

SHOULD AUCTIONS BE TRANSPARENT?

By

Dirk Bergemann and Johannes Hörner

August 2010
Revised August 2017

COWLES FOUNDATION DISCUSSION PAPER NO. 1764R2



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Should Auctions be Transparent?*

Dirk Bergemann[†] Johannes Hörner[‡]

August 30, 2017

Abstract

We investigate the role of market transparency in repeated first-price auctions. We consider a setting with independent private and persistent values.

We analyze three distinct disclosure regimes regarding the bid and award history. In the minimal disclosure regime each bidder only learns privately whether he won or lost the auction. In equilibrium the allocation is efficient and the minimal disclosure regime does not give rise to pooling equilibria. In contrast, in disclosure settings where either all or only the winner's bids are public, an inefficient pooling equilibrium with low revenues exists.

*We gratefully acknowledge financial support from NSF SES 0851200 and ICES 1215808. We would like to thank seminar participants at the NBER Market Design Meeting, University of Mannheim, University of Munich and the University of Pennsylvania for their constructive comments. Lastly, thanks to Yi Chen for excellent research assistance.

[†]Department of Economics, Yale University, New Haven, CT 06511, dirk.bergemann@yale.edu

[‡]Department of Economics, Yale University, New Haven, CT 06511, johannes.horner@yale.edu

1 Introduction

1.1 Motivation

Information revelation policies vary widely across auction formats. In the U.S. procurement context, as a consequence of the “Freedom of Information Act,” the public sector is generally subject to strict transparency requirements that require full disclosure of the identity of the bidders and the terms of each bid. In auctions of mineral rights to U.S. government-owned land, however, only the winner’s identity is revealed. In many markets, only the winner’s bid *and* identity are disclosed. This is the case, for instance, in the mussels sealed-bid auction documented by Kleijnen and Schaik (2007), and also happens in some European procurement auctions (for instance, in the London bus routes auctions, see Cantillon and Pesendorfer (2006)). This is also, by the very design of the auction, the bidders’ feedback in Dutch or English auctions. In some markets, even less information is disclosed. Over the last two years, the online auction site eBay has progressively moved toward a less transparent auction format. Bidders’ identities are no longer disclosed, although it remains possible to determine the list of items won in the last month by any of the (anonymous) user identities. Auction houses like Christie’s or Sotheby’s often preserve the anonymity of the winning bidder, and sometimes of the transaction amount. Another instance of an opaque mechanism is the sponsored search auction of Google in which the algorithm based on which the winner is determined is publicly unknown.

Undoubtedly, the choice of such feedback policies reflects a variety of considerations, such as security, privacy, risk of corruption, or risk of collusion. This paper focuses on the impact of these policies on bidders’ strategies, efficiency and revenue. We consider infinitely repeated first-price auctions, with persistent, independent private values, and multi-unit demand. We shall consider three information policies. With *unobservable bids*, bidders are only privately informed at the end of each round whether they have won the auction or not. With *observable bids*, all bids are disclosed at the end of each round. Finally, with *winner-only observable bids*, only the winner’s bid (and, although it plays no role, his identity) is publicly disclosed at the end of each round. In terms of the

disclosure policies, we therefore focus on the provision of past bidding information for future auction events. This should be viewed as distinct from the provision of bidding information *within* a single auction, a topic of central interest for the comparison of the English or Dutch auctions to standard sealed bid auctions, see Milgrom and Weber (1982). Thus, our concern for transparency is narrowly focused on the amount feedback given about past auction events. But importantly, the rules of the bidding game, namely the first price auction, are commonly known among the bidders, and thus we are not concerned with transparency with respect to the rules of the game.

As we investigate a bidding game with an infinite horizon, as in any infinitely repeated game, collusive equilibria exist if bidders are sufficiently patient, even under the most restrictive feedback policy.¹ Bid rotation, for instance, is always a possibility. To evaluate the intrinsic performance of each policy absent any tacit but explicit collusion, we focus on Markov equilibria, in which strategies only depend on bidders' beliefs.

Providing more information to the bidders about the competing bids has conflicting effects. If more information about bids is disclosed, bidders have an incentive to submit low bids to mimic bidders with low valuations, and induce high valuation bidders to lower their bid so as to win more easily in later periods. This is the familiar ratchet effect that suggests that more information is bad for revenue. On the other hand, if less information is provided, a winning bidder has an incentive to lower his bid to learn more about his opponents' bids. If these bids were observable, such a discovery process would be futile, but when the bids of the losers' remain undisclosed, it becomes valuable: to determine how low he can bid and still win, a past winner has an incentive to depress his bids. This learning effect suggests that less information is bad for revenue. While both effects are present in our model, we shall see that the first one clearly dominates the second as the discount factor tends to one.

Of the three policies, the policy of unobservable bids is the most challenging to analyze, but it is also probably the most interesting one. Because bidding histories are private, higher-order beliefs arise naturally. A past winner's belief about his opponents' value naturally depends on the bid with

¹At least as long as feasible allocations exist that guarantee each player his minmax payoff, see Section 4.

which he has won (winning with a very high bid, for instance, is not very informative). But losers have not observed the winning bid, and so, based on their losing bid, they must form beliefs not only about the winner's private value, but also about the winner's beliefs about the other bidders' values. In turn, because the losing bids are not observed, the winner must therefore form beliefs about the losers' beliefs about his belief, etc. Therefore, the relevant state space is the rather formidable universal belief space introduced by Mertens and Zamir (1985). To have any hope at making some progress, we restrict attention throughout to binary valuations. Even then, establishing equilibrium existence, let alone uniqueness, is rather difficult.

Fortunately, it is possible to explicitly construct a Markov equilibrium. In this equilibrium, high-valuation bidders always bid strictly more than low-valuation bidders, so that the allocation is efficient. The high-valuation bidder who wins in the initial period cautiously decreases his bids over time, trading-off the opportunity of winning with a slightly lower bid with the risk of losing and, more importantly, of generating mutual knowledge that he is not the only high-valuation bidder, which leads to higher future bids. As we show, a high-valuation bidder who loses in the initial period does not need to increase his later bids. In equilibrium, such a bidder can expect the winner to come down with his bids over time. In fact, it turns out that submitting bids that are constant over time is optimal for such a bidder. We provide closed-form expressions for the equilibrium strategies, which allows us to study expected revenue and perform comparative statics. In particular, we show that, as the discount factor goes to one, or equivalently, if auctions are repeated frequently enough, this revenue approaches the revenue of the optimal auction (without reserve price).

In contrast, when bidders have more feedback, a low-revenue pooling equilibrium might exist, which is impossible under unobservable bids. With winner-only observable bids, this pooling equilibrium is not unique. Indeed, there always exists a separating equilibrium whose revenue also tends to the maximal revenue as the discount factor goes to one. In contrast, with observable bids, the existence of a pooling equilibrium rules out the possibility of a separating equilibrium.

In the comparison of the minimal, the maximal and the intermediate disclosure policy, we establish in Proposition 1 that the policy of minimal information disclosure renders a pooling equilibrium

impossible. In contrast, under the more transparent information policies, a low revenue pooling equilibrium always existed as long as the number of bidders and the discount factor were not too low, as established in Proposition 3 and 5. In fact, under the maximal disclosure policy, we established that with a sufficiently large number of competitors an efficient, separating equilibrium ceased to exist, see Proposition 2. In contrast, under minimal disclosure, we proved the existence of an efficient separating equilibrium by construction, even though the bidders ceased to have common knowledge of beliefs after the initial bidding period. The combination of these results then lends support to minimal information disclosure, both from an efficiency and from a revenue maximizing point of view.

The equilibria in the distinct disclosure regimes can alternatively be cast in terms of the “ratchet effect”. The private information of each bidder is completely persistent and hence each bidder attempts to use (and hence hide or reveal) his private information optimally over the course of the bidding game. In the canonical analysis of the ratchet effect, the strategic environment is described by the relationship between the agent (with private information) and the principal (without commitment power). In the current auction environment, the seller is the principal, and is actually committed to the first price auction format. By contrast to the canonical model, there are many competing agents, the bidders, which lack in commitment in the sense that each bidder optimally adapts his present bid to the past bidding data. Thus, the ratchet effect appears as each bidder seeks to optimally use his private information against his competitors. The pooling strategy, and the resulting pooling equilibrium, is then the result of the ratchet effect in the competitive bidding environment. Yet, in the competitive environment, the ratchet effect is weakened by the competition among the bidders. Each bidder can increase his probability of winning only by increasing his bid and hence reveal additional information about his valuation. Importantly, the strength of the ratchet effect is affected by the disclosure regime. If the disclosure regime does not reveal all the past bidding data, the inference and hence the effectiveness of each strategy is impacted.

We interpret these findings as consistent with the common wisdom that more transparency is likely to hurt revenue. For instance, OECD guidelines for public procurement state that “disclosing information such as the identity of the bidders and the terms and conditions of each bid allow

competitors to detect deviations from a collusive agreement, punish those firms and better coordinate future tenders” (see OECD (2008)). Note, however, that, as mentioned, our findings do not rely on explicit collusion. The observable bids format is *inherently* more collusive than the unobservable bids format. Our analysis thus supports empirical findings, such as those of Albaek, Mollgard, and Overgaard (1997), and experimental findings, such as those of Cason, Kannan, and Siebert (2009), showing how finer public feedback may lead to lower revenues (and in experiments, pooling behavior). But our analysis also indicates that such findings must be interpreted with care, as the lower revenues need not be evidence of explicit collusion, but rather, of necessary adjustments in light of a new environment. In the words of Dave McCormick, Senior Vice President of FreeMarkets Inc., “suppliers are finding that, in a transparent environment where competitors can see each others’ bids, the price for goods is being driven down.” (Wilson (2000)). A second caveat is that it is not necessary to suppress all information to obtain efficiency. With winner-only observable bids, an efficient, high-revenue equilibrium exists as well. Simply, this equilibrium is not unique, and given this multiplicity, it should then come as no surprise that the impact of feedback might be limited in some settings, as, for instance, in Cramton and Schwartz (2002). Our bidding dynamics under unobservable bids are also consistent with the experimental results of Selten and Buchta (1999), who find that winner’s bids tend to be lowered or not to be increased, while losers’ bids tend to increase or not to be lowered over time.

While we view procurement auctions that frequently adopt first price rather than second price formats as a leading example, our work is also motivated by sponsored search auctions. The first generation of auctions in this environment were first price auctions, which have then gradually been replaced by generalized second price auctions in the past decade. In a second price auction and its generalizations, our present concerns over the strategic revelation of information become less relevant. However, the recent development in the sponsored search auctions, led by Google over the last five years, increasingly discloses less and less about the nature of the auction run. As of now, the only guaranteed aspect of the sponsored search auction is that each bidder pays no more than his bid. Google no longer claims to be running a generalized second price auction, and algorithms that take

into account current and past bids are being employed for the determination of the winning bid price (cost-per-click).² Thus, the bid of the advertiser determines its Ad Rank, and therefore its eventual payment. Other sponsored search auctions are more explicitly following first price auctions, such as Facebook or Microsoft Bing Ads auction, which explicitly takes the bid of the advertiser to calculate the cost-per-click.³ These current practices suggest that elements of the first price auction are reappearing in the sponsored search auction.

We should be clear that other considerations affect the choice of transparency. Collusion and corruption do not only involve buyers, but also auctioneers. It is intuitively clear that too much opaqueness facilitates corruption of the auctioneer by the buyers. More generally, the auctioneer must be trusted to follow the auction rules that he adopts. Yet it is not hard to see how, even in first-price auctions, an auctioneer could take advantage from naive bidders by allocating the unit to the low bidder in a given period, in order to make the high bidder more aggressive in his future bids. It is less clear whether such manipulation can be profitable if bidders understand the auctioneer's incentives. Google's choice not to fully disclose its rules for sponsored search auctions, mentioned in the introduction, raises then an interesting puzzle.

