Manipulative Auction Design*

Philippe Jehiel†

10th January 2008

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 observed in previous auctions to report to bidders. 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.

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, A. Wolinsky, and seminar participants at Roy seminar (Paris), Northwestern University, and the conference in honor of Ehud Kalai (Jerusalem, 2007) 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 in an asymmetric private values setting involving each time new bidders of observable characteristics $i = 1, 2, ...n$. That is, each bidder when choosing his bid knows his valuation but not that of the other bidders, 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 would be able to play a best-response to the actual distribution of bids of the other bidders - 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 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.[4]

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. Specifically, our framework 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

[1] A bidder with characteristic i need not even know the distribution from which his own valuation is drawn.

[2] This looks like the norm rather than the exception in many auction setups. For example, in treasury auctions, the bids whether in discriminatory or uniform auctions are kept anonymous.

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

[4] This is assuming that bidders reason as if there were no asymmetries between bidders, which given the feedback they receive is the simplest conjecture they can entertain.

distribution of bids across the various auction formats is being disclosed at the learning stage.[5] The equilibrium obtained (which stands for the limiting outcome of a learning process with corresponding feedback disclosure) is called a feedback equilibrium,[6] and it requires that bidders play a best-response to the aggregate distribution of bids, as given by the feedback they receive.

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). It can also be viewed as a selection of self-confirming equilibrium (with an appropriate signal structure) in which the conjecture considered by the bidders is the simplest theory consistent with the feedback they receive. 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 a multi-game experimental setting in which the feedback about subjects' past behaviors is not cleanly separated between games.

Importantly, 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 choice. 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.[7] Other than that, the designer is free to choose the feedback (partition) she likes best and she is free to use whatever auction format in which larger bids translate into greater chances of winning the object. In some parts of the paper, we also require the bidders to approve the terms of the deal after the outcome of the auction is known. That is, we require ex-post participation constraints.

The first question we ask is as follows. Suppose the main objective of the designer is welfare maximization whereas the auxiliary goal is revenue. Can the designer do better

[5] In the latter case, one should have in mind that bidders are not necessarily informed of the payment rule of other bidders.

[6] 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.

[7] This can be viewed as being in the interest of the designer as otherwise bidders would presumably discount the information provided to them.

than using a second-price auction (or equivalently an ascending-price auction)?

In the classic rationality setup, 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 our manipulative auction design framework, we show that the designer can sometimes do better, thereby illustrating a failure of the revenue equivalence theorem in a setup with partial information feedback. 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 characteristic 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 with asymmetric bidders (see Maskin and Riley (2000)).[8]

The second question we address is as follows. Suppose that the designer is solely interested in revenue maximization. Can the designer generate more revenues than in the classic optimal auction (Myerson (1981) or Riley and Samuelson (1981))?

We show that this is always so whether or not parties can veto the transaction after the outcome of the auction is known. In other words, 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.[9]

[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 thus there is no room for resale.

[9] 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

More precisely, when bidders do not enjoy quitting rights, we show that the designer can always design an auction mechanism and a feedback device so that she makes arbitrarily large revenues. And when bidders have ex-post quitting rights, the designer's best revenue lies strictly above Myerson's optimal auction revenue and strictly below the full information optimal revenue that would be obtained if the designer knew the valuations of all bidders. In the latter case, the manipulation exploited by the designer has the effect of reducing the informational rent left to the bidders.

It may be mentioned that while considering the information feedback as a design choice 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.[10]

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 Sections 5 and 6. Section 7 contains further results in particular about the complete information case and the incentive for the designer to use shill bidders. Section 8 suggests some further applications of the approach and some avenues for future research. Missing proofs can be found in the Appendix.

---

knowledge, she could then use the feedback she likes best, as analyzed in this paper.

[10]Our results are very different in nature from those results showing that the designer is better off revealing as much information as she can when the seller's private information and bidders' valuations are affiliated (see Milgrom and Weber (1982)). In our setup, providing partial feedback about past bids leads bidders to play a non-Nash equilibrium, which has no counterpart in Milgrom and Weber's analysis. Yet at a more general level, the idea of choosing a design that reduces the informational rent left to bidders is common to the literature on the linkage principle and the manipulative auction design setup considered here.

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 $[c, d]$ and (continuous) density $f_i(\cdot)$ where we assume that $f_i(v) > 0$ for all $v \in [c, d]$. We assume that $d > 0$, and in some parts of the paper we assume that $c < 0$ where 0 is the seller's valuation. Bidders do not *a priori* know the densities $f_i(\cdot)$ (not even their own $f_i(\cdot)$). 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 $pv - t$.

