revenues even if the difference of revenues is negligible in the limit.[21] Clearly, more work is required to sign the revenue comparison between the two auction designs in the general case with more than two bidders. But, the above result indicates important circumstances (when there are few competitors) when the first-price auction with bidder-anonymous feedback partition should be preferred.

The above insight is distinct even though related to Myerson's insight about how to increase revenues in asymmetric auctions. An important implication of Myerson's analysis is that in the asymmetric case, competition between bidders should be biased in favor of weak bidders in order to increase revenues. The net effect of such biased auctions is that some inefficiencies are induced letting the weak bidder sometimes win the object. As Proposition 1 shows, the use of the bidder-anonymous feedback partitions allows somehow to symmetrize (a bit) the competition without sacrificing on efficiency. Of course, this is achieved by moving away from the full rationality paradigm, which the use of partial feedback permits.

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 $\beta^*(v) = \frac{\int_c^v xf^*(x)dx}{F^*(v)}$ where $F^*(v) = (F_1(v)F_2(v))^{\frac{1}{2}}$ and $f^*(v) = \frac{dF^*(v)}{dv}$. Such bidding strategies would always generate higher revenues than in the second price auction, as in Proposition 1.[22]

[20] When the distributions are very asymmetric, the difference of revenues can be quite substantial. For example, in the two-bidder case considered in Proposition 1, assume that the distribution of valuations of one bidder is concentrated around d whereas the distribution of valuations of the other bidder is concentrated around c. In this case, the first-price auction with anonymous bidder feedback partition provides a revenue gain of $\frac{d-c}{2}$, which should be compared with the revenue c of the second-price auction. Clearly, as d gets large relative to c, the revenue gain can be quite substantial in such asymmetric setups.

[21] Adding just a few more bidders in the first-price auction with anonymous bidder feedback partition would reverse the revenue ranking.

[22] 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

4 Revenues

Assume now that the designer seeks to maximize revenues. Our first main observation is:[23]

Proposition 2 *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 (denoted by R^M).*

The intuition for Proposition 2 is as follows. Myerson's optimal auction can always be implemented in such a way that every bidder has a (weakly) dominant strategy and ex-post quitting rights of bidders are fulfilled (think of the second-price auction in the symmetric regular case). One can now think of an auction design in which this auction format - call it MD - is mixed with a little bit of first-price auction with 0 reserve price (FPA), and bidders get only to know the aggregate distribution of bids over the two auction formats. In format MD, the strategies are the same as in the standard case (because bidders have a weakly dominant strategy in MD). In format FPA, bidders play a best-response to the aggregate distribution of bids over the two formats. For many choices of MD, this construction need not deliver revenues higher than R^M.[24] But, there are many variants of MD in which submitted bids are first transformed before the original format is being applied. For a suitable choice of such a variant, the construction leads bidders in FPA to bid very aggressively because they are led to think that by shading too much their bid the chance of winning in FPA gets too small. In the limit, bidders

with symmetric bidders and density of valuations $f^*(\cdot)$, one can establish that

$$\begin{aligned} \overline{R}^* - R^{SPA} &= \int_c^d \left[\sqrt{F_1(v)} - \sqrt{F_2(v)}\right]^2 dv > 0 \\ R^* - \overline{R}^* &= 0 \end{aligned}$$

[23]By inspecting the proof of Proposition 2, one can see that all feedback equilibria (not employing weakly dominated strategies) of the auction design considered there are such that the designer obtains higher revenues than in Myerson's optimal auction. Thus, the conclusion of Proposition 2 would hold under the stronger full implementation requirement (provided one restricts oneself to equilibria not employing weakly dominated strategies).

[24]In the case of uniform distributions, a mix of second-price auctions and first-price auctions would have no effect on revenues.

may be induced to bid very close to their valuation. Given that such bidding strategies in FPA induce a revenue close to the full information optimal revenue $R^F = E(\max_i(v_i, 0))$, and given that $R^F > R^M$, the result of Proposition 2 follows.

Proposition 2 establishes that the designer can do better than using Myerson's optimal auction (with fine feedback), but how much can she gain? Clearly, the best revenue that the designer can extract in auction designs with ex-post quitting rights can never exceed the full information optimal revenue R^F in which for all realizations $(v_i)_i$ of the valuations, the seller would extract a revenue equal to $\max_i(v_i, 0)$. This trivially follows from the observation that a winner of the auction would always object if he were asked to pay more than his valuation. As it turns out, the designer's best revenue lies strictly below R^F whenever bidders' valuations can lie below the seller's valuation.

Proposition 3 *The largest revenue that the designer can achieve in a manipulative auction design with ex-post quitting rights is strictly smaller than the full information optimal revenues R^F if $\Pr(v_i < 0) > 0$ for all i.*

In the proof of Proposition 2 some (Myerson-optimal) mechanism MD implementable in dominant strategy was mixed with a little bit of FPA, and the revenue obtained in FPA was shown to be close to R^F. However, such a construction required that the weight put on FPA was set sufficiently small. As one increases the frequency of FPA, the manipulation loses its force, and, of course, in the limit as the designer almost always picks FPA, one gets the standard revenue generated in the first-price auction, which following Myerson's analysis cannot be larger than R^M.