1.2 Related Literature

There are a number of recent contributions in auction theory that consider similar information environments. Landsberger, Rubinstein, Wolfstetter, and Zamir (2001) analyze the first-price auction in a static environment when only the ranking of the valuations is common knowledge. Their analysis is motivated by the information revealed through the interaction in repeated bidding environments. The main focus of their paper is the analysis of the specific asymmetric auction environment that results when two bidders, possibly starting with the same common prior over the valuation, receive additional information about their ranking with respect to their competitor. In a model with a

²For one of the few public statements, see: <https://support.google.com/adwords/answer/6297>.

³See <http://advertise.bingads.microsoft.com/en-us/blog/27821/bing-ads-auction-explained-how-bid-cost-per-click-and-quality-score-work-together>.

continuum of valuations, they establish the existence and uniqueness of a pure strategy Bayes-Nash equilibrium. They also show, by example, that the equilibrium bidding strategies can typically not be expressed as an analytic function, due to a singularity in the bidding function at the lower end of the valuations.

Fevrier (2003) extends the analysis of Landsberger, Rubinstein, Wolfstetter, and Zamir (2001) from 2 to n bidders. He then compares the revenue generated by the sale of two identical units of an object in the sequential auction over two periods to the revenue when the units are sold as a bundle in a single auction. Fevrier (2003) establishes that the revenue in the static auction of the bundle yields a higher revenue than the sequential auction with or without the announcement of the winner in the first period. Yao (2007) analyzes the equilibrium in a two-period model when the winning bidder and the winning bid is revealed after the first period. In particular, he finds that the initial bids in the two-period model are uniformly lower than the bids in the static first-price auction. Tu (2007) compares the revenue properties of a number of auction formats and disclosure policies in a two-period setting. In particular, he establishes revenue comparisons when the valuation of the buyers are uniformly distributed. Notably, he finds that with respect to a number of possible disclosure policies in the first-price auction, announcing the winning bid yields a higher revenue than announcing the winning and the losing bid. In turn, the revenue from the public announcement of the bids yields a higher revenue than the announcement of the determination of winner and loser in the past auction. Thomas (2010) also compares different disclosure policies within a two period model with two possible types for each bidder. By imposing a reserve price at the low valuation, he restricts attention to separating equilibria in which low valuation bidders never submit bids at all. This restriction completely removes the possibility of pooling equilibria in the analysis.

The role of the disclosure policy and the information flow is central also in the recent work of Kaya and Liu (2015) who consider a model of sequential bargaining with a sequence of short-lived seller. Similarly to our auction environment, they come to the conclusion that unobservability lowers the price that the competing sellers obtain and increases the speed of trading.⁴

⁴The impact of the disclosed (or publicly observable) information is also central in the recent work on optimal

The remainder of the paper is organized as follows. Section 2 presents the model and the rules of information disclosure. Section 3 considers the equilibrium bidding strategies with unobservable bids. It explicitly constructs an equilibrium in separating strategies and establishes comparative statics. Section 4 considers the environment with observable bids. Section 5 analyzes the equilibrium when only the bid of the winner is observable. Section 6 concludes and Section 7, the Appendix, collects the remaining proofs.

2 The Model

2.1 Three Variations on a Theme

There are $n + 1 \geq 2$ bidders, or players, competing in an infinite sequence of auctions. In every period $t = 0, 1, \dots$, a single unit is sold via a first-price sealed bid auction. There is no reserve price, and ties are broken randomly. Bidders have quasilinear preferences that are additively separable across periods. A player's valuation u^i , or type, is constant across time and private information. Valuations are binary: bidder i 's type is either high and equal to \bar{u} or low, equal to \underline{u} . Types are drawn independently across bidders, and the probability that bidder i 's valuation is \bar{u} is $q \in (0, 1)$. The same $n + 1$ bidders participate in all these auctions, and both the number of bidders, and the type distribution, are common knowledge among bidders. These assumptions, discussed in the conclusion, are quite restrictive (in particular, the binary valuations of the bidders), but it will become clear that relaxing them appears difficult in the case in which bids are not observable.⁵

Thus, the reward r_t^i in period t of player i with valuation u^i is equal to $u^i - b_t^i$ if he bids b_t^i and he

dynamic auctions where the private information of each bidder is given by a multi-dimensional type, say his valuation and the deadline by which the bidder needs to acquire the object. For example, in Dizdar, Gershkov, and Moldovanu (2011) and in Pai and Vohra (1979), the partial disclosure of information can increase the revenue of the auctioneer as it reduces the dimensionality of the incentive constraints. The disclosure in these papers directly pertains to the private information rather than the strategy (bid) of the agents and hence results in distinct implications of information disclosure.

⁵With more than two values, the Markov assumption also loses its bite in the arguably simpler case of observable bids, *i.e.*, the equilibrium outcome is no longer unique –not much of a surprise given the leeway in specifying out-of-equilibrium beliefs. See Hörner and Jamison (2008) for a discussion of such examples in a related environment.

wins the object, or to 0 if he does not win. Bidders discount future periods with a common discount factor $\delta < 1$. The realized payoff of a bidder is the average discounted sum of his rewards:

$$\sum_{t=0}^{\infty} (1 - \delta) \delta^t r_t^i.$$

Our purpose is to compare different information policies available to the auctioneer. In all cases, every individual bidder is privately informed, at the end of any given period, whether he has won the unit in that period or not. We compare three scenarios:

1. In the *unobservable* case, bids are not disclosed. The identity of the winner is not disclosed either. Of course, if $n + 1 = 2$, a bidder can infer who won from his own information (whether he won or lost), but this is no longer the case with more bidders.
2. In the *observable* case, the auctioneer discloses who bid how much. This is the case of perfect monitoring, and bidders accordingly update their beliefs about the valuations of others.
3. In the *winner-only observable* case, the bid of the winning bidder is announced. Although this turns out to be irrelevant for our analysis, we also assume that the winner's identity is disclosed. Nothing else is disclosed.

In a repeated game such as ours, even with incomplete information, there is a myriad of equilibria. For instance, there are collusive equilibria that involve bid rotation, and a winning bid of zero in every period, which are easy to support if $\underline{u} > 0$, independently of the structure of the uncertainty. Because we are not interested in collusion *per se*, we focus on Markov equilibria, in which players' strategies only depend on payoff-relevant information.

What information is payoff-relevant in our environment is tricky. In the observable case, players' beliefs (about others' values) are public after every history, and we take these beliefs as the state variable. A similar definition is possible in the winner-only observable case. This, however, is difficult in the unobservable case. For instance, a winner infers from his bid how high his opponents' bids

could be, and this affects his beliefs about their valuations. His beliefs, however, are no longer common knowledge, because his bid is not. Because a loser can only deduce a lower bound on the winner's bid from his own bid, the loser has beliefs about the winner's beliefs, and they are certainly payoff-relevant from his viewpoint. We are therefore led to consider the universal type space (see Mertens and Zamir (1985)) as the natural state space for our definition of Markov strategies.

But first, let us define precisely the bidders' information and strategies in each scenario.

2.2 Histories and Strategies

Even under complete information, it is often convenient to introduce infinitesimal bids in order to avoid complications linked to real numbers: if bidder 1 is known to be of value \underline{u} , and bidder 2 is known to be of value \bar{u} , it is natural, in the one-shot game, to focus on the equilibrium in which bidder 1 bids \underline{u} and bidder 2 bids "as little as possible" above \underline{u} . Of course, there is no such bid in the field of real numbers. To avoid this difficulty, one can resort to richer strategies, as in Blume (2003), to endogenous tie-breaking rules, as in Jackson, Swinkels, Simon, and Zame (2002), or to arbitrarily fine but discrete bid grids, as in Chwe (1989). As a convention, we shall follow here Engelbrecht-Wiggans (1983) and Hörner and Jamison (2008) and assume that there is such a bid \underline{u}_+ , which costs just as much as \underline{u} , but that is strictly larger, while being strictly smaller than any real number $b > \underline{u}$.

A private history of player i up to period t is a sequence $(b_0^i, k_0^i, \dots, b_{t-1}^i, k_{t-1}^i)$, consisting of the bids $b_{t'}^i \in \mathbb{R}_+ \cup \{\underline{u}_+\}$ that he made in period t' , and of his personal outcome in that period: $k_{t'}^i = 0$ if bidder i did not win the object in period t' , and $k_{t'}^i = 1$ if he won the object. A private history of player i up to period t is denoted $h_t^i \in H_t^i := ((\mathbb{R}_+ \cup \{\underline{u}_+\}) \times \{0, 1\})^t$.

In the unobservable case, this is the only information available to player i , and a (behavior) strategy σ^i is then simply a countable sequence of transition probabilities $\sigma_t^i : \{\underline{u}, \bar{u}\} \times H_t^i \rightarrow \Delta(\mathbb{R}_+ \cup \{\underline{u}_+\})$, mapping bidder i 's valuation and private history into a distribution over bids.

In the observable case, bidder i knows the entire sequence of bids in each period up to t . The public history up to period t is thus a sequence $((b_0^1, \dots, b_0^n), j_0, \dots, (b_t^1, \dots, b_t^n), j_t)$, with $(b_{t'}^1, \dots, b_{t'}^n) \in$

$(\mathbb{R}_+ \cup \{\underline{u}_+\})^n$, and $j_{t'} \in \{1, \dots, n\}$, where $b_{t'}^{j_{t'}} := \max_i b_{t'}^i$. Of course, the identity $j_{t'}$ of the winner in period t' can be inferred from the ordered bid tuple (unless there is a tie). The public history up to t is denoted $h_t \in H_t := ((\mathbb{R}_+ \cup \{\underline{u}_+\})^n \times \{1, \dots, n\})^t$. (Set $H_0^i := \{\emptyset\}$, $H_0 := \{\emptyset\}$.)

In the winner-only observable case, a public history up to t is a sequence $(b_0^w, j_0, \dots, b_{t-1}^w, j_{t-1})$ of winning bids $b_{t'}^w$ in period t' , and of the identity of the winning bidder in that period: $j_{t'} \in \{1, \dots, n\}$ refers to the winning bidder in that period. The set of public histories H_t in this case is equal to $((\mathbb{R}_+ \cup \{\underline{u}_+\}) \times \{1, \dots, n\})^t$. For consistency, we set $H_t := \{\emptyset\}$, all t , in the unobservable case, so that we may talk about the three scenarios in a unified way.

A behavior strategy σ^i for player i in the observable case or the winner-only observable case is again a countable sequence of transition probabilities $\sigma_t^i : \{\underline{u}, \bar{u}\} \times H_t^i \times H_t \rightarrow \Delta(\mathbb{R}_+ \cup \{\underline{u}_+\})$, mapping bidder i 's valuation, along with the private and public history up to period t , into a distribution over bids.

2.3 Solution Concept

A strategy profile $\sigma = (\sigma^i)_i$ defines a probability distribution \mathbb{P}_σ over infinite histories in the obvious way, and we can therefore define player i 's payoff under the strategy profile σ as the expectation of his realized payoff relative to this distribution

$$V^i(\sigma) = \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} (1 - \delta) \delta^t r_t^i \right].$$

Fix some strategy profile σ . Given the common prior on bidders' valuations, and given any pair of private and public histories (h_t^i, h_t) that are in the support of the distribution \mathbb{P}_σ , Bayes' rule determines bidder i 's beliefs about the other bidders' valuations and their private histories h_t^j , $j \neq i$. This in turn defines a conditional distribution $\mathbb{P}_{\sigma | (h_t^i, h_t)}$ over the sequence of future rewards, and we

can define the continuation payoff of player i after (h_t^i, h_t) as

$$V^i(\sigma|(h_t^i, h_t)) = \mathbb{E}_{\sigma|(h_t^i, h_t)} \left[\sum_{t'=t}^{\infty} (1 - \delta) \delta^{t'-t} r_{t'}^i \right].$$

We may then define a perfect Bayesian equilibrium (or PBE, for short) as a strategy profile σ in which players' strategies σ^i are sequentially rational after every pair (h_t^i, h_t) given their beliefs, and these beliefs are consistent with Bayes' rule if this pair is in the support of \mathbb{P}_σ .