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. Each auction format M_k takes the following form:

- Bidders $i = 1, ...n$ simultaneously submit a bid $b_i \in \left[0, \overline{b}\right]$.
- 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 $b_i = 0$.

Throughout the paper, we assume that for each i and k, $\varphi_i^k(b)$ is a non-decreasing function of b_i and a non-increasing function of b_j, $j \neq i$. That is, in line with the auction interpretation, the higher the bid of bidder i (both in absolute and relative terms) the greater the chance that i wins the object.

When auction format M_k is selected, 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 $\widetilde{b}_{-i}$ in M_k, his perceived expected utility in M_k is:

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

Observe that to perform this calculation, bidder i need not know the functions $\varphi_j^k(b)$, $\tau_j^k(b)$ that apply to bidders $j \neq i$. This is the benchmark knowledge assumption one should have in mind when interpreting the feedback equilibrium (to be introduced later on).

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 player i's bid in format M_k when i's valuation is v_i.[11]

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 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. 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 $\varphi_i^k(b)$ and $\tau_i^k(b)$) the aggregate distribution of bids in every α_i (see the learning interpretation in Section 3). 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:

[11] 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 $\overline{\beta}_{-i}^k = (\overline{\beta}_j^k)_{j \neq i}$, $\overline{\beta}_j^k$ *is the aggregate distribution of bids 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.*[12]

Remark. It should be mentioned that, in the current formulation, the feedback received by bidders is about the distribution of individual bids and not about the distribution of profiles of bids.[13] 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.

We will consider various objectives for the designer. The first objective will be a lexicographic criterion with welfare ranked first and revenues ranked second. The second objective that we will consider is revenues. In all cases, we will assume that the designer is risk neutral (and as already mentioned that she has no intrinsic value for the object).

The participation of bidders to the auction should be voluntary. Our assumption that a bidder submitting a 0 bid makes no payment ensures that participation constraints are always (perceived to be) satisfied at the interim stage before bidders see the outcome of the mechanism. We will also consider the scenario in which participation should also be approved ex post after bidders see the outcome of the mechanism, thereby providing bidders with ex post quitting rights (see Compte and Jehiel (2007) for elaborations on quitting rights in mechanism design).

When we consider ex-post quitting rights, we restrict our analysis to auction formats in which payments are only made by the winner (anyway a loser would refuse to make

[12]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 $(j',k') \in \alpha_i(j,k)$. Note that the densities f_j are not assumed to be known to the bidders. The aggregate distribution of bids is got accessible to bidders through learning, see the next section for further elaborations on this interpretation.

[13]Accordingly, every bidder i treats every bidder j's distribution of bids, $j \neq i$, as being independent of each other.

a positive payment) and, in line with the auction interpretation, we further assume that the payment made upon winning is a non-decreasing (and non-negative) function of the bids submitted by bidders.[14]

We note that, in the standard rationality paradigm, the optimal revenue-maximizing auction can always be implemented while satisfying bidders' ex-post participation constraints (think, for example, of the second price auction with optimally set reserve price in the symmetric regular case). One of our main insights will be to show that by an appropriate choice of manipulative auction design, the designer can achieve strictly larger revenues while ensuring the ex post participation constraints. The best revenues so obtained then lies in between the complete information optimal revenues and the incomplete information optimal revenues as analyzed in Myerson (1981) (see Section 7).

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.[15] Bidders are assumed not to be aware of $f_i(\cdot)$ (even though they can observe the characteristic of other bidders as the designer does). To simplify the exposition of the interpretation, we assume that the designer has got sufficient experience to know $f_i(\cdot)$.[16]

In such a context, the designer can change her auction format from one period to another while ensuring that the frequency with which format M_k is used 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

[14] That is, over the range $\varphi_i^k(b) > 0$, we assume that $\frac{\tau_i^k(b)}{\varphi_i^k(b)}$ is increasing in b_i and b_{-i}.

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

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

aggregate distribution of past bids (the empirical frequencies) in every α_i.[17] If she does so every period, and if behaviors stabilize, it must be to a feedback equilibrium provided bidders consider the simplest theory that is consistent with the feedback they receive.[18] 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 $\boldsymbol{A}$.[19]

The approach 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, and in line with the literature on bounded rationality, bidders are assumed to consider the simplest theory consistent with the feedback they receive.[20]

About the choice of feedback partitions, 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, thereby imposing that 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 (ABEE) introduced in Jehiel (2005), further developed in Jehiel and Koessler (2007) and Ettinger and Jehiel (2007). The feedback

[17] 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.