What Proposition 3 establishes is that within the class of mechanisms under study one can never reach R^F whatever the manipulation. To get an intuition for this result, think of a symmetric scenario in which the auction design $\boldsymbol{A}$ uses the format-anonymous feedback partition. To get close to R^F, it would be required that in all auction formats M_k used in $\boldsymbol{A}$ every bidder pays a price close to his valuation when he wins. Consider those auction formats in $\boldsymbol{A}$ for which the feedback equilibrium bid of a bidder with valuation $\frac{d}{2}$ is above the median bid of bidders with the same valuation $\frac{d}{2}$ across the various formats in $\boldsymbol{A}$. A bidder with valuation $v = d$ can consider deviating to $\beta_i^k(\frac{d}{2})$ in such formats. He should expect to win at least half of the time whenever $\max_{j \neq i} v_j < \frac{d}{2}$, and pay at most $\frac{d}{2}$ (so that the ex-post quitting right of a bidder with valuation $\frac{d}{2}$ is satisfied).[25] It follows that this bidder must perceive to get a payoff at least as large as $\frac{1}{2}\Pr(\max_{j\neq i} v_j < \frac{d}{2})(d-\frac{d}{2})$, which is strictly positive.[26] But, by following his equilibrium strategy, bidder i with valuation d cannot perceive to make a non-negligible profit if a revenue close to R^F is to be obtained. This follows from the monotonicity of the payment rule and the observation that to get close to R^F a bidder with positive valuation should win when all other bidders' valuations are negative and pay a price close to his valuation. The above observations together lead to a contradiction, thereby yielding the desired conclusion.

Comment. If the designer were allowed to commit to offering positive payments to losers and if the payments from the winner were not assumed to be monotonic in bids, then the designer could get a revenue close to R^F while still preserving the ex-post participation constraints of bidders.[27] Our restriction on mechanisms (i.e., not allowing positive payments to losers and imposing that payments from winners be monotonic in bids) can then be thought of as resulting from the regulatory desire to protect bidders from manipulation.

[25] The half of the time comes from the fact in at least half the formats $M_{k'}$, $\beta_i^{k'}(\frac{d}{2}) < \beta_i^k(\frac{d}{2})$.

[26] This is so because in all other events, this bidder must perceive to make non-negative profits given his quitting rights.

[27] To see this, consider a symmetric two-bidder scenario in which bidders' valuations are identically distributed on (c,d). Consider an auction design with format-anonymous feedback partition and two formats $\underline{M}$ and $\overline{M}$ used in equal proportion. In format $\underline{M}$, the equilibrium bids will lie in $[0,d]$; in format $\overline{M}$, the equilibrium bids will lie in $\{0\}\cup[d,2d]$. In each format, a bidder with negative valuation bids 0 in equilibrium. In both $\underline{M}$ and $\overline{M}$, the bidder with highest bid wins the auction if this bid is strictly positive. In $\underline{M}$, if $b_i \in (0,d)$ for $i=1,2$, the winner pays his own bid and the loser receives no transfer. In $\overline{M}$, if $b_i \in \{0\}\cup[d,2d]$ for $i=1,2$, the winner i^* pays $b_{i^*}-d$ and the loser receives no transfer. The idea is to augment the transfers in $\underline{M}$ and $\overline{M}$ to cover all bid profile configurations even for bid realizations that will never occur in the respective formats. So in $\underline{M}$, a (losing) bidder submitting $b_i \in (0,d)$ will be offered a promise of transfer $\underline{h}(b_i)$ if $b_j \in (d,2d)$ and in $\overline{M}$, a (winning) bidder submitting $b_i \in (d,2d)$ will be offered a transfer $\overline{h}(b_i)$ if $b_j \in (0,d)$. By suitable choices of $\underline{h}$ and $\overline{h}$, one can ensure that for $v_i > 0$ bidding $\underline{\beta}(v_i)=v_i$ in $\underline{M}$ and bidding $\overline{\beta}(v_i)=v_i+d$ in $\overline{M}$ is a feedback equilibrium. [For example, in the uniform distribution case, $\underline{h}(b)=\frac{b^2}{2d}-\frac{bc}{d}$ and $\overline{h}(b)=\frac{(b-d)^2}{2d}-\frac{(b-d)c}{d}$. These functions are determined so that the expected perceived transfers correspond to those that would be made in a SPA with 0 reserve price.] With such bidding strategies, the expected revenues generated in each format are R^F, and thus the designer gets R^F in expectation.

5 Discussion

5.1 Related literature

This paper is related to several strands of literature. First, the feedback equilibrium is very 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 (2009). 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 I 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 pursued 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).

A feedback equilibrium (as well as an analogy-based expectation equilibrium) can also be viewed as a self-confirming equilibrium in which bidder i would receive as signal the aggregate distribution of bids in α_i for the various $\alpha_i \in P_i$ (see Fudenberg and Levine (1998) for a general presentation of the self-confirming equilibrium). From this perspective, a feedback equilibrium (as an analogy-based expectation equilibrium) is a selection of self-confirming equilibrium in which the theory adopted by bidders is the focal (or simplest) one in which the distribution of bids is assumed to be the same across the various elements (j, k) that belong to the same feedback class (see subsection 2.2 for further interpretation).[28]

To my knowledge, there is only paper that adopts a mechanism design perspective using the self-confirming equilibrium rather than the Nash equilibrium. This is the paper by Esponda (2008a). He considers first-price auctions in which the *same* bidders get involved over sequences 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

[28]The experiment in Huck et al. (2008) gives support to this selection, even though not in an auction setup.

Nash equilibrium.

Apart from the obvious difference that Esponda considers first-price auctions with symmetric bidders and affiliated signals whereas I consider general auction formats with arbitrary yet independent distributions of valuations, Esponda's result shares some similarities with the insight that partial feedback may help achieve greater revenues in first-price auctions (see Proposition 1 above).[29] Yet, there are notable differences between his framework and mine that I now discuss. First, Esponda considers a setting in which the same bidders keep participating in the auctions whereas I have in mind situations in which (as on ebay) new bidders arrive each time. This difference in turn explains why in my setting the feedback of bidders is not conditional on their own valuation whereas in Esponda's setting it is.[30] Second, Esponda's solution concept is the self-confirming equilibrium whereas I rely on the feedback equilibrium (a selection from the set of self-confirming equilibria, 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). It should also be mentioned that Esponda's 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 in his setup.[31] By contrast, my insight about the first-price auction concerns the case of asymmetric bidders with independent distributions, and the selection imposed by the feedback equilibrium ensures a strict superiority of providing partial feedback (see Proposition 1).