As mentioned, we are not interested in characterizing all PBE. It is natural to focus on Markov equilibria, in which players' strategies are measurable with respect to their beliefs. However, we have seen that, at least in the unobservable case, attention cannot be restricted to first-order beliefs; even if player i conditions on j being of the high type, he cannot infer player j 's first-order beliefs from his own private history h_t^i , because player j 's high type might randomize over bids, and what determines j ' first-order beliefs is the realization of these bids, *i.e.* player j 's private history. We are thus led to adopt as state space for player i the universal belief space Θ^i (see Mertens and Zamir (1985)), which is compact and metric. Given the strategy profile σ , a pair of histories (h_t^i, h_t) in the support of \mathbb{P}_σ defines a belief $\theta^i \in \Theta^i$ via Bayes' rule. Player i 's strategy is Markov if it is measurable with respect to these beliefs. It is natural to further impose that player i 's strategy is also measurable with respect to his belief θ^i off-path as well, even if these beliefs are no longer determined by Bayes' rule. A Markov strategy is then summarized by a measurable map $\sigma^i : \Theta^i \rightarrow \Delta(\mathbb{R}_+ \cup \{\underline{u}_+\})$, and a Markov equilibrium (hereafter, MSE) is a PBE in Markov strategies. In the observable case, these hierarchies of beliefs are trivial, and they turn out to simple in the case of winner-only observable bids as well. They are, however, more complicated in the unobservable case.

Because of the arbitrariness of the specification of beliefs off-path, these beliefs can be used to threaten players, so that the Markov restriction does not reduce the set of equilibria as much as one would like to. Consider for instance the observable case with two bidders. Fix some history after which it is commonly believed that the two bidders have low valuations. It would be natural, then, to conjecture that in a Markov equilibrium, after such a history, both bidders set their bid equal to

\underline{u} in every period. But any lower common bid would do as well, as long as the equilibrium specifies that any higher bid will lead to a belief revision. For instance, if bidder i observes j bidding more, we could specify that i now believes that j has a high valuation after all, and then bids \underline{u} thereafter. This deters any deviation. To prune such artificial equilibria, we impose the following refinement.

Refinement A:

1. *After any history $((h_t^i)_i, h_t)$, low-type bidders bid \underline{u} in every period.*
2. *After any history $((h_t^i)_i, h_t)$ such that it is common knowledge among at least two high-type bidders that they are both high-type bidders, those two high-type bidders bid \bar{u} thereafter.*

The first restriction is a combination of two assumptions. First, a low-type bidder does not use a weakly dominated strategy, such as bidding strictly more than \underline{u} . Second, all bids are at least as high as the lowest commonly known value (while (1.) does not impose that the high-type bidder bids at least \underline{u} , it is easy to see that it will imply it). Note that the second part of the refinement does not require that the two high-type bidders know their respective identities. Rather, it suffices that it be common knowledge among them that they exist. Still, Refinement **A** will not ensure uniqueness, but it will help narrow down the set of candidate equilibria considerably.

Note that we have now pinned down, by assumption, the equilibrium behavior of the low-type bidder. Therefore, the difficulty lies in identifying the behavior of the high-type bidder.

2.4 The Static Auction

We conclude this section with a brief review of the static first-price auction, with $n + 1$ bidders, $i = 1, \dots, n + 1$, two possible valuations each, $u^i \in \{\underline{u}, \bar{u}\}$, and identical and independent priors given by $1 - q$ and q , respectively. With discrete, here binary, valuations, the unique equilibrium of the first-price auction involves randomization by the high-valuation bidder (see Maskin and Riley (2003)). His bid has to balance the probability of a winning bid against the price paid conditional

on winning. The unconditional distribution of bids from each of his competitors is denoted by $F(b)$. Every bid in the support of the random bidding strategy must maximize the expected payoff

$$F(b)^n (\bar{u} - b).$$

Indifference of the high-valuation bidder requires that the right-hand side be independent of b , *i.e.*

$$F(b) = (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}, \quad (1)$$

where the support of the distribution is given by $[\underline{u}, \bar{u} - (1 - q)^n(\bar{u} - \underline{u})]$. The distribution displays a mass point at the lower extremity of the support, where $F(\underline{u}) = (1 - q)^n$ reflects the fact that the low-valuation bidder makes a deterministic bid equal to his valuation. In contrast, the high-valuation bidder continuously randomizes on $(\underline{u}, \bar{u} - (1 - q)^n(\bar{u} - \underline{u}))$. The low-valuation bidder receives zero net payoff, while the high-valuation bidder receives a positive payoff, given by $(1 - q)^n (\bar{u} - \underline{u})$. In the first-price auction, each type's payoff is equal to his payoff from the second-price auction. Yet, with discrete types, the revenue equivalence theorem fails, and revenue might differ across mechanisms that are efficient and yield no surplus to the low-type bidder. We shall encounter such mechanisms. However, the allocation from the first-price auction maximizes revenue among efficient mechanisms.

3 Unobservable Bids

We begin our analysis with the case of unobservable bids.

3.1 On the Impossibility of Pooling

An equilibrium is pooling if, on the equilibrium path, bidders of different valuations use the same bidding strategy, so that, equivalently, beliefs do not change. Note that, if the strategies of the bidders act to separate types, then a high-valuation bidder will (eventually) win against a low-valuation

bidder. In the process of separation, the high-valuation bidder also reveals his true valuation and consequently might be forced into an eventual competition with another high-valuation competitor that will leave both of them with no surplus. From this point of view, a pooling equilibrium might seem desirable for the bidders, especially when the probability that a bidder has a high valuation is high, and when bidders are patient. Indeed, such equilibria will arise under other information structures, as we shall see. Pooling implies that the surplus must be shared, in particular with low-valuation bidders. The benefit is that the price remains low.

Consider then such a candidate pooling equilibrium. As the bids are not observable, any loss can be attributed to pure chance (given the random tie-breaking) and does not lead to a revision of the prior. This opens the door for a high-valuation bidder to bid slightly more than the pooling bid and to win the current auction for sure. As beliefs of the agents are unchanged, the current benefit comes without a future cost, and this represents a profitable deviation.

Proposition 1 (Impossibility of Pooling).

For all q, n, δ , a pooling Markov equilibrium does not exist with unobservable bids.

Refinement **A** is not necessary for Proposition 1. To see this, note first that a pooling equilibrium must involve pure strategies, because it is not possible, given single-crossing, that low- and high-type bidders are simultaneously indifferent over two bids (*i.e.*, over two distinct probabilities of winning: the probability of winning in the continuation equilibrium must be independent of this bid, by the Markov assumption, and by the fact that the observed bid does not affect beliefs in a pooling equilibrium). This pooling bid must (at least in some period) be no larger than \underline{u} , else the low-type bidder would make negative profits. Pick any such period, and apply the argument above.

Having established the impossibility of pooling in a Markov equilibrium, we now proceed to construct a specific separating equilibrium.

3.2 The Separating Equilibrium: Preview

In a separating equilibrium, low- and high-valuation bidders' strategies have disjoint supports, which allows some learning to take place. Of course, with unobservable bids, this learning might be incomplete. For instance, a high-valuation bidder that wins in the initial period only infers that his opponents have bid less than he did, but that does not allow him to ascertain his opponents' valuation for sure. He simply updates his beliefs given his winning bid. So do the losing bidders, given their losing bids. Further, the losing bidders revise their beliefs about the winner's beliefs, but because they do not know the winning bid, this leads to a subtle updating process.

We shall circumvent these difficulties as follows. Recall that low-valuation bidders bid \underline{u} throughout, so the focus is on the high-valuation bidders. The separating equilibrium we shall construct has the following properties:

1. In the initial period, high-valuation bidders continuously randomize over the support $[\underline{u}_+, \bar{b}_0]$, for some $\bar{b}_0 > \underline{u}_+$. This partitions the set of bidders according to their status after the initial period, as “winner” and “losers.” We shall refer to a bidder as the winner entering period t if he won in all periods up to t , and as a loser if he lost in all those periods.
2. In later periods, *as long as the (initial) winner has never lost*, as a function of their initial bid, (high-valuation) bidders submit bids that decrease over time.⁶ Both the (high-valuation) winner and the loser always bid strictly more than \underline{u} , but, depending on his initial bid, the winner might bid \underline{u}_+ . More precisely, for every period, there exists a range of bids in $[\underline{u}_+, \bar{b}_0]$, that includes \underline{u}_+ , such that, if the winner has won in the initial period with a bid in this range, then his bid in period t and beyond is \underline{u}_+ . The support of the bid distribution of the winner and (high-valuation) losers is common, *i.e.*, the highest bid that a loser could conceivably make in a given period, *i.e.*, the bid a loser would make if he lost in the initial period with a bid of \bar{b}_0 , coincides with the bid the winner would make if he had initially won with \bar{b}_0 .

⁶We insist that, although for convenience we describe later bids as functions of earlier bids, they are truly functions of the bidders' beliefs, which on the equilibrium path happen to be pinned down by their initial bid.

We note that such an equilibrium would have the desirable feature that, as soon as a high-valuation bidder who always won so far loses in some period $t > 0$, it would become common knowledge among two bidders that there are two high-valuation bidders.⁷ To see this, note that the high-valuation loser who then wins knows that there exists another high-valuation bidder, because he lost in the initial period with a bid strictly above \underline{u} . But as the winner eventually loses with a bid strictly above \underline{u} , this winner learns that there is another high-valuation bidder, and thus, that there is another bidder who knows that there are two high-valuation bidders. Because they both know that the winner has lost in period t , this establishes common knowledge (among them) that there are two high-valuation bidders. By Refinement **A**, bids then jump up to \bar{u} , which ends the game for all practical purposes. We may then focus on the histories in which the “winner” of the initial auction has never lost afterwards.

In such an equilibrium, the process of belief updating is simple. Assume there are two bidders. Consider the high-valuation winner’s inference problem. Given that the loser is using a monotone strategy, the winner’s belief can be summarized by a cut-off bid. Namely, the winner can derive an upper bound on the bid that the loser might submit in the current period, which is the highest bid consistent with the loser’s equilibrium strategy, given that all his bids were below the winner’s bids until then. While the winner knows that the loser will not bid above this cut-off, his private information gives him no further information regarding the relative likelihood of lower bids. Therefore, a belief revision for the winner amounts to truncating (from above) the corresponding distribution. Updating proceeds similarly for the high-valuation loser. His private history provides him with a lower bound on the bids that the high-valuation winner might submit. Therefore, a belief revision for the loser amounts to truncating (from below) the corresponding distribution.

The next subsection shows how to explicitly solve for the equilibrium strategies. The reader mostly interested in the qualitative findings might elect to skip it.

⁷Note, however, that with more than two bidders all but two bidders will never learn for sure whether the initial winner has already lost, in which case there is no longer any scope for them to win.

3.3 Deriving the Equilibrium Strategies

Fix a separating Markov equilibrium. Let F_t denotes the cumulative and *unconditional* distribution function (c.d.f.) summarizing the equilibrium strategy in period t of a player who always lost up to period $t - 1$, and G_t the *unconditional* c.d.f. summarizing the equilibrium strategy in period t of a player who always won up to period $t - 1$. That is, F_t captures the winner's belief about the bid distribution of any given loser in period t , if he had submitted in all previous periods bids with which he was sure to win, and were thus uninformative. Similarly, G_t describes the belief about the winner's distribution of a loser who would have bid less than \underline{u} throughout.