[18] 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.

[19] 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.

[20] While it would be possible to interpret bidders' beliefs through the Bayesian machinery by relying on well chosen subjective priors, we consider the approach in terms of simplicity as preferable to the subjective prior approach. This is because in the subjective prior paradigm, it is not clear where the subjective prior would come from, especially if one has in mind that the prior should, in principle, be determined by what is known about similar auctions.

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).

We note that Huck et al. (2007) provides experimental support to ABEE. In the experiment of Huck et al., 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.[21] 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). It should be mentioned that in Huck et al.'s experiment, subjects were not informed of their opponent's payoff structure, and this is why we prefer interpreting our results assuming that bidder i in format M_k is not informed of the functions $\varphi_j^k(\cdot)$, $\tau_j^k(\cdot)$ for $j \neq i$.

The only other paper we are aware of that considers information feedback in auctions as a design choice is by Esponda (2007). 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 Nash equilibrium.

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

[21] 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.

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.[22] 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). We also note 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.[23] 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) *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

[22] 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.

[23] 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.

distribution of bids with no mention of the characteristics of the bidders who generated the various bids.

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 $i \in I$, $\boldsymbol{P}_i = \left\{ \bigcup_{k \in K} \{(j,k)\} \right\}_{j \in I}$. 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.

4 Preliminaries

We make a few preliminary observations. 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. 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 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, and thus no matter how the auction design is further specified.

Third, one of the auction designs that we will study 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 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.[24]

Claim 2: 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 $\Gamma^{fa}(\boldsymbol{A})$.

5 Efficiency and revenues

In a number of applications, the designer may be interested both in efficiency and 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 $v_i = c$ 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.

[24] 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.

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 device).

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. 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* $\overline{F}(v) = \int_c^v \overline{f}(x)dx$.[25]

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 $\overline{F}(v) = \frac{F_1(v)+F_2(v)}{2}$. It follows that the outcome is efficient in our auction design.

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,

[25] $\beta(\cdot)$ 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.

$\beta_1(v) \neq \beta_2(v)$ 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.[26]

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

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

[26]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 $v+\varepsilon$.

written as

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

$$= \int_c^d -\frac{v}{2}\left(f_1(v)-f_2(v)\right)\left(F_1(v)-F_2(v)\right)dv$$

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

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

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

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

$$\int_c^d \beta(v)\left[f_1(v)F_2(v)+f_2(v)F_1(v)-2\frac{f_1(v)+f_2(v)}{2}\cdot\frac{F_1(v)+F_2(v)}{2}\right]dv$$

$$= \int_c^d -\frac{1}{2}\beta(v)\left(f_1(v)-f_2(v)\right)\left(F_1(v)-F_2(v)\right)dv$$

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

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

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

What is the intuition for the above result ? Not providing feedback about the characteristic i of the bidders who submitted past bids leads the bidders in the current auction 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 enhances revenues because it makes the strong bidder feels the weak bidder is less weak than he really is.[27]

It should be noted that the first-price auction with anonymous-bidder feedback partition need not always generate more revenues than the second-price auction 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 d and a third bidder whose 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). On the other hand, if there is only one bidder whose distribution of valuations is concentrated around d while the 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). 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 neither at the interim stage nor at the ex post stage when the outcome of the auction is known. This is because bidders bid below their valuation (bidding above one's own valuation is dominated) and thus they cannot make losses. Thus, even if bidders have a right to quit the mechanism ex-post after the outcome of the auction is known, no bidder would use this option. The manipulation exploited by the designer has here the effect of reducing the informational rent left to the bidders while preserving the veto rights of the bidders. The most extreme revenue gain provided by the above exploitation is when the distribution of valuations of one bidder is concentrated around d whereas the distribution of valuations

[27]This is essentially what step 2 formalizes. Step 3 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.

of the other bidder is concentrated around c in which case the first-price auction with anonymous feedback partition provides a revenue gain of $\frac{d-c}{2}$, which should be compared to the revenue c of the second-price auction.[28] Clearly, as d gets large relative to c, the revenue gain can be quite substantial in such asymmetric setups.