Some other equilibrium approaches that move away from Nash equilibrium (and thus permits erroneous expectations) have been proposed in the recent past. These include the

[29] The coarse feedback considered in Esponda's paper has no effect on revenues in a symmetric private values setup. Moreover, the insight that in asymmetric first-price auctions, the bidder-anonymous feedback partition can help achieve higher revenues (Proposition 1) has obviously no counterpart in Esponda's symmetric setup.

[30] 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 feedback about the highest bid only need not result in a revenue gain.

[31] I view this as potentially 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.

cursed equilibrium of Eyster and Rabin (2005) and the behavioral equilibrium of Esponda (2008b). These approaches shed new light on the winner's curse and on the possibility of trade in adverse selection problems. Yet, in private values setting such as the one considered here, these approaches coincide with Nash equilibrium and as such are not closely related to the present study (see Jehiel and Koessler (2008) for a discussion of the link between the cursed equilibrium and the analogy-based expectation equilibrium).

There is a strand of literature concerned with learning in mechanism design. This strand includes among others Cabrales (1999) and Cabrales and Serrano (2007) (see the latter for a more comprehensive review of that strand of literature). A typical question addressed by this literature is about equilibrium selection when there are several Nash equilibria and whether the choice of mechanism may induce good convergence properties of the corresponding learning process. The approach pursued here is complementary to this strand. It offers a different perspective by suggesting how the use of coarse feedback may result in convergence to non-Nash equilibria, i.e. feedback equilibria. The approach pursued here also assumes to start with that a steady state has been reached. In line with the literature just mentioned, it would be of interest to study the convergence properties of the learning process that was suggested to motivate the present approach.

Finally, there have been various approaches to study how a mechanism designer should deal with various behavioral biases assumed on agents. These include among others the work of Eliaz (2002) who assumes that a fixed number of agents may have a crazy behavior, the work of Eliaz and Spiegler (2007) who assume that agents may have erroneous subjective beliefs and the recent work of Crawford et al. (2008) who assume that bidders behave according to the level k-mode of thinking.[32] Beyond the obvious observation that the bias considered in this paper is of a different nature than the ones considered by these authors, I believe the current approach differs from these other approaches in an important way. Somehow the bias in the expectation formation that appears in a feedback equilibrium is induced by the choice of feedback partition made by the designer. It is thus as if the cognitive limitations of the bidders were endogenously created by the designer rather than being there to start with.

[32] In a related vein, Matsushima (2008) considers an implementation problem when agents rely on two levels of eliminations of dominated strategies.

5.2 Complete information

In the above analysis, some uncertainty about bidders' valuations was assumed. When each bidder i's valuation can take a single value v_i, the designer can extract a revenue equal to $R^F = \max_i(v_i, 0)$ in the classic rationality setup.

In the above setup with ex-post quitting rights, no manipulation can allow the designer to extract more than R^F given that a bidder will never accept to pay more than his valuation if he wins the object (and he will never accept to pay anything if he does not win the object). Thus, some private information is required for manipulation to be of effective use to the designer.

5.3 Interim participation constraints

In the above analysis, auction formats with ex-post quitting rights were considered. If participation constraints are only required at the interim stage before bidders know the outcome of the auction and if relatedly the designer can also require payments froms losers, then the designer can generate much larger revenues if bidders play according to a feedback equilibrium.

Proposition 4 *Suppose there are at least two bidders and that bidders cannot withdraw from the auction later on. Then by a suitable choice of auction design the designer can make arbitrary large amounts of money.*

The idea of the proof which is detailed in the appendix is as follows. By choosing several formats and by using a format-anonymous feedback partition for say bidder 1, the designer can mislead bidder 1 in his understanding of the distribution of bids of other bidders $i \neq 1$. She can then propose a bet to bidder 1 whose monetary outcome is contingent on the realization of b_i, $i \neq 1$, in such a way that the bet sounds profitable from the viewpoints of both bidder 1 and the designer. By increasing the stakes of the bet, bidder 1 will still agree on the terms of the bet given our assumption of risk neutrality, which translates into potentially arbitrarily large revenues for the designer.

The above argument bears strong resemblance with the observation that with subjective prior beliefs the logic of the no trade theorem breaks down.[33] Of course, here since the designer is assumed to know the correct distributions of bids, one makes the further inference that it is the designer (and not the bidder) who benefits from the bet. Another key difference with the literature on subjective priors is that the erroneous perception of the bidders is viewed here as resulting from the feedback manipulation of the designer and not from the subjective character of bidders' prior beliefs.

Comment. To the extent that bidders know that the designer is more informed than they are about the distributions of bids, one might argue that in the context of the above manipulation bidder 1 might be suspicious, thereby deciding to stay outside the auction room rather than playing according to a feedback equilibrium.[34] This is to be contrasted with auction designs with ex-post quitting rights as considered in the main part of this paper in which staying outside the auction room is always a bad idea (nothing worse can happen by participating). In the class of auction designs with ex-post quitting rights, it is not clear what else (i.e. other than playing according to a feedback equilibrium) a player could do.[35]

5.4 Shill bidding

In the above analysis, 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 a manipulative mechanism design setup it does, as I now illustrate.

[33] See however Morris (1994) for an exploration of when the no trade theorem continues to hold in the subjective prior paradigm.