The distinction between winner and loser is immaterial initially, and thus $G_0 = F_0$. Given the properties of the separating equilibrium we seek, it must be that $F_t(\underline{u}) = 1 - q$ (recall that $1 - q$ is the prior probability of a low valuation). By contrast the high-valuation bidder who won until t might bid \underline{u}_+ with discrete probability, *i.e.* $G_t(\underline{u}) = 1 - q$ but $G_t(\underline{u}_+) \geq 1 - q$.⁸

In general, a player's beliefs are pinned down by his entire private history. It turns out that the last bid (along with the bidder's status as winner or loser) is a sufficient statistic for this belief, at least on the equilibrium path, on which we focus for now. Thus, we denote by $V_t(b)$ the continuation value of the winner with a high valuation \bar{u} , given that his last bid was b . Similarly, we denote by $W_t(b)$ the continuation value of a loser with a high valuation \bar{u} , given that his last bid was b . We emphasize that this is just a convenient short-hand for the player's beliefs.

The derivation below is performed for the case of two bidders. This makes the exposition somewhat easier. Results, however, are stated for $n + 1$ bidders, with proofs in Appendix.

3.3.1 The Loser's Bidding Strategy

We start by determining the equilibrium bid distribution F_t of the loser for all periods $t \geq 1$. The bid distribution F_t of the loser is determined by the indifference condition of the winner. His

⁸To avoid clutter, we shall just omit the distinction between $G_t(\underline{u})$ and $G_t(\underline{u}_+)$, with the convention that $G_t(\underline{u}) - (1 - q)$ is the probability assigned to the bid \underline{u}_+ , as the probability assigned to \underline{u} is $1 - q$ throughout.

continuation value is given by the optimality equation

$$V_t(b) = \max_{\beta} \left\{ \frac{F_t(\beta)}{F_{t-1}(b)} [(1 - \delta)(\bar{u} - \beta) + \delta V_{t+1}(\beta)] \right\}, \quad t \geq 1. \quad (2)$$

The winner receives the object in period t with a bid $\beta > \underline{u}$ if and only if the loser makes a bid below β .⁹ The ratio $F_t(\beta)/F_{t-1}(b)$ is the conditional belief of the winner, and is obtained by truncation of his original, unconditional belief, as explained above. The unconditional probability of a bid below β is given by $F_t(\beta)$. The winner received the object in the preceding period with a bid b , hence he can condition his bid β today on the information that the loser made a bid below b yesterday (this, as it turns out, is finer information than the one contained in his earlier bids). If, for instance, he makes the bid $\beta_t(b)$ that the loser would submit after bidding b in period $t - 1$, he would win with probability $F_t(\beta_t(b))/F_{t-1}(b) = 1$, because, given monotonicity, $F_t(\beta_t(b)) = F_{t-1}(b)$, by definition of $\beta_t(b)$. The winner has no incentive to bid more than $\beta_t(b)$, since this bid suffices to win for sure.

In the case of a winning bid β , the winner receives the object today at the price β and maintains his status as winner for at least another period. By contrast, if he loses the auction today, then it is common knowledge among the bidders that they both have a high valuation. Hence, by Refinement **A**, all future bids will be equal to \bar{u} and exhaust all surplus from the bidders' point of view.

We define $Y_t(b)$ to be the expected future utility from a bid b in the preceding period, so

$$Y_t(b) := F_{t-1}(b) V_t(b). \quad (3)$$

This allows to rewrite the value function of the winner as

$$Y_t(b)/(1 - \delta) = \max_{\beta} \{F_t(\beta) (\bar{u} - \beta) + \delta Y_{t+1}(\beta)/(1 - \delta)\}, \quad t \geq 1, \quad (4)$$

⁹Implicitly, here and in the winner's problem, we restrict the domain of the choice variable β to the range of values that will preserve the feature that the last bid is a sufficient statistic for the entire past, and for which the ratio $F_t(\beta)/F_{t-1}(b)$ is less than one, *i.e.* such that $\beta \leq \beta_t(b)$. We will then verify that the strategy profile obtained in this manner is an equilibrium.

so that the unconditional distribution F_{t-1} of the preceding period $t - 1$ no longer appears. Note that the right-hand side no longer depends on b , so that $Y_t(b)$ does not either. That is, $Y_t(b)$ is constant, and the last term on the right-hand side, $Y_{t+1}(\beta)$, must be as well. Thus, the first term of the right-hand side must be constant over the support of F_t , and so

$$F_t(b) = \frac{\varphi_t}{\bar{u} - b}, \quad t \geq 1,$$

for some constant φ_t . Since the equilibrium we seek to identify satisfies $F_t(\underline{u}) = 1 - q$ for all $t \geq 1$, the constant φ_t is given by $\varphi_t = \varphi := (1 - q)(\bar{u} - \underline{u})$, independently of t for $t \geq 1$. We record below the equilibrium strategy for the loser in case that there are $n + 1$ bidders, as an immediate generalization of the formula above (see (37) in Appendix).

Lemma 1 (The Loser's Bid Distribution).

The loser's bid distribution is given by, for all $t \geq 1$,

$$F_t(b) = F(b) := (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}, \quad (5)$$

on the support $[\underline{u}, \bar{u} - (1 - q)^n(\bar{u} - \underline{u})]$. Thus, the loser makes a constant bid from $t \geq 1$ onward.

Equation (4) provides a simple difference equation for the sequence $\{Y_t\}_t$ of unconditional payoffs, namely, $Y_t/(1 - \delta) = \varphi + \delta Y_{t+1}/(1 - \delta)$, whose unique bounded solution is

$$Y_t = \varphi = (1 - q)(\bar{u} - \underline{u}), \quad (6)$$

which is independent of time and of the past bid b . With the solution to the loser's bid distribution and the unconditional payoffs, given by (5) and (6), we obtain the conditional value of the winner, $V_t(b)$, by using the equation (3) as: $V_t(b) = \bar{u} - b$. Using the recursion of the value function given by (2), it follows that the continuation value of the winner is a martingale, *i.e.* $\mathbb{E}[V_{t+1}] = V_t(b) = \bar{u} - b$.

3.3.2 The Winner's Bidding Strategy

Next, we derive the winner's unconditional bid distribution G_t , which in turn is determined by the loser's optimization problem. The value function W_t of the loser, as a function of his last bid, which here as well encapsulates his belief (on the equilibrium path) satisfies the optimality equation

$$W_t(b) = \max_{\beta} \left\{ \frac{G_t(\beta) - G_{t-1}(b)}{1 - G_{t-1}(b)} (1 - \delta)(\bar{u} - \beta) + \delta \frac{1 - G_t(\beta)}{1 - G_{t-1}(b)} W_{t+1}(\beta) \right\}, \quad t \geq 1. \quad (7)$$

To understand the loser's payoff, we must distinguish between two events. The contemporaneous bid $\beta > \underline{u}$ can either win the current auction, and hence yield a reward of $\bar{u} - \beta$ (after which bids jump to \bar{u}) or it can be too low, in which case the loser remains in his loser's status, until the subsequent period $t + 1$ when he can expect a continuation value $W_{t+1}(\beta)$. The unconditional probability of winning (or losing) with a bid β becomes a conditional probability by conditioning on the previous event, in which the loser lost with b , so that the winner's bid must have been at least as high. As before, it is useful to restate this equation with the help of an auxiliary function. Let $X_t(b)$ be the expected continuation value from losing with a bid b in period $t - 1$, or

$$X_t(b) := (1 - G_{t-1}(b)) W_t(b). \quad (8)$$

With this definition, we rewrite (7) to get

$$X_t(b) / (1 - \delta) = \max_{\beta} \{ (G_t(\beta) - G_{t-1}(b)) (\bar{u} - \beta) + \delta X_{t+1}(\beta) / (1 - \delta) \}, \quad t \geq 1. \quad (9)$$

The value function, described in terms of the unconditional expected values, is again more accessible than the conditional values. But we observe that the past bid b of the loser continues to appear on the right-hand side of the equation. First-order conditions are then

$$G'_t(\beta) (\bar{u} - \beta) - (G_t(\beta) - G_{t-1}(b)) + \frac{\delta}{1 - \delta} X'_{t+1}(\beta) = 0, \quad t \geq 1. \quad (10)$$

Note also that, from the envelope theorem applied to (9),

$$X'_t(b)/(1 - \delta) = -G'_{t-1}(b)(\bar{u} - \beta). \quad (11)$$

To make further progress, some calculations will be required, and they will necessitate to distinguish according to whether t is equal to, or larger than 1. As we learned earlier (see (5)), the equilibrium bid of the loser is constant across periods for $t \geq 1$, or $b_t = b_{t+1}$, so that the first-order condition must hold for the choice $\beta = b$. It follows that we can describe the bidding behavior in terms of contemporaneous bid b alone for all periods $t > 1$, *i.e.*, from (10),

$$G'_t(b)(\bar{u} - b) - (G_t(b) - G_{t-1}(b)) + \frac{\delta}{1 - \delta} X'_{t+1}(b) = 0. \quad (12)$$

The property of constant bids across periods only arose in the continuation game after an initial winner and initial loser had been determined. The relationship between the initial bid b and the bid β after the determination of the “winner” and “loser” position respectively has yet to be established, which is why we assume first that $t \geq 2$. After forwarding the time index from t to $t + 1$ in (11), and using the fact that $\beta = b$, we can eliminate $X'_{t+1}(\beta)$ from (10), to obtain

$$(1 - \delta)G'_t(b)(\bar{u} - b) = G_t(b) - G_{t-1}(b), \quad t \geq 2. \quad (13)$$

Because the support of F_t and G_t must coincide, it holds that $G_t(\bar{u}) = 1$. Thus, we have an ordinary differential equation and a boundary condition that can be solved for G_t , if G_{t-1} is given.

Let us turn to G_1 . We have already observed that the relationship between the contemporaneous bid β in period $t = 1$ and the preceding bid b in period $t = 0$ is more intricate than in later periods. Recall that the optimality equation in $t = 1$, derived earlier in (10), states that

$$G'_1(\beta)(\bar{u} - \beta) - (G_1(\beta) - G_0(b)) + \frac{\delta}{1 - \delta} X'_2(\beta) = 0. \quad (14)$$

While it is no longer the case that $\beta = b$ in $t = 1$, the hypothesis of monotone bidding strategies allows us to relate the bid b in $t = 0$ to the bid β in $t = 1$, noting that

$$G_0(b) = F_0(b) = F_1(\beta) = (1 - q)(\bar{u} - \underline{u}) / (\bar{u} - \beta), \quad (15)$$

where the final equality had been established in (5). Thus, the equations (13) and (14), along with $G_t(\bar{u}) = 1$, all $t \geq 1$, allow us to solve recursively for the distributions G_t , $t \geq 1$. The solution of (13) and (14) is the special case for two bidders of the formula given in the following lemma, where as before, $G_t(\underline{u}) - (1 - q)$ is the probability assigned by the high-valuation winner to the bid \underline{u}_+ .

Lemma 2 (The Winner's Bid Distribution).

The winner's bid distribution is given by, for all $t \geq 1$,

$$G_t(b) = \frac{1}{\delta^t} F(b) + F(b)^{\frac{1}{1-\delta}} \sum_{\tau=0}^t \frac{1 - \delta^{\tau-t}}{\tau!} \left(\ln F(b)^{-\frac{1}{1-\delta}} \right)^\tau, \quad (16)$$

on the support $[\underline{u}_+, \bar{u} - (1 - q)^n(\bar{u} - \underline{u})]$.

Recall that the distribution $F(b)$ refers to the loser's bid distribution obtained in Lemma 1. It follows from this formula that the bids of the winner are decreasing over time from $t \geq 1$ onward.