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 $\beta^*(v) = \frac{\int_c^v x f^*(x)dx}{F^*(v)}$ where $F^*(v) = (F_1(v)F_2(v))^{\frac{1}{2}}$ and $f^*(v) = \frac{dF^*(v)}{dv}$. Such bidding strategies would always generate higher revenues than in the second price auction, as in Proposition 1.[29]

6 Revenues

We now assume that the designer seeks to maximize revenues. We first consider the case in which the participation constraints are only required at the interim stage and then move on to consider the case in which bidders can decide to withdraw from the auction after learning the outcome of the auction.

6.1 Unrestricted domain of mechanisms

If participation constraints are only required at the interim stage before bidders know the outcome of the auction, we observe that the designer's revenue can be arbitrary large.

[28] 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.

[29] 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

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

Proposition 2 *Suppose there are at least two bidders. 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 fairly simple. 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 $i \neq 1$. He 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.[30] Of course, here since the designer is assumed to know the correct distributions of bids, we make 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 about adopting the simplest conjecture consistent with the feedback he receives, especially if the stakes of the bet are very large (bidder 1 may well realize that basing his estimates on the simplest conjecture consistent with the observed feedback may be harshly exploited by the designer even if bidder 1 need not understand how). Taking such considerations into account would require amending the solution concept of feedback equilibrium, which should be the subject of future research. Somehow the next subsection bypasses this difficulty by considering scenarios in which the stakes are necessarily limited due to the ex-post quitting rights of the bidders. In such cases as in the scenario studied in Section 5, the feedback equilibrium seems quite reasonable (as bidders have no clue on how they could avoid being exploited - for

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

example, not participating in the auction would always be a bad idea as bidders always make non-negative profits ex-post).[31] Alternatively, one may think of the ex-post quitting rights scenario considered below as a regulatory constraint imposed on designers to better protect bidders from manipulation.

6.2 Mechanisms with ex-post quitting rights

We now assume that bidders must approve the transaction after the outcome of the auction is known. As mentioned in Section 2, we limit ourselves to auction formats in which only the winner of the auction makes a payment, and we further assume that the payment made by the winner of the auction is a non-decreasing function of the bids submitted by the bidders.[32] We refer to such settings as auction designs with ex-post quitting rights.

As already mentioned, the optimal auction of Myerson (1981) can always be implemented within the class of auction designs with ex-post quitting rights. As we now show, the designer can always extract strictly larger revenues.

Proposition 3 *The largest revenue that the designer can achieve in a manipulative auction design with ex-post quitting rights is strictly larger than the revenue generated in Myerson's optimal auction (denoted by R^M).*

The intuition for Proposition 3 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

[31] Recently, Lehrer (2007) has proposed another 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 2. 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. We find the latter conclusion unrealistic.

[32] Clearly a loser would object to make a positive payment. The fact that the seller is not allowed to make a positive payment to the loser can be viewed as resulting from the ex-post quitting right of the seller vis a vis the losing bidder. Alternatively, it can be thought of as a regulatory constraint.

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 will not necessarily deliver revenues higher than R^M.[33] 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 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 3 follows.

Remark. By inspecting the proof of Proposition 3, we 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 3 would hold under the stronger full implementation requirement (provided one restricts ourselves to equilibria not employing weakly dominated strategies).

Proposition 3 shows 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 we now show, in manipulative auction designs with ex-post quitting rights the designer's best revenue lies strictly below R^F whenever bidders' valuations can take negative values.

Proposition 4 *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 whenever $\Pr(v_i < 0) > 0$ for all i.*

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

In the proof of Proposition 3 we have mixed some (Myerson-optimal) mechanism MD implementable in dominant strategy with a little bit of FPA, and we have observed that the revenue obtained in FPA could get 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 4 establishes is that within the class of mechanisms under study one can never reach R^F whatever the manipulation. To get an intuition of 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 average 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).[34] 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.[35] 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.[36] Our restriction on mechanisms (i.e., not allowing

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

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

[36] To see this, consider a symmetric two-bidder scenario in which bidders' valuations are identically

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.[37]

7 Further results

7.1 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 designer can extract a revenue equal to $R^F = \max_i(v_i, 0)$ in the classic rationality setup.

Consider now our manipulative auction design setup. If no ex-post quitting right is assumed then arbitrary large revenues can be obtained using mechanisms of the type considered in Proposition 2. Even though the setup is deterministic, the designer can introduce some randomness by considering different auction formats. By providing partial feedback she may next mislead bidders into thinking that high stake bets are valuable whereas in reality they are not as in the proof of Proposition 2 (see the appendix).