[34] Alternatively, if one has in mind that there is a risk that bidder 1 would play according to a feedback equilibrium even for large stakes, one may think of the ex-post quitting rights scenario considered in the main part of the paper as a regulatory constraint imposed on designers to better protect bidders from manipulation.

[35] Recently, Lehrer (2007) has proposed a selection of self-confirming equilibrium based on the most pessimistic conjecture (rather than the simplest conjecture as in this paper). Such an approach would lead bidders not to accept bets as considered in the proof of Proposition 4. But, it would also lead bidders not to take part in any auction of the sort analyzed throughout this paper as long as feedback is partial and there is a slight cost to participate in auctions. I find the latter conclusion unrealistic.

Specifically, the designer is now assumed to be allowed to hire shill bidders who have no intrinsic value for the object. If the cost of hiring shill bidders is zero (this cost may in practice be high if shill bidding is illegal), the designer can get a revenue close to the full information revenue R^F (still assuming that bidders have ex-post quitting rights).

Proposition 5 *Suppose the cost of hiring shill bidders is zero. Then the designer can get a revenue close to* R^F *in the optimal auction design.*

The proof of Proposition 5 follows the logic used to prove Proposition 2. By inviting m shill bidders, the designer can make them bid as she wishes, say each according to a distribution of bids $g(.)$ with support on $(0, d)$. Consider now a variant of the first-price auction defined as follows. Only the bids of the real bidders $i = 1, ...n$ matter in this format, and the rules of the auction restricted to these bids are the same as the first-price auction. That is, the bidder $i \in I = \{1, ...n\}$ with highest bid wins the object as long as this bid is strictly positive and he pays his bid b_i. Other bidders make no payment. Consider now the bidder-anonymous feedback partition in the above auction format. It is readily verified that as m grows to infinity, bidders will submit a bid that is approximately a best-response to the distribution of bids $g(\cdot)$ of each of the other relevant bidder among $i \in I = \{1, ...n\}$. That is, each bidder with valuation $v > 0$ will submit a bid close to $\beta(v) \in \arg\max_b (v - b)G^n(b)$ where $G(\cdot)$ is the cumulative of $g(\cdot)$. By considering a cumulative $G(\cdot)$ of the form $G(v) = \left(\frac{v}{d}\right)^q$ with q large enough, one easily obtains that $\beta(v)$ gets close to v, thereby providing a proof of Proposition 5 (see more details in the proof of Proposition 2).

5.5 Random number of bidders

In the above analysis, the set of bidder I was deterministic (this was also to simplify the comparison with Myerson's optimal auction paper). How are our results affected if the set of bidders I is stochastic?

For concreteness, consider a symmetric regular case in which the valuation of every bidder i is drawn from the same distribution with density $f(\cdot)$ and $v \rightarrow v - \frac{1-F(v)}{f(v)}$ is increasing. When bidders are risk neutral as assumed in this paper, the best revenue in the classic rationality setup is achieved by having a regular auction (say a second price or first-price) auction with reserve price R set such that $R - \frac{1-F(R)}{f(R)} = 0$. This is so because such a format would achieve the best revenue even if the number of bidders were known to the designer and no matter what this number is (see McAfee and McMillan (1987) for the treatment of risk aversion when the number of bidders is stochastic).

Can the designer achieve larger revenues in a manipulative auction design setup when the number of bidders is stochastic? The answer is yes and this is shown in the same way as Proposition 2 was proven, that is, by mixing a little bit of first-price auction with mostly a well chosen auction format that is strategically equivalent to the second-price auction with reserve price R, and by considering the format-anonymous feedback partition. Similarly, Proposition 3 extends to the stochastic number of bidder case.

In a vein similar to that of Proposition 1, one might also be interested in scenarios in which the number of bidders would vary from one auction to the next and bidders would observe the number of competitors they face in the current auction. If a second-price auction is considered, observing the number of competitors has no effect on the optimal strategy, but if a first-price auction is considered such an observation will affect the optimal strategy (as it is indicative of how the chance of winning depends on the bid shading). In the latter case, one may also consider the effect of providing as feedback the distribution of past bids without telling how many bidders were present in the auction room when the bid was submitted. How does the revenue generated in such a first-price auction design compare to the revenue of the second-price auction? Addressing such a question is left for future research.

5.6 Cheating on feedback

An important assumption made throughout the paper is that the feedback reported by the designer must be correct. There are important reasons why I believe this is a natural assumption. First, in most countries it would be illegal to report false pieces of information (this should be contrasted with the kind of manipulation considered in this paper in which every released information is correct even if partial). Second, even if there is no legislation about the correctness of feedback, it is likely to be in the interest of sellers to report truthful feedback, as otherwise if bidders realize feedback may be erroneous there is no reason why bidders would trust the feedback that is transmitted to them. Along this line, it may be argued that it is in the interest of an auction house such as ebay to being committed to never report false feedback to bidders (and it seems clear that no one would dispute the correctness of the feedback provided by ebay to bidders).

5.7 Other forms of feedback

Even if the class of feedback considered in this paper is quite large, some forms of feedback are not covered.[36] A key difficulty should be addressed though if one wishes to consider more general classes of feedback. That is, one need to define an appropriate/focal notion of best-response to the feedback received by the bidder, thereby leading to an appropriate notion of equilibrium. While this may be defined in some cases beyond the class of feedback considered here (see the discussion at the end of Section 4), I believe this would be often problematic for general forms of feedback as the feedback would not easily translate into a focal conjecture about other players' strategies. And considering the whole set of self-confirming equilibria is unlikeley to give sharp predictions given in particular that the Nash equilibrium always belongs to this set whatever the feedback.