3.3.3 The Bidding Strategy in the Initial Period

We are left to determine the bidding strategy in the initial period $t = 0$ (at this stage, the distinction between winner and loser does not yet appear). Each high-valuation bidder maximizes

$$\max_b \left\{ F_0(b) (\bar{u} - b) + \frac{\delta}{1 - \delta} Y_1(b) + \frac{\delta}{1 - \delta} X_1(b) \right\}. \quad (17)$$

The bid b in the initial period determines the expected reward, as well as the continuation value, conditional on being the winner, $Y_1(b)$, or the loser, $X_1(b)$, where we maintain the notation, introduced in (3) and (8), that already accounts for the probability of each event. If high-valuation bidders are

indifferent over some interval, the profit from (17) must be independent of b over this interval. Thus, plugging $X_1(b), Y_1(b)$ into (17), straightforward algebra gives that

$$F_0(b)(\bar{u} - b) - \delta(1 - q)(\bar{u} - \underline{u}) \ln F_0(b) = K_0, \quad (18)$$

for some constant K_0 . This implicitly defines F_0 and can be explicitly solved using the Lambert function. The initial bid distribution at $t = 0$:

$$F_0(b) = (1 - q) \left(-\delta \frac{\bar{u} - \underline{u}}{\bar{u} - b} W_{-1} \left(-\frac{e^{-\frac{1}{\delta}} (\bar{u} - b)}{\delta (\bar{u} - \underline{u})} \right) \right)^{1/n}, \quad (19)$$

where W_{-1} is the branch -1 of the Lambert function.¹⁰ The support of this distribution is given by

$$[\underline{u}, \bar{u} - (1 - q)^n (\bar{u} - \underline{u}) (1 - n\delta \ln(1 - q))].$$

Theorem 1 (Separating Equilibrium).

The initial bid distribution $F_0(b)$ and the continuation distributions $F_t(b)$ and $G_t(b)$ of loser and winner respectively form a separating Markov equilibrium.

The equilibrium strategies have only been described on path so far. The description of the equilibrium in terms of the off path behavior is completed in the proof of Theorem 1. It uses the fact that the initial randomization of bids *in conjunction with* the unobservable nature of the bids extends the information contained in the on-path to the off-path bids. We conclude this section by commenting on uniqueness. Our separating equilibrium is the unique such equilibrium (there cannot be gaps, or atoms in the high-valuation bidder's distributions, by standard arguments), and as mentioned above, the limited monitoring gives each high-valuation bidder a strong incentive to bid more than \underline{u} . However, ruling out equilibria in which this is not the case appears daunting, because

¹⁰The Lambert W function is the inverse function of $f(x) = xe^x$. The function f is not injective. For $x \in \mathbb{R}$ the function is defined only for $x \geq -1/e$, and is double-valued on $(-1/e, 0)$. The alternate branch on $[-1/e, 0)$ with $x \leq -1$ is denoted $W_{-1}(x)$ and decreases from $W_{-1}(-1/e) = -1$ to $W_{-1}(0_-) = -\infty$.

any candidate equilibrium in which high-valuation bidders bid \underline{u} in some periods loses the property that uncertainty is resolved once the initial winner loses. Consequently, we would no longer be able to elude considering continuation games characterized by less tractable belief hierarchies. Currently, we are only able to show that the high-valuation bidder cannot be willing to bid \underline{u} indefinitely.

3.4 The Separating Equilibrium: Summary

Let us sum up the main findings that were either mentioned in passing, or that follow from the equilibrium solution described in Theorem 1. First, the bids in the initial period are lower than in the static first-price auction. This is puzzling at first glance. After all, for every possible bid b , the winner of the initial bidding game, has a higher continuation payoff than the loser. So it may seem that winning the initial auction is like winning the static auction, but with an additional prize provided by a more favorable continuation value. If the continuation value of winning or losing were independent of the current bid, then each bidder would bid more aggressively initially as it would look like the static auction, but with a prize larger than the flow payoff \bar{u} . But the analysis of the initial bid, given by (17), demonstrates that continuation values depend on the information provided by the initial bid. The payoff contribution from winning, $Y_1(b)$, is constant in b . Since the probability of winning is increasing in the bid, this implies that the continuation value $V_1(b)$, *conditional* on winning with b , is actually decreasing in b . By contrast, the payoff contribution $X_1(b)$ from losing is decreasing in b . Here, the continuation value $W_1(b)$ from losing at higher b is not increasing sufficiently fast to offset the lower probability of losing with a higher bid. Thus the initial bidding is less aggressive than in the first price-auction, and initial bids are depressed.

The winner's bid decreases over time, except at the very top, where it is constant. Figure 1 shows how bids decrease over time with two bidders. Bids b are on the abscissa, the probability $G_t(b)$ is on the ordinate. Higher curves correspond to later periods. That is, the probability assigned to the bid not exceeding a given value goes up over time, which means that over time bids go down.

The reason why the winner decreases his bid over time is clear: he cautiously explores how low

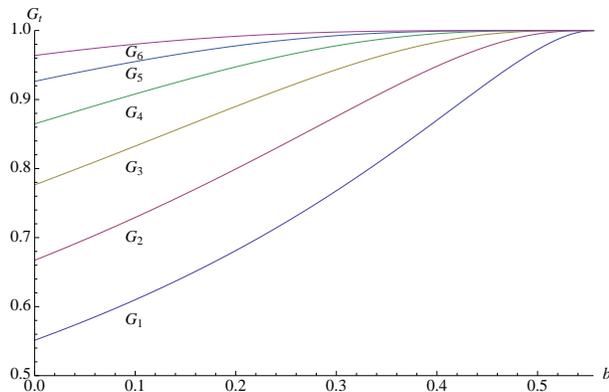


Figure 1: Bid distribution for $n = 2$ in periods $t = 1, \dots, 6$, $q = 1/3$, $\delta = 9/10$, $\bar{u} = 1 = 1 - \underline{u}$ (bottom $t = 1$, top $t = 6$)

he can get while still winning. At some point, he is sufficiently confident that his opponents have low valuations to submit a bid \underline{u}_+ , which conclusively establishes whether or not this is the case. Formally, for fixed δ , the distribution G_t converges pointwise to $\delta_{\underline{u}_+}$, the Dirac distribution that assigns probability one to \underline{u}_+ , as $t \rightarrow \infty$.

Given that the winner lowers his bid over time, the loser has no reason to raise his. Although he has always lost so far, which should push him towards higher bids, he knows that the winner is coming down with his bids, so that by not raising his bid, he will win his unit perhaps later, but at a lower cost. The equilibrium balances these forces, and a constant bid is best.

The total discounted revenue of this dynamic auction is close to, but strictly below, the theoretical maximum (in the absence of reserve prices) given by the static auction. To see this, note that the outcome is efficient (a low-valuation bidder never gets the unit if there is a high valuation present), and that the payoff of a high-valuation bidder can be computed by considering what happens if he always makes the highest bid. In that case, he will win all units, and the price he will pay for this is equal to $\bar{u} - (1 - q)^n(\bar{u} - \underline{u})$, as in the static auction, except in the initial period, where it is lower. How much lower depends on the discount factor: if the bidders are very impatient (δ near zero), then the initial bid distribution is close to the distribution in the static auction. This is not the case if they are very patient (δ near one), but in that case, the relative importance of the first period in the auctioneer's revenue is negligible (assuming that he shares the same discount factor).

Let us mention a few comparative statics, which follow from the characterization of the bid distribution. The results are proved in the supplementary Appendix. Regarding the discount factor, it is immediate to check that, for fixed, t , $\lim_{\delta \rightarrow 1} G_t(b) = F(b)$, which is equal to the distribution of the losers' bid, and is independent of t . More generally, for fixed t , the distribution G_t is decreasing in δ . That is, as is intuitive, the higher the discount factor, the slower the pace at which the winner lowers his bid over time. We summarize these observations in the next corollary.

Corollary 2 (Comparative Statics). *For every q, n, δ :*

1. *The bids (of high valuation bidders) jump up from $t = 0$ to $t = 1$. Thereafter, they are constant for losers, and decreasing over time for the winner.*
2. *For a fixed t , the losers' bids are independent of δ and the winner's bid is increasing in δ .*
3. *For a fixed t , bids increase with n , and tend to a two-point distribution on $\{\underline{u}, \bar{u}\}$ as $n \rightarrow \infty$.*
4. *Expected revenue tends to the revenue of the static auction as $\delta \rightarrow 1$.*

4 Observable Bids

We shall now turn to the case in which all bids are observable. As soon as a high-valuation bidder submits an equilibrium bid that is not in the low-valuation bidder's distribution support, the game simplifies. In this continuation game, if it is commonly known that two bidders have high valuations, then play is trivial by virtue of Refinement **A**. This occurs immediately, in particular, if the equilibrium is separating, that is, if bidders with different valuations use bidding distributions with different supports. We shall see, however, that such equilibria do not exist, unless players are sufficiently few, or high valuations sufficiently unlikely, as we establish in the next subsection. This provides a counterpoint to the non-existence of pooling equilibria when bids are unobservable.

4.1 On the Difficulty of Separating

An equilibrium is separating if the bid \underline{u} (which is the equilibrium bid of the low-valuation bidder) is not an equilibrium bid for the high-valuation bidders, so that, on path, all information is disclosed immediately. Note that this is a prerequisite for efficiency and revenue maximization.

Suppose then that a high-valuation bidder does not assign positive probability to the bid \underline{u} in the initial period. Obviously then, he will bid more. By bidding \underline{u}_+ , he gets $(1 - q)^n(\bar{u} - \underline{u})$. By deviating and bidding \underline{u} in this period, followed by \underline{u}_+ , a high-valuation bidder gets

$$(1 - \delta)(1 - q)^n(\bar{u} - \underline{u})/(n + 1) + \delta((1 - q)^n + (1 - \delta)nq(1 - q)^{n-1})(\bar{u} - \underline{u}). \quad (20)$$

To understand (20), note that, by bidding \underline{u} in the current period, he gets a (flow) reward only if all his opponents have low valuations as well, and even then, he wins only with probability $1/(n + 1)$ an object that is worth $\bar{u} - \underline{u}$ to him. This is the first term. In the following period, however, he will be believed to be a low-valuation bidder, and this allows him to win one unit at a price arbitrarily close to \underline{u} provided that there are not two or more high-valuation bidders, hence the second term.

On the other hand, separation yields the same payoff as the static auction, $(1 - q)^n(\bar{u} - \underline{u})$. This is because the continuation payoff of a high-valuation bidder is independent of the specific bid above \underline{u} that he submits, so that he has the same incentives as in the static auction, which gives him a payoff $(1 - \delta)(1 - q)^n(\bar{u} - \underline{u})$. He further gets $\delta(1 - q)^n(\bar{u} - \underline{u})$ from the second period onward if it turns out that all other bidders have low valuations.

Proposition 2 (Difficulty of Separating).

For all positive q and δ , there exists \bar{n} such that for all $n > \bar{n}$, a separating Markov equilibrium does not exist with observable bids.

Comparing the two payoffs, we find that separation is *not* an equilibrium if, and only if

$$q \geq \underline{q} := \frac{1}{1 + (n + 1)\delta}. \quad (21)$$

This condition, expressed in terms of the prior probability of a high valuation is satisfied if there are sufficiently many bidders and/or if the discount factor is sufficiently high. Proposition 2 then has to be contrasted with Proposition 1, where we showed that with unobservable bids pooling is never an equilibrium, and constructed a separating equilibrium for all possible values of q , δ and n .

What then is the equilibrium of the game? As will be the case, it might occur that only one bidder reveals himself to have a high valuation, while all the other n bidders submit a bid \underline{u} , pooling thereby with the low-valuation type. As always, we must first understand this continuation game of one-sided incomplete information before we can solve the original game. We turn now to the game of one-sided incomplete information, as a preamble to the general analysis.