If however ex-post quitting rights are assumed, then no manipulation can allow the designer to extract more than R^F given that a bidder will never accept to pay more than

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.

[37] Alternatively, one may postulate that the use of promise of positive payments to losers (as well as the use of non-monotonic payment rules from winners) would raise the suspicion of bidders who would be more reluctant to participate in such auctions.

his valuation if he wins the object (and he will never accept to pay anything if he does not win the object). In mechanisms with ex-post quitting rights, manipulation is pointless when there is no private information.

7.2 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 does, as we now illustrate.

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, we observe that the designer can get a revenue close to the full information revenue R^F even if bidders have ex-post quitting rights (and the restrictions on mechanisms as considered in subsection 6.2 apply).

Proposition 5 *Suppose the cost of hiring shill bidders is zero and that bidders have ex-post quitting rights. 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 3. 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 with reserve price 0 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 with reserve price 0. 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 3).

8 Discussion 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' distributions of bids.

If the designer is unconstrained in her use of mechanisms and manipulations, she can extract arbitrary large revenues by misleading bidders about the distribution of others' bids (through the use of coarse feedback partitions) and offering large stake bets to these bidders based on such wrong beliefs.

Such an insight suggests that there is a role for regulatory interventions that would provide some sort of protection to bidders. We have suggested the idea of providing bidders with ex-post quitting rights, that is, a right to undo the transaction after it has been made. In this case, the designer cannot extract more than the full information revenues. As it turns out, the designer cannot approach this full information revenue upper bound unless she can use shill bidders freely. In order to better protect bidders, our analysis suggests that the use of shill bidders should be made less attractive either through an explicit ban or by forcing the designer to treat all bidders alike (which would make the use of shill bidders more costly).

Another possible regulatory idea would be to force the designer to make public the rules of the auction that applies to all bidders.[38] Some of the manipulations used in the paper would sound less plausible in such a scenario as bidders could make further inferences based on other bidders' incentives in the various formats.[39] We note however

[38] We note that the rules used by Google for the ad position auctions are far from being transparent as far as how bidders are ranked as a function of their bids because the ranking eventually depends on the expected number of clicks, which Google does not announce publicly.

[39] For example, if it is known that the designer uses FPA and SPA both with frequency 0.5, bidders can identify the distribution of valuations assuming bidders play a Nash equilibrium in each format, even

that the manipulation highlighted in Proposition 1 seems more robust to such regulatory constraints.[40]

There are several interesting directions for future research that we now wish to mention. 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).[41] 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. We note that our analysis of manipulative auction design could also be extended to study the case of stochastic number of bidders and to the case of correlated distributions of valuations or affiliated signals over valuations. In addition, while in our setting only the designer can provide feedback about the distribution of past bids, one could imagine that bidders themselves try to search for additional feedback to better combat manipulation. Finally, our analysis and main results even though formulated for the auction setup can be applied to other contexts. 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 winners 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.

---

if bidders are only told the aggregate distribution of bids over the two formats. From such distributions of valuations, they can next infer the distribution of bids in FPA and play the corresponding Nash equilibrium.

[40] Besides, in the context of Proposition 1, since the designer prefers the FPA with anonymous-bidder feedback partition to the SPA whatever the distributions of valuations, it seems hard for bidders to make further inference from the choice of the auction design.

[41] For example, in many electricity auctions over the world, only the clearing price is revealed initially while the aggregate distribution of bids is revealed only months later if at all (due to bidders' reluctance to have this information revealed).

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).[42] 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 $\Gamma^{ba}(\boldsymbol{A})$. 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 $\beta_i^k(v_i) \in \arg\max_{b_i} u_i^k(v_i, b_i; \overline{\beta}_{-i}^k)$ 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 2:

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 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

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

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.**

Proof of Proposition 3:

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 $\max_{j \neq i} \left[c_i^{-1}(c_j(v_j)), c_i^{-1}(0)\right]$. 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.[43] **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,

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

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, $\bigcup\limits_{(j,k)} \{(j,k)\}$ 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.

To this end, let ψ be defined such that[44]

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

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

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

by bidding $b < b^*$.

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

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 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 4:

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 $\frac{1}{4}(1-m)\Pr(v_j < \frac{d}{3})$,[45] and he must be paying at most

[45] 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

$\frac{d}{2}$ whenever he wins.[46] 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.**

which are below $\beta_k^i(\underline{d})$ with a probability at least $\frac{1}{2}$.

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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