5.8 The replacement assumption

To motivate the feedback equilibrium, we have assumed that new bidders participate in each auction (see subsection 2.2). It would clearly be of interest to cover also situations in which each individual bidder remains active longer. When bidders remain active arbitrarily long, we should expect convergence to Nash equilibrium in our private values setup given that in each format M_k bidders could learn from their own past observations the distribution of other bidders's bids. However, if each individual bidder does not remain active for ever, then some outcome in between the Nash equilibrium and the feedback equilibrium (as defined in this paper) should be expected. More work is required to model such intermediate cases in a satisfactory way.

[36] Our framework is broad enough however for the purpose of showing the potential gain of disclosing partial feedback.

6 Conclusion

I believe the above abstract setup is useful to understand a number of applications. The idea that bidders would form their beliefs about others's bidding strategies by looking at the history of past bids should sound familiar to anyone who has considered buying or selling on ebay (under the item "completed listings" one has access to the history of previous bids in auctions of similar objects that took place within a month).[37] The feedback provided on ebay is partial in the sense that one has never access to the characteristics (such as gender, age etc) of the bidders and the same feedback appears whether or not a buyout option prevailed, as long as the option was not exerted (technically speaking, whether or not there is a buyout option should be interpreted as corresponding to different auction formats).

What this paper has emphasized is the use of feedback policy as a new instrument in mechanism design. My main interest was in understanding the effect of the feedback policy on efficiency and revenues.

The main contribution of this paper has been to show that there is a role for a strategic use of feedback disclosure in mechanism design. On the one hand, first-price auctions with bidder-anonymous feedback partition generate more revenues that second-price auctions when there are few competitors. Thus, providing coarse feedback in first-price auctions may be thought of as a new way of promoting more competition in asymmetric auctions that avoids the cost of reducing efficiency. On the other hand, the insight that with coarse feedback one can generate more revenues than in Myerson's optimal auction is suggestive that the lack of transparency that is often observed in promotion-like contests may be desirable for organizations. More work is required to understand more generally how much can be gained with the use of coarse feedback in mechanism design.

[37]See Reynolds and Wooders (2008) for an analysis of auctions with buyout options. I am grateful to John Wooders for introducing me with the various formats and feedback available on ebay.

Appendix

Proof of Claim 1: Consider a symmetric feedback equilibrium β of $\boldsymbol{A}$ (where $\beta(v)$ refers to the equilibrium bid of any bidder with valuation v).[38] By definition, bidder i plays a best-response to the distribution of bids of other bidders that has assigns density $\frac{\sum\limits_{j \in I} f_j(v)}{n}$ to the bid $\beta(v)$. But, this is the definition of a Bayes Nash equilibrium of $\Gamma^{ba}(\boldsymbol{A})$. The converse part is also immediate. **Q. E. D.**

Proof of Claim 2: Consider an equilibrium β of $\Gamma^{fa}(\boldsymbol{A})$. In $\Gamma^{fa}(\boldsymbol{A})$, 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 $\Gamma^{fa}(\boldsymbol{A})$, when the payment method is k, bidder i plays a best-response $\beta_i^k(v_i) \in \arg\max\limits_{b_i} u_i^k(v_i, b_i; \overline{\beta}_{-i}^k)$ 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 Proposition 1:

Step 1: Consider the first-price auction with anonymous bidder feedback partition. There exists a unique feedback equilibrium defined as follows: for $i = 1, 2$, $\beta_i(v) = \beta(v) = \frac{\int_c^v x\overline{f}(x)dx}{\overline{F}(v)}$ where $\overline{f}(x) = \frac{f_1(x)+f_2(x)}{2}$ and $\overline{F}(v) = \frac{F_1(v)+F_2(v)}{2}$. Bidders never quit ex-post and the outcome is always efficient, i.e. the bidder who values the good most gets the object.

Proof of step 1. Consider a feedback equilibrium $\beta_i(\cdot)$ for $i = 1, 2$. Standard incentive compatibility considerations imply that $\beta_i(\cdot)$ must be a non-decreasing function of the valuation (as otherwise a higher valuation type of bidder i would perceive to win the object with a probability strictly lower than a lower valuation type, which is ruled out by incentive compatibility). Thus, the bid functions $\beta_i(\cdot)$ must be continuous almost everywhere.

Suppose we have a non-symmetric equilibrium (that is not equivalent almost everywhere to a symmetric equilibrium). This implies that for a positive measure of v,

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

$\beta_1(v) \neq \beta_2(v)$ and both $\beta_1(v)$ and $\beta_2(v)$ are best-responses for a bidder with valuation v to the aggregate distribution of bids. There must then be a neighborhood of v within which a positive measure of v has this property. Yet, this implies that we can make another selection of the best-response correspondence that violates the monotonicity of $\beta_i(\cdot)$, thereby showing a contradiction.[39]

The rest of the argument follows from Claim 1 (see Section 4). Indeed, any symmetric feedback equilibrium must be a Nash Bayes equilibrium of the FPA with symmetric bidders and density $\overline{f}(v)$ and vice versa. Given the analysis of the FPA with symmetric bidders, we may conclude as desired. (The fact that bidders never exert their ex-post quitting rights follows from the rules of FPA. No bidder finds it optimal to bid above his valuation and thus when he wins a bidder finds it optimal to accept the deal.) **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 (the identity between the last two expressions can be obtained as a consequence of the allocation equivalence):

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

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

Step 2: $\overline{R} - R^{SPA} = \int_c^d \frac{1}{4}\left(F_1(v) - F_2(v)\right)^2 dv$

[39] Suppose $\beta_1(v) < \beta_2(v)$. By continuity $\beta_1(v+\varepsilon) < \beta_2(v)$ and $\beta_2(v+\varepsilon) > \beta_1(v)$. The definition of a feedback equilibrium implies that $b_1(v) = \beta_2(v)$ and $b_1(v+\varepsilon) = \beta_1(v+\varepsilon)$ with all other bids unchanged should also be part of an equilibrium. But, such bids would violate the incentive compatibility conditions and as a result cannot maximize (over bids) the corresponding expected payoffs of bidder 1 with valuations v and $v+\varepsilon$.