4.2 The Game of One-Sided Incomplete Information

We consider here the game in which one bidder, say bidder 1, is commonly known to have a high valuation, while each of the other n bidders is believed to have a high valuation with probability q . Accordingly, bidder 1 is uninformed, while all other bidders are informed.

Let F_t^U denote the bid distribution of the uninformed player, and F_t^I the common bid (unconditional) distribution of the other players. This represents a slight abuse of notation, as in a Markov equilibrium, the state variable, *i.e.* the beliefs, determine the distributions, not the period. But as there is a one-to-one correspondence between time and beliefs on the equilibrium path, using time as an index facilitates the exposition.¹¹ Because any equilibrium bid different from \underline{u} by an informed bidder establishes common knowledge that there is a second high-valuation bidder, after which the game becomes trivial, we must only understand how play proceed along histories in which all informed bidders have bid \underline{u} in every period so far. Let q_t denote the probability that (any of) the informed bidder's valuation is high at the beginning of period t , given any history on the equilibrium path in which all informed bidders bid \underline{u} in all periods up to (and including) $t - 1$, and set $q_0 := q$. Note that the past bids of the uninformed bidder do not affect this belief, so that this is really a

¹¹To prevent any confusion, we avoid the notation F, G introduced in the unobservable case.

function of time only. Let $T \in \mathbb{N}_0 \cup \{+\infty\}$ denote the length of the longest such history when at least one informed bidder has a high value; i.e., there exists no history on the equilibrium path in which all high-value informed bidders submit bids equal to \underline{u} for $t > T$ periods.¹² Finally, let \bar{b}_t denote the highest bid in the bid support in period $t \leq T$ (conditional on a history in which all informed bidders bid \underline{u} throughout, a qualification we shall omit from now on). While this is a private-values setting, this game and its solution bear strong similarities with the common-values game of Hörner and Jamison (2008). Thus we keep the analysis concise.

Because the informed bidder always bids at least \underline{u}_+ (he has nothing to lose from doing so, given the Markov assumption), a high-valuation informed bidder will not submit the bid \underline{u} forever (as he would then lose forever). Hence $T < \infty$. By standard arguments, the informed bidders must randomize in period t between the bid \underline{u} (at least as long as $t < T$) and mixing over the interval $(\underline{u}, \bar{b}_t)$, for some $\bar{b}_t > \underline{u}$. Further, the uninformed bidder must bid \underline{u}_+ with positive probability, for otherwise the low-valuation bidder would be unwilling to submit bids arbitrarily close to, but above \underline{u} (such a revelation would be in vain, as it would yield a zero reward and continuation payoff). Because the low-valuation bidder bids \underline{u} for sure, Bayes' rule tells us that the probability that an informed bidder is of the low type in period $t + 1$, given such a history, is given by

$$1 - q_{t+1} = \frac{1 - q_t}{F_t^I(\underline{u})}, \quad q_0 = q, q_{T+1} = 0. \quad (22)$$

Because the reward from every bid in the interval $b \in (\underline{u}, \bar{b}_t]$ must be the same (for the informed bidder, this is because the continuation payoff is then 0; for the uninformed bidder, this follows from the Markov assumption), we have, for all $t \leq T$,

$$F_t^U(b)F_t^I(b)^{n-1}(\bar{u} - b) = \bar{u} - \bar{b}_t = F_t^I(b)^n(\bar{u} - b), \quad (23)$$

and so $F_t^U(b) = F_t^I(b) =: F_t(b)$ for all b, t , where $F_t^U(\underline{u})$ is the probability assigned by the uninformed

¹²If $T = +\infty$, there exists arbitrarily long such histories. If $T < +\infty$, then we must have $F_T^I(\underline{u}) = 1 - q_T$, since the high-valuation bidder must bid strictly more than \underline{u} in that period. Hence, $q_{T+1} = 0$.

bidder to \underline{u}_+ . Finally, because an informed bidder is indifferent between bidding just above \underline{u} in period $t < T$ and bidding \underline{u} followed (if no informed bidder bid more than \underline{u} in period t) by a bid just above \underline{u} in period $t + 1$, we have, for all $t < T$,

$$F_t^n(\underline{u}) = \delta F_t^{n-1}(\underline{u}) F_{t+1}^n(\underline{u}), \text{ or } F_t(\underline{u}) = \delta F_{t+1}^n(\underline{u}). \quad (24)$$

The equality must be replaced by an inequality in period T , *i.e.*, $F_T(\underline{u}) \geq \delta$. Equations (22)–(24) then allow us to solve for the equilibrium.¹³ Given a prior belief $q < 1$, it follows from (24) that, as $\delta \rightarrow 1$, $T \rightarrow \infty$ and $\delta^T \rightarrow 1$ for $T = T(q)$. That is, uncertainty is resolved arbitrarily fast relative to δ , although the time it might take (in fact, the expected time it takes) grows without bound. We can summarize our findings as follows.

Lemma 3 (Bidding with One-Sided Information and Observable Bids).

1. *The equilibrium bid distribution $F_t(b)$ is increasing in t for all b (*i.e.*, the uninformed bidder's bids decrease on average).*
2. *(At least) one informed high-valuation bidder reveals his type by period $T < \infty$, where $\lim_{\delta \rightarrow 1} T = \infty$, yet $\lim_{\delta \rightarrow 1} \delta^T = 1$.*

While this equilibrium exhibits interesting features, only the resulting payoffs matter for the analysis in the game in which all players are symmetrically informed. The payoff of a high-valuation informed bidder can be computed from the strategy of always bidding \underline{u} until period $T - 1$, and \underline{u}_+ in period T (assuming no informed bidder bid more than \underline{u} until then). The payoff of this is then

$$V^I(q) := (1 - \delta)\delta^T F_T(\underline{u})^{n^{T+1}} (\bar{u} - \underline{u}) \rightarrow 0$$

¹³To see this, note that by “telescoping” (22) for $t = 0, \dots, T$, we get $1 - q = \prod_{t=0}^T F_t(\underline{u})$, which combined with (24) gives $F_T(\underline{u}) = 1 - q_T$ as a function of T and q . The requirement that $F_{T-1}(\underline{u}) \leq \delta \leq F_T(\underline{u})$ yields for $n > 1$ $T = \max\{\tau : \delta^{\frac{n^{\tau+1}-1}{(n-1)^2} - \frac{\tau+1}{n-1}} \geq 1 - q\}$, (for $n = 1$, following the same steps, $T = \max\{\tau : \delta^{\frac{\tau(\tau+1)}{2}} \geq 1 - q\}$) which pins down q_T . (Note that $T \rightarrow \infty$, $\delta^T \rightarrow 1$ as $\delta \rightarrow 1$ follows.) Eq. (23) for $b = \underline{u}$ then gives \bar{b}_t and so $F_t(b)$.

as $\delta \rightarrow 1$. Writing out the payoff V^U of the uninformed player is a little messier, and so is omitted here. (See Online Appendix A for details, including an asymptotic expansion as $\delta \rightarrow 1$.) Note that, from period T onward, in case no informed bidder separated so far, his continuation payoff is $F_T(\underline{u})^n(\bar{u} - \underline{u})$. Since $\delta^T \rightarrow 1$ as $\delta \rightarrow 1$, it follows that $V^U(q) \rightarrow (1 - q)^n(\bar{u} - \underline{u})$: asymptotically, the payoff to the uninformed agent is only positive in the event in which all his opponents have low valuations. This is good for the auctioneer in this game, who thus gets (asymptotically) a maximal revenue.

4.3 The Game of Symmetric Information

Building on these findings, we may now return to the game with symmetric information. In particular, when is pooling an equilibrium outcome when all bidders' valuations are equally unknown?

A pooling equilibrium must involve all bidders submitting the bid \underline{u} (given Refinement **A**). The payoff to a high-valuation bidder is then $(\bar{u} - \underline{u})/(n + 1)$. The best deviation for a high-valuation bidder involves bidding \underline{u}_+ , which garners $(1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q)(\bar{u} - \underline{u})$, assuming that such a deviation is ascribed to a high-valuation player, so that the game with one-sided incomplete information ensues. Therefore, pooling is an equilibrium if and only if

$$\frac{1}{n + 1} \geq 1 - \delta + \delta V^U(q)/(\bar{u} - \underline{u}). \quad (25)$$

Because V^U is monotonically decreasing in q , and bounded above by $\bar{u} - \underline{u}$, this gives a lower bound \bar{q}° to the values of q for which such an equilibrium exists, and it is easy to see that $\bar{q}^\circ > \underline{q}^\circ$.

Proposition 3 (Possibility of Pooling).

For all positive q , there exists $(\bar{\delta}, \bar{n})$ such that for all $\delta > \bar{\delta}$ and $n > \bar{n}$, a pooling Markov equilibrium does exist with observable bids.

Let us now focus on $\delta \rightarrow 1$. Because $V^U \rightarrow (1 - q)^n(\bar{u} - \underline{u})$, we get

$$q \geq \bar{q}^o \rightarrow 1 - (n + 1)^{-1/n}. \quad (26)$$

The left-hand side of (25) tends to what a high-valuation bidder can secure, if low-valuation bidders do not bid more than \underline{u} . Indeed, a high-valuation bidder can secure $(1 - q)^n(\bar{u} - \underline{u})$ by always bidding \underline{u}_+ . Thus, the pooling equilibrium exists whenever it yields an individually rational payoff to the high-valuation bidder. As we shall see, the same holds when only the winner's bid is observed.

Recall that a separating equilibrium exists whenever $q \leq \underline{q}^o$. This leaves us with the (non-empty) interval $(\underline{q}^o, \bar{q}^o)$. In that case, the equilibrium must involve semi-pooling. That is, the high-valuation bidder puts positive probability on \underline{u} , but he also continuously randomizes over some interval of higher bids. In the event that all realized bids are \underline{u} , bidders assign a growing probability to the event that their opponents have a low valuation, so that, at some point, beliefs are such that a separating equilibrium exists. In the supplementary Appendix, we show that such (not necessarily unique) semi-pooling equilibria exist in this intermediate range of values for q , that a semi-pooling equilibrium cannot end up in pooling in finite time, and that no such equilibrium can exist if $q \leq \underline{q}^o$, *i.e.*, it cannot exist when a separating equilibrium exists. Semi-pooling equilibria may also exist for $q \geq \bar{q}^o$ (that is, they co-exist with pooling equilibria); in fact, we prove in the Online Appendix A that they do for any such prior when δ is sufficiently close to one, and that their revenue converges to the revenue from the separating equilibrium (as $\delta \rightarrow 1$). We summarize this discussion in the following theorem, see also Figure 2.

Theorem 3. *A (Markov) equilibrium always exists. Furthermore, if*

1. $q \in [0, \underline{q}^o]$, *the unique equilibrium is separating;*
2. $q \in (\underline{q}^o, 1]$, *no separating equilibrium exists. Furthermore, if*
 - (a) $q \in (\underline{q}^o, \bar{q}^o]$, *all equilibria are semi-pooling;*

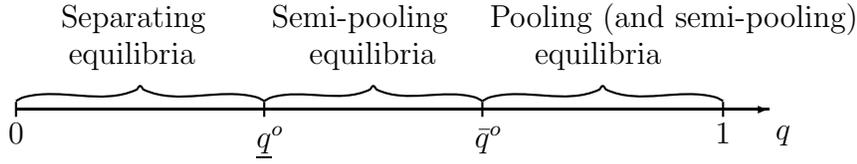


Figure 2: Markov equilibria in the observable case, as a function of q .

(b) $q \in (\bar{q}^o, 1]$, a pooling equilibrium exists.

To summarize, as $\delta \rightarrow 1$, the revenue converges to the optimal revenue (without reserve price) if and only if $q < 1 - (n + 1)^{-1/n}$, a decreasing function of n (that tends to 0 as $n \rightarrow \infty$). Otherwise, because bidders all use the same low bid, the auctioneer's revenue is equal to \underline{u} .

5 Winner-Only Observable Bids

Finally, we consider the case in which the bid and the identity of the winner are disclosed after each auction.¹⁴ Again, in order to solve the game in which bidders have symmetric information, we start with the case in which exactly one bidder is known to have a high valuation, because such informational structures arise as continuations of the game with symmetric information.

5.1 The Game of One-Sided Incomplete Information

We already analyzed the game with one-sided incomplete information in the environment with observable bids. There, every bid by the informed bidder strictly above \underline{u} revealed that the bidder has a high valuation. Thus, the evolution of the posterior belief had a simple binary structure. Either the bids of the informed agent were all equal to \underline{u} , and then the posterior declined from q_t to q_{t+1} , or at least one of the informed bids was above \underline{u} , and then uncertainty was resolved. In the current

¹⁴While we assume so, it is not necessary that the winner's identity be disclosed: if it is not, then a bidder who has not won might not know whether the winning bids were submitted by the same, or different bidders; but if two different bidders won with bids above \underline{u} , it would still be the case that the two of them would commonly know that there are two high-valuation bidders, and bids would then be \bar{u} from that point on.

environment where only the winning bid is observable, the updating process depends on the realized bid of the uninformed bidder. His bid, as long as it is winning, provides an upper bound for losing bids. Therefore, the level of his bid determines the rate at which updating occurs. In contrast to the observable case, it is therefore convenient to describe the strategies in terms of the commonly known uninformed bidder's posterior belief about the informed bidder's valuation.

So suppose that a player, bidder 1, say, is known to have valuation \bar{u} . His opponents, on the other hand, have privately known, and independently drawn valuations. The probability with which each of these bidders has a high valuation is denoted q , as before. Bidder 1 is the uninformed bidder, while the other bidders are informed. We begin with a few observations.

Because bidder 1's valuation is known, he has nothing to lose from breaking any tie in his favor. By winning, he not only gets an immediate reward, but he also increases the probability that the game "goes on" (if another bidder outbids him, then it is known that there are two high-valuation bidders, and so their bids are \bar{u} from then on). Thus, the uninformed bidder bids at least \underline{u}_+ .

Any informed high-valuation bidder has no incentive to bid \underline{u} either. Such a bidder will never be able to win more than one unit, because, by the previous observation, he needs to bid at least \underline{u}_+ to win, and doing so would reveal his valuation as well. So the best he can hope for is winning at a price arbitrarily close to, but above \underline{u} . So bidding, say, $\underline{u} + \varepsilon$, for $\varepsilon > 0$ small enough, does strictly better than bidding \underline{u} . That is, denoting by F_q^I the common informed bidder's distribution, given that each of them is believed to have a high valuation with probability q , it holds that $F_q^I(\underline{u}) = 1 - q$.

The bids of the uninformed bidder affect how much he learns about the informed bidders: he is more likely to win with a higher bid, but such a bid is less informative about the probability that at least one of the informed bidders has a high valuation. His bid distribution F_q^U is therefore also indexed by the belief q . The uninformed bidder must randomize over some interval $[\underline{u}_+, \bar{b}_q]$, for some bid \bar{b}_q , given the common belief q . If he did play a pure strategy, the informed bidders would outbid him by a very small amount, so that this could not be optimal. He cannot be the only player submitting bids in this range, so the high-valuation informed bidder must do so as well.

We are now ready to solve for the equilibrium. Observe that by bidding \bar{b}_q , the uninformed bidder

prevents any learning, since he wins for sure and only his own bid is observed. While in equilibrium he randomizes in every period, one optimal strategy consists in making this same bid forever. At the opposite end, by bidding \underline{u}_+ , he ensures that he learns perfectly his opponent's type, since any informed bidder with a high valuation bids strictly more. Hence, denoting by $V^U(q)$ the uninformed bidder's payoff given belief q , we have, for all q ,

$$V^U(q) = (1 - q)^n(\bar{u} - \underline{u}) = \bar{u} - \bar{b}_q. \quad (27)$$

Note that, by Bayes' rule, the probability that the uninformed bidder assigns to any of his opponents having a low valuation, conditional on winning with a bid b , is given by $(1 - q)/F_q^I(b)$, where q is his prior belief. Therefore, we must have, more generally,

$$(1 - q)^n(\bar{u} - \underline{u}) = F_q^I(b)^n(1 - \delta)(\bar{u} - b) + \delta(1 - q)^n(\bar{u} - \underline{u}),$$

where we use the first equality of (27) to eliminate V^U . It follows that

$$F_q^I(b) = (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}. \quad (28)$$

Therefore, any informed bidder must bid as in the static auction (and as the losers in the repeated auction with unobservable bids). Let us turn now to the informed bidders' problem. By bidding \bar{b}_q , such a bidder wins once, and "ends" the game. So his payoff $V^I(q)$ must satisfy

$$V^I(q) = (1 - \delta)(\bar{u} - \bar{b}_q) = (1 - \delta)(1 - q)^n(\bar{u} - \underline{u}), \quad (29)$$

where we have used (27) to obtain the second equality. More generally, his payoff from bidding b consists of two terms. With probability $F_q^U(b)F_q^I(b)^{n-1}$, he is the highest bidder and wins. If he loses, then he only gets a positive continuation payoff if the uninformed bidder wins. That is, the

uninformed bidder must bid some $\beta > b$, and all other informed bidders must bid less than β . Thus,

$$V^I(q) = (1 - \delta)(1 - q)^{n-1} (\bar{u} - \underline{u})^{\frac{n-1}{n}} \left(F_q^U(b) (\bar{u} - b)^{\frac{1}{n}} + \delta \int_b^{\bar{b}_q} (\bar{u} - \beta)^{\frac{1}{n}} dF_q^U(\beta) \right),$$

where we use (28) and (29). Plugging into (29) gives, for all b in the support of F_q^I ,

$$(1 - q)(\bar{u} - \underline{u})^{\frac{1}{n}} = F_q^U(b) (\bar{u} - b)^{\frac{1}{n}} + \delta \int_b^{\bar{b}_q} (\bar{u} - \beta)^{\frac{1}{n}} dF_q^U(\beta).$$

Because this is an identity with respect to b , and because the first and last terms are differentiable in b , the second term must be as well. Taking derivatives yields

$$\frac{dF_q^U(b)/db}{F_q^U(b)} = \frac{1}{n(1 - \delta)} \frac{1}{\bar{u} - b}.$$

We can integrate, and use that $F_q^U(\bar{b}_q) = 1$ (where \bar{b}_q is determined by (27)), to get

$$F_q^U(b) = \left((1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{\frac{1}{n}} \right)^{\frac{1}{1-\delta}}, \quad (30)$$

where as usual $F_q^U(\underline{u}) = (1 - q)^{\frac{1}{1-\delta}}$ is the probability attached to the bid \underline{u}_+ . The uninformed bidder bids more aggressively than the informed bidders, as winning carries the additional benefit of prolonging the game. This benefit becomes predominant as the discount factor is high, and the uninformed bidder is then likely to make a bid that is near the upper end of the support. Equations (28) and (30) characterize the equilibrium of the game with one-sided incomplete information, with payoffs given by (27) and (29). We summarize this discussion in the next proposition.

Proposition 4 (Bidding with One-Sided Incomplete Information and Winner-Only Observable Bids).

The game with one-sided incomplete information admits a unique equilibrium, in which:

1. *The informed bidders bid as in the static game, while the uninformed bidder bids more aggres-*

sively.

2. *The uninformed bidder's bids decrease (weakly) over time. The time until the uninformed bidder loses is finite (a.s.), but admits no finite bound.*

It is informative to contrast the bidding behavior in the observable and the winner-only observable environment. Intuitively, in the observable bid regime, the informed bidders should bid more cautiously as any bid above \underline{u} will reveal that the bidder has a high valuation. The informed bidder is therefore concerned that a bid higher than \underline{u} would reveal his valuation without necessarily winning the object. The bidding strategy of an informed bidder in the observable environment is less aggressive. In particular, from (22), with observable bids, the probability of a low bid is

$$F_t^I(\underline{u}) = \frac{1 - q_t}{1 - q_{t+1}} > 1 - q_t,$$

which is larger than in the winner-only observable bid environment, in which it is given by $F_t^I(\underline{u}) = 1 - q_t$. In fact, given the characterization of the bidding distribution given by (23) and (28), it follows that the upper bound of the bids is lower, and more generally that the bid distribution with observable bids is first-order stochastically dominated by the bid distribution with winner-only observable bids. Given the more defensive posture of the informed bidders in the observable environment, it then follows that the uninformed bidder is also bidding less aggressively. In fact, the bid distribution of the uninformed bidder in the observable environment is also first-order stochastically dominated by the bid distribution with winner-only observable bids, as can be directly inferred by comparing (23) with (30). As a result, the revenue is higher in the winner-only observable case.¹⁵

¹⁵This is immediate from the expected payoffs, given that both auctions are efficient. V^U achieves the obvious lower bound on the uninformed bidder's payoff in the winner-only observable case (see (27)), so it is lower than in the observable case. Similarly, the informed bidder's payoff is lower in the winner-only observable case (compare the upper end of the bid supports in the initial period, which determine the informed bidder's revenue).

5.2 The Game of Symmetric Information

We start with the necessary and sufficient conditions for a pooling equilibrium to exist. Assume all bidders submit the bid \underline{u} with probability one. By doing so, a high-valuation bidder gets $(\bar{u} - \underline{u})/(n + 1)$. By bidding \underline{u}_+ instead, which is the best alternative, he gets

$$(1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q) = (1 - \delta)(\bar{u} - \underline{u}) + \delta(1 - q)^n(\bar{u} - \underline{u}), \quad (31)$$

where the equality uses (27). We can then establish the following result.

Proposition 5 (Possibility of Pooling).

For all positive q , there exists $(\bar{\delta}, \bar{n})$ such that for all $\delta > \bar{\delta}$ and $n > \bar{n}$, a pooling Markov equilibrium does exist with winner-only observable bids.

By comparing the two payoff streams above, we see that pooling is an equilibrium whenever

$$\frac{1}{n + 1} \geq 1 - \delta + \delta(1 - q)^n, \quad (32)$$

i.e. the belief q should exceed some threshold q^w . It is easy to see that $q^w > \bar{q}^o$.¹⁶ Note that, as $\delta \rightarrow 1$,

$$q^w \rightarrow 1 - (n + 1)^{-1/n}, \quad (33)$$

the same threshold as with observable bids.