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

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

where the last equality is obtained by integration by parts (noting that $F_1(v) - F_2(v) = 0$ for $v = c$ and d). **Q.E.D.**

Step 3: $R - \overline{R} = \int_c^d \frac{1}{4} \frac{d\beta(v)}{dv} \left(F_1(v) - F_2(v) \right)^2 dv$

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

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

where the last equality is obtained by integration by parts (noting that $F_1(v) - F_2(v) = 0$ for $v = c$ and d). **Q. E. D.**

Observe that $\frac{d\beta(v)}{dv} > 0$ for all v. Hence, Proposition 1 follows from steps 1, 2, 3. **Q. E. D.**

Proof of Proposition 2:

We start with the following observation:

Step 1: Myerson's optimal auction can be implemented while satisfying the ex-post quitting rights of the bidders in a direct truthful mechanism in which reporting the truth is a weakly dominant strategy for every bidder.

Proof. This is easily shown by simple adaptation of the second-price auction to the optimal auction of Myerson. In the asymmetric regular case, the functions $c_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ are increasing in v_i, and the optimal auction 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 required to pay $\max_{j \neq i} \left[c_i^{-1}\left(c_j(v_j)\right), c_i^{-1}(0)\right)]$. It is easily checked that this payment is always less than v_{i^*} by the monotonicity of $c_i(\cdot)$. A similar construction can be achieved in the general non- necessarily regular in which intervals of valuations are treated alike.[40] **Q. E. D.**

The rest of the argument goes as follows. Consider a monotonic bijection ψ from $[c, d]$ into itself, and let M^{ψ} be the mechanism obtained from the mechanism M^D identified in step 1 as follows: in M^{ψ}, every bidder i submits a bid b_i and mechanism M^D is applied to the profile of announcements $(\psi(b_i))_{i=1}^n$. Clearly, M^{ψ} falls in the class of admissible mechanisms and reporting $\psi^{-1}(v_i)$ for bidder i with valuation v_i is a weakly dominant strategy. Besides, M^{ψ} achieves Myerson's optimal auction revenues and no bidder is willing to exercise his ex-post quitting rights in M^{ψ}.

Consider now the following auction design. Format M^{ψ} is used with probability $1 - \varepsilon$ and the first-price auction FPA with 0 reserve price is used with probability ε. Besides, bidders get only to know the aggregate distribution of bids of all bidders across both formats. That is, we consider the bidder-anonymous and format-anonymous feedback partition in which for all i, $\bigcup_{(j,k)} \{(j, k)\}$ forms the unique feedback class of P_i. We will show that for a suitable choice of ε and ψ this auction design generates strictly more revenues that Myerson's optimal auction. First, we observe that the revenue generated in this auction design can be written as $(1 - \varepsilon)R^{\psi} + \varepsilon R^*$ where R^{ψ} is the expected revenue generated in this auction design when M^{ψ} prevails and R^* is the corresponding expected revenue when FPA prevails. It is clear that R^{ψ} is equal to Myerson' s optimal auction revenue R^M, since the behaviors in M^{ψ} are unaffected by the rest of the auction design given that bidders have (weakly) dominant strategies in M^{ψ}. Thus, it suffices to show that $R^* > R^M$ for suitable choices of ε and ψ.

[40] One can easily perturb the format so as to make incentives strict in all cases (even in the non-regular case).

To this end, let ψ be defined such that[41]

$$\prod_{i=1}^{n} F_i(\psi(b)) = \left(\frac{b-\widehat{c}}{d-\widehat{c}}\right)^{\frac{n}{n-1}m}$$

for some m that will be chosen sufficiently large later on.

In the limit case in which $\varepsilon = 0$, the (perceived) optimal bid in FPA for a bidder with valuation v would be $\arg\max_b (v-b)\left(\frac{b-\widehat{c}}{d-\widehat{c}}\right)^m$ as $\left(\frac{b-\widehat{c}}{d-\widehat{c}}\right)^m$ would represent the perceived probability that all other bidders' bids are below b. This expression is maximized at b^{opt} such that

$$b^{opt} - \widehat{c} = \frac{m}{m+1}(v-\widehat{c})$$

Let b^* be such that $b^* - \widehat{c} = \frac{m}{m+2}(v-\widehat{c})$ and consider $\varepsilon > 0$. A bidder with valuation v will perceive to get at most

$$(1-\varepsilon)(v-b^*)\left(\frac{b^*-\widehat{c}}{d-\widehat{c}}\right)^m + \varepsilon(v-\widehat{c}) \tag{1}$$

by bidding $b < b^*$.

By bidding b^{opt}, a bidder with valuation v will perceive to get at least:

$$(1-\varepsilon)(v-b^{opt})\left(\frac{b^{opt}-\widehat{c}}{d-\widehat{c}}\right)^m \tag{2}$$

Hence, whenever (2) is larger than (1) we can be sure that a bidder with valuation v bids no less than $\widehat{c} + \frac{m}{m+2}(v-\widehat{c})$. The difference between (2) and (1) writes

$$\Delta(v) = (1-\varepsilon)\frac{(v-\widehat{c})^{m+1}}{(d-\widehat{c})^m}\left[\frac{1}{m+1}\left(\frac{m}{m+1}\right)^m - \frac{2}{m+2}\left(\frac{m}{m+2}\right)^m\right] - \varepsilon(v-\widehat{c})$$

Given that $\frac{1}{m+1}\left(\frac{m}{m+1}\right)^m - \frac{2}{m+2}\left(\frac{m}{m+2}\right)^m > 0$, this allows us to obtain that:

Step 2: $\forall \underline{v} > \widehat{c}$, $\forall m$, $\exists \overline{\varepsilon} > 0$ such that $\forall \varepsilon < \overline{\varepsilon}$, $\forall v > \underline{v}$, $\Delta(v) > 0$.

From step 2 and the above considerations, we infer that for all $v > \underline{v}$, $b^{FPA}(v) - \widehat{c} > \frac{m}{m+2}(v-\widehat{c})$ in the above auction design as defined by ψ and $\varepsilon < \overline{\varepsilon}$. The corresponding

[41] It is clear that such ψ exists and of course satisfies $\psi(c) = c$ and $\psi(d) = d$.

value of R^* converges to the full information revenue R^F as m converges to infinity and $\underline{v}$ converges to $\widehat{c}$. It follows that one can find m large enough, $\underline{v}$ close enough to $\widehat{c}$ and $\varepsilon > 0$ so that $R^* > R^M$. This completes the proof of the proposition. **Q. E. D.**

Proof of Proposition 3:

Consider an auction design assumed to deliver an expected revenue that is ε-close to R^F. We will show that this is not possible for ε small enough.

To simplify the notation, we consider the case of two symmetric bidders $i = 1, 2$ and we allow only for auction designs with format-anonymous feedback partitions. The argument easily generalizes to the n asymmetric bidder case with arbitrary feedback partitions (by restricting attention to those formats that are pooled together into one feedback class of say bidder i).

Observe that in all formats, whenever $v_j < 0$, we must have $\beta_j^k(v_j) = 0$ given the payment rules of the auction. We further let $\gamma = \Pr(v_j < 0, j \neq i) > 0$, and we let $\Delta^k(v)$ denote the expected revenue loss incurred by the designer in format M_k when bidder i has valuation v as compared with the full information case.

Let m be large enough and let $e < \frac{1}{2}\Pr(\frac{d}{3} < v_j < \frac{d}{2})$. Define $\underline{d} > \frac{d}{3}$ such that $e = \Pr(\underline{d} < v_i < \frac{d}{2})$. Finally, let $f = \min(\Pr(v_i > \frac{3d}{4}), e)$, and $\eta = \frac{m\varepsilon\gamma}{f}$. We define

$$\Gamma = \left\{ k \text{ s.t. } \exists v > \frac{3d}{4} \text{ and } \exists v' \in (\underline{d}, \frac{d}{2}),\ \Delta^k(v) < \eta \text{ and } \Delta^k(v') < \eta \right\}$$

Given that the auction design delivers an expected revenue that is ε close to R^F, it is readily verified that $\sum_{k \in \Gamma} \lambda_k \geq 1 - \frac{1}{m}$.

Perceived equilibrium payoff:

We note that for $v > \underline{d}$, if $\Delta^k(v) < \eta$, then bidder with valuation v should in format M_k win whenever $b_j = 0$ (which happens with probability γ) and pay at least $v - \frac{\eta}{\gamma}$. By monotonicity of the payment rule, this implies that in M_k bidder i with valuation v perceives in equilibrium to get at most

$$\frac{\eta}{\gamma} = \frac{m\varepsilon}{f}$$

(the payment when i wins, bids b_i and $b_j > 0$ must be at least as large as when i wins, bids b_i and $b_j = 0$).

We also note that $\Delta^k(v) < \eta$ implies that in M_k bidder i with valuation v should win against some $v_j \in (\underline{d}, \frac{d}{2})$ with probability at least $1 - \frac{\eta}{fv} = 1 - \frac{m\varepsilon\gamma}{f^2 v}$. We will choose ε small enough so that this probability is no smaller than $\frac{1}{2}$.

Perceived equilibrium from downward deviation:

One can rank the various $k \in \Gamma$ by decreasing order of $\beta^k(\frac{d}{3})$, and let $\overline{r}$ denote the maximum r such that the sum of λ_k over the first $r-1$ formats in Λ is strictly below $\frac{1}{2}\sum_{k\in\Gamma}\lambda_k$. We denote by Λ^{sup} the formats in Λ which correspond to the first $\overline{r}$ formats in this induced order.

Consider $k \in \Lambda^{\text{sup}}$ and let $v' \in (\underline{d}, \frac{d}{2})$, $\Delta^k(v') < \eta$. Consider any $v > \frac{3d}{4}$ and let v submits a bid $b_i = \beta_i^k(v')$. Bidder i with valuation v in format M_k must perceive to be winning with probability at least $\frac{1}{4}(1-m)\Pr(v_j < \frac{d}{3})$,[42] and he must be paying at most $\frac{d}{2}$ whenever he wins.[43] Given that $v > \frac{3d}{4}$ (and thus $\frac{3d}{4} - \frac{d}{2} = \frac{d}{4}$), overall such a deviation makes bidder i feel he can get at least $\frac{1}{2}(1-m)\frac{d}{4}\Pr(v_j < \frac{d}{3})$ in M_k.

Given that ε can be chosen so that $\frac{m\varepsilon}{f} < \frac{1}{2}(1-m)\frac{d}{4}\Pr(v_j < \frac{d}{3})$ we get a contradiction to the definition of a feedback equilibrium (since a bidder should obviously feel that his perceived payoff obtained by following his equilibrium strategy is no smaller than his perceived payoff obtained by following any other strategy). **Q. E. D.**

Proof of Proposition 4:

We consider the following formats M_1 and M_2 both used with probability $\frac{1}{2}$. In both formats, the good is never allocated whatever the bids, $\varphi_i^k(b) = 0$ for all i, b and $k = 1, 2$. In format, M_1, bidder 1 wins ε if $b_1 = 1$ and 0 otherwise. In format M_2, bidder 1 wins ε if $b_1 = 2$ and 0 otherwise. In format M_2, bidder 2 pays $\frac{A}{2} > 0$ if $b_1 = 2$ and $b_2 = 1$ and

[42] This is because in the format anonymous feedback partition, all the bids $\beta_j^{k'}(v_j)$ with $v_j < \frac{d}{3}$ and $k' \in \Lambda \setminus \Lambda^{\text{sup}} \cup \{\overline{r}\}$ must be below $\beta_i^k(v')$ and by construction $\Lambda \setminus \Lambda^{\text{sup}} \cup \{\overline{r}\}$ has a probability at least $\frac{1}{2}(1-m)$.

Moreover in M_k, i with valuation v should win against some $v' \in (\underline{d}, \frac{d}{2})$ with a probability at least $\frac{1}{2}$ (see above), and thus by the monotonicty of $\varphi_i^k(b)$ with respect to b_j he should also win against all bids which are below $\beta_k^i(\underline{d})$ with a probability at least $\frac{1}{2}$.

[43] This is because he is mimicking type v' who never pays more than v' when he wins.

receives A if $b_1 = 1$ and $b_2 = 1$, and receives nothing otherwise, i.e.

$$\tau_2^2(b) = \begin{cases} \frac{A}{2} \text{ if } (b_1, b_2) = (2,1) \\ -A \text{ if } (b_1, b_2) = (1,1) \\ 0 \text{ if } (b_1, b_2) \neq (2,1), (1,1) \end{cases}$$

and the feedback partition is the anonymous-format feedback partition.

Clearly, in this auction design, bidder 1 will bid $b_1 = 1$ in M_1 and $b_1 = 2$ in M_2. Given that $\lambda_1 = \lambda_2 = \frac{1}{2}$, and the format-anonymous feedback partition is being used, bidder 2 will believe that in M_2, bidder 1 bids $b_1 = 1$ or 2 each with probability $\frac{1}{2}$. Based on this belief, bidder 2 will find it optimal to bid $b_2 = 1$ in M_2 (because $\frac{1}{2}(A - \frac{A}{2}) > 0$).

In such a feedback equilibrium, the designer gets a revenue equal to $-\varepsilon$ in M_1 and $\frac{A}{2} - \varepsilon$ in M_2 so an overall expected revenue of $\frac{A}{4} - \varepsilon$. Since A can be chosen arbitrarily large, we get the desired result. **Q. E. D.**

References

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

[2] Cabrales, A. (1999): 'Adaptative Dynamics and the Implementation Problem with Complete Information,' *Journal of Economic Theory* **86**, 159-184.

[3] Cabrales, A. and R. Serrano (2007): 'Implementation in Adaptive Better Response Dynamics,' mimeo.

[4] Compte, O. and P. Jehiel (2007): 'On Quitting Rights in Mechanism Design,' *American Economic Review*, **97**, 137-141.

[5] Crawford, V. T. Kugler, Z. Neeman and A. Pauzner (2008): 'Behaviorally Optimal Auction Design: Examples and Observations,' forth. *Journal of European Economic Association.*

[6] Eliaz, K. (2002): 'Fault Tolerant Implementation,' *Review of Economic Studies*, **69**, 589-610.

[7] Eliaz, K. and R. Spiegler (2007): 'A Mechanism Design Approach to Speculative Trade,' *Econometrica* **75**, 875-884.

[8] Esponda, I. (2008a): 'Information Feedback in First Price Auctions', *Rand Journal of Economics*, **39**, 491-508.

[9] Esponda, I. (2008b): 'Behavioral Equilibrium in Economies with Adverse Selection,' *American Economic Review*, **98**, 1269-91

[10] Ettinger, D. and P. Jehiel (2009): 'A Theory of Deception,' forthcoming *American Economic Journal: Microeconomics*

[11] Eyster, E. and M. Rabin (2005): 'Cursed Equilibrium,' *Econometrica*, **73**, 1623-1672.

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

[13] Hafalir I. and V. Krishna (2007): 'Asymmetric Auctions with Resale,' *American Economic Review*, **98**, 242-258.

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

[15] Huck, S., P. Jehiel, and T. Rutter (2008): 'Learning Spillovers and Analogy-based Expectations: A Multi-Game Experiment,' mimeo.

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

[17] Jehiel, P. and F. Koessler (2008): 'Revisiting Games of Incomplete Information with Analogy-based Expectations', *Games and Economic Behavior*.

[18] Lehrer, E. (2007): 'Equilibrium with Partially Specified Strategies,' mimeo.

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

[20] Matsushima, H. (2008): 'Detail-free Mechanism Design in Twice Iterative Dominance: Large Economies,' *Journal of Economic Theory* **141**, 134-151.

[21] McAfee, P. and J. McMillan (1987): 'Auctions with a Stochastic Number of Bidders,' *Journal of Economic Theory* **43**, 1-19.

[22] Milgrom, P. (2004): *Putting Auction Theory to Work,* MIT Press.

[23] Milgrom, P. and R. Weber (1982): 'A Theory of Auctions and Competitive Bidding,' *Econometrica*, **50**, 1089-1122.

[24] Morris, S. (1994): 'Trade with Heterogeneous Prior Beliefs and Asymmetric Information,' *Econometrica*, **62**, 1327-1347.

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

[26] Reynolds S. and J. Wooders (2008): 'Auctions with a Buy Price' forth. *Economic Theory.*

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