Consider now a separating equilibrium, *i.e.* assume that, in the initial period, all high-valuation bidders submit bids that are at least \underline{u}_+ . The bid distribution F_q on $[\underline{u}, \bar{b}_q]$ must then satisfy

$$F_q(b)^n \left((1 - \delta)(\bar{u} - b) + \delta V^U \left(1 - \frac{1 - q}{F_q(b)} \right) \right) + \delta \int_b^{\bar{b}} V^I \left(1 - \frac{1 - q}{F_q(\beta)} \right) n F(\beta)^{n-1} dF_q(\beta) = (1 - \delta)(\bar{u} - \bar{b}) + \delta V^U(q),$$

¹⁶This is because the uninformed bidder's payoff in the game with observable bids, V^U , exceeds $(1 - q)^n$ (a lower bound to his payoff). Because V^U is decreasing in q , and (32) and (25) only differ by a term $((1 - q)^n$ replacing $V^U(q)$ in the winner-only observable case), the result follows.

where the two terms of the left-hand side are the payoffs from winning and losing with a bid $b > \underline{u}$, and the right-hand side is this payoff computed for $b = \bar{b}$. Simplifying (using (27) and (29)) gives

$$F_q(b)^n(\bar{u} - b) + \delta n(1 - q)^n(\bar{u} - \underline{u}) \int_b^{\bar{b}} \frac{dF_q(\beta)}{F_q(\beta)} = \bar{u} - \bar{b}.$$

Since $dF_q(\beta)/F_q(\beta) = d \ln F_q(\beta)$, we obtain, using $F_q(\underline{u}) = 1 - q$,

$$\frac{\bar{u} - b}{\bar{u} - \underline{u}} \left(\frac{F_q(b)}{1 - q} \right)^n = 1 + \delta \ln \left(\frac{F_q(b)}{1 - q} \right)^n, \quad (34)$$

Equation (34) determines F_q . A closed-form solution is given in terms of the (branch -1 of the) Lambert function W_{-1} , *i.e.*

$$F_q(b) = (1 - q) \left(-\frac{\delta(\bar{u} - \underline{u})}{\bar{u} - b} W_{-1} \left(-\frac{e^{-1/\delta}(\bar{u} - b)}{\delta(\bar{u} - \underline{u})} \right) \right)^{1/n},$$

and the largest bid is $\bar{b}_q = \bar{u} - (1 - q)^n(1 - \delta n \ln(1 - q))(\bar{u} - \underline{u})$. Note that this initial bid distribution is the same as in the unobservable case! In fact, the continuation payoffs from winning and from losing are identical as well. In particular, the continuation payoff from winning is independent of the level of the winning bid in the winner-observable environment, as it is the unobservable environment. It also means that a high-valuation bidder makes the same overall payment for winning in all periods by submitting the highest bid in the bidding support in all periods. Note also that bidding supports coincide for later periods. Thus, his payoff, and given that the equilibrium is efficient, the revenue, is the same in both separating equilibria. The similarities in the payoffs (and the strategies) in the winner-only observable bid and the unobservable bid environment can be traced back to the similarities in the players' information. In either game form, the winning seller does not know whether his competitors have low or high valuation, and he does not know their past bids either. The losing bidders in either game form know that they face a high-value bidder. In addition, in the winner-only observable environment they know the winner's past bids. In either game form, the losing bidders have the same opportunities, namely to win once against the winning bidder.

In fact, a stronger equivalence holds: the unconditional distribution of bids is the same in the unobservable case as in the separating equilibrium of the winner-only observable case.¹⁷ We should however emphasize that the equilibrium play across the two environments is quite distinct. In the unobservable environment, the bidders randomize only initially and then pursue a deterministic policy as a function of their initial randomized choice, whereas in the winner-only observable environment bidders keep on randomizing in all periods, until all uncertainty is resolved.

Note that the separating equilibrium exists for all parameters. If the other high-valuation bidders separate, it is strictly better for a high-valuation bidder to bid \underline{u}_+ rather than \underline{u} . In either case, the bidder only wins if all the opponents have low valuation. If they do not, both bids are equivalent. If they all do, then bidding \underline{u}_+ is strictly better than \underline{u} , for by making the former bid, the bidder learns that all other bidders have low valuation, and so he always wins at the price \underline{u}_+ .

If the prior belief is high enough for pooling equilibria to exist, semi-pooling equilibria also exist, in which high-valuation bidders assign positive probability to \underline{u} , but also continuously randomize over some range $(\underline{u}, \bar{b}_q)$. The Online Appendix A provides an exhaustive analysis of such equilibria.¹⁸ Such equilibria must converge (in finite time or not) to pooling, so that the belief must be above the threshold ensuring existence of pooling equilibria (characterized by (32)), and their revenue is bounded away from the revenue of the separating equilibrium. The next theorem summarizes these findings, as does Figure 3.

Theorem 4. *A (Markov) equilibrium always exists. Furthermore,*

1. *A unique separating equilibrium always exists. If $q < q^w$, this is the unique equilibrium;*
2. *If $q > q^w$, a (unique) pooling equilibrium exists, as well as semi-pooling equilibria.*

¹⁷This follows by computation. The unconditional distribution in the winner-only observable case in a given period is a convolution over those distributions in earlier periods, as earlier bids affect beliefs. Details available upon request.

¹⁸For each prior $q \geq q^w$, countably many semi-pooling equilibria exist. Exactly one such equilibrium involves an infinite sequence of distinct posterior beliefs (conditional on the highest bid being \underline{u}), which converges to q^w ; all other semi-pooling equilibria specify switching at an arbitrary time from this sequence to the belief q^w at which pooling ensues. The later the switch, the lower the revenue.

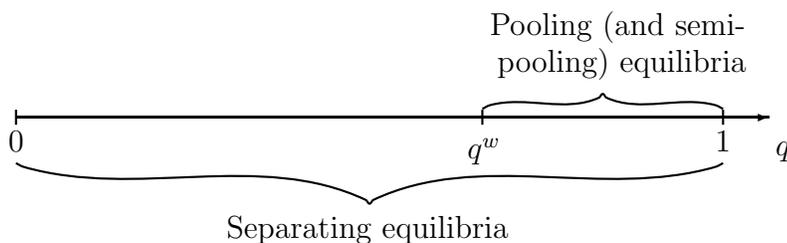


Figure 3: Markov equilibria in the winner-only observable case, as a function of q .

The informational environment of the winner-only observable bids leads to qualitative characteristics of the equilibria *between* the unobservable and the observable environment. As in the unobservable environment, a separating equilibrium is guaranteed to exist for all values of q , δ , and n . But as in the observable environment, a pooling equilibrium also exists for every prior probability q as long as δ and n are sufficiently large.

To conclude, the revenue comparison between observable bids and winner-only observable bids is not clear-cut, shrouded in part by the multiplicity of equilibria. There is a sense in which winner-only observable auctions make bidders more aggressive. To wit, the threshold above which pooling exists is higher ($q^w > \bar{q}^o$), and a separating equilibrium always exists, unlike in the observable case. In addition, the comparison is unambiguous when information is one-sided. But this is precisely what mitigates aggressive bidding in the winner-only observable case: true, bidding aggressively and failing to win does not lead to costly information revelation in that case, but if revelation *does* occur, more aggressive bidding in the one-sided case ensues, which holds up aggressive bidding in the first place. The success of the winner-only observable format in raising revenue, and so depressing payoffs in the one-sided case backfires, putting a damper on bidding when information is still symmetric.

6 Discussion and Conclusion

The objective of the present paper was to investigate the role of disclosure policies in the context of repeated auctions and procurements. We compared the minimal disclosure policy, namely unob-

servable bids, with the maximal disclosure policy, namely observable bids. We also considered an intermediate form of disclosure policy, in which only the winning bid was disclosed. We now discuss the robustness of the analysis (and its result) with respect to a number of natural generalizations and variations of the stylized model.

Finite vs. Infinite Horizon: So far, we have analyzed the bidding behavior for all disclosure environments in an infinite horizon setting. It is then natural to ask whether we might not have obtained similar results in a finite horizon setting, perhaps even in a two period setting with considerably less effort. We thus briefly discuss how the equilibrium analysis would be affected by imposing a finite horizon (Details in a Supplementary Appendix B available from the authors). The resulting comparison is also informative as much of the literature discussed in the introduction obtains results for the special case of a finite horizon with two periods. To anticipate the discussion below, we find that the assumption of finite horizon crucially misses the possibility of pooling (or even semi-pooling) equilibria that arise naturally under infinite horizon, but vanish under finite horizon. As we established that the pooling equilibria only arise in the observable or the winner-only observable bid environment our conclusion regarding the benefits of disclosing bid information therefore differ markedly from the previous literature, and lead us to suggest that disclosing past bids may be detrimental to the revenue to the seller.

We begin our discussion with the case of unobservable bids. Here we derived the Markov equilibrium in the infinite horizon model. The Markov equilibrium for a finite horizon (with or without discounting) is in fact simply the truncated version of the infinite horizon model and all the results, in particular Proposition 1 on the impossibility of pooling and Theorem 1 on the characterization of the separating equilibrium remain valid in their exact form.

However, the restriction to a finite horizon does impact the bidding behavior with observable bids. It remains true that is difficult to obtain a separating equilibrium, and the critical threshold \underline{q}^o remains valid, but importantly a (complete) pooling equilibrium as in Theorem 3 now fails to exist for any finite horizon. To see this, suppose in the penultimate period both types were to bid \underline{u} . Then

the beliefs would not be updated at all, remaining at the common prior q . Now if a high value bidder were to deviate to bid \underline{u}_+ , then he would win the object for sure while beliefs about all other bidders would remain at the *ex ante* prior q . As we have already shown, a static auction with such one-sided information involves the same unconditional distribution of bids and same expected payoff as the symmetric information one. So deviating is unambiguously a better strategy and therefore imply the impossibility of pooling. Thus any finite horizon model would neglect the existence of pooling equilibria, and thus would fail to recognize the revenue gain from unobservable bids.

With observable bids by the winner, the impossibility to sustain pooling equilibria remains, and thus Proposition 5 ceases to hold. In fact, even the semi-pooling equilibria cease to exist when only the winning bid is observable. The difference to the observable bid environment is that now the losers' bids are private, and thus bidding \underline{u}_+ does not risk to reveal one's type when another high type bidder made the winning bid. This further undermines the ability of bidders to disguise their valuation in a semi-pooling equilibrium. Hence the only remaining equilibrium is the separating equilibrium.

The Role for Reserve Prices: We have argued that the disclosure environment impacts the revenue of the seller, and by consequence the net utility of the bidder in an important way. But information is only one of many instruments available to the seller to increase his revenue, and thus we might wonder whether a more traditional instrument such as a reserve price would not be a more powerful instrument than the disclosure policy. We therefore briefly analyze how the use of reserve prices would impact the analysis of the disclosure regimes. For the purpose of the discussion here, we shall restrict our attention to optimally chosen *stationary* reserve prices, and thus will not consider reserve prices that can change over time.

In a static setting, we know that the revenue of a standard first price auction is generally not optimal, and can be increased by suitably chosen reserve prices. In the current binary setting only a reserve price equal to the high valuation can possibly increase the revenue. Indeed with a high reserve price, the seller can extract all the surplus from the bidders but faces the risk that none of

them meets the reserve price. By comparing the revenue of the auction format with and without a reserve price, we can easily establish that it is optimal to impose a high reserve price if and only if

$$q^r \geq \frac{\underline{u}}{(n+1)\bar{u} - n\underline{u}}. \quad (35)$$

After all, we established in Section 2.4 that the revenue without a reserve price equals the revenue of the second price auction. Now, the above inequality is more likely to get satisfied when (a) q is higher, (b) n is larger, and (c) the difference $\bar{u} - \underline{u}$ is larger.

With this benchmark from the static setting, we can now understand the impact of the reserve prices in the dynamic setting. We shall restrict our attention to the limiting case as $\delta \rightarrow 1$ to be able to directly use our previous results. By Corollary 2, we then know that with unobservable bids, the expected revenue tends to the revenue of the static auction. By contrast, with observable bids, we know from Theorem 3 that a pooling equilibrium exists, and from condition (26) we know that as $\delta \rightarrow 1$, the critical threshold converges to:

$$\bar{q}^o = 1 - (n+1)^{-1/n}. \quad (36)$$

Moreover, Proposition 5 establishes for winner-only observable bids that as $\delta \rightarrow 1$, the threshold for the pooling equilibrium q^w converges to \bar{q}^o as well, see (33). Thus, we can conclude that for all q such that $\bar{q}^o < q < q^r$, the optimal policy would be to *not* impose reserve prices and to *not* disclose any bidding information. After all, the flow revenue in any pooling equilibrium is limited to the low valuation \underline{u} . By contrast in the unobservable bid environment, it is either the expected value of the second order statistic, which is strictly larger than \underline{u} . A comparison between the two critical thresholds q^r and \bar{q}^o establishes that there is always a range of values $[\underline{u}, \bar{u}]$ such that the imposition of high reserve price \bar{u} is not optimal for a range of prior probabilities q .

Proposition 6 (Irrelevance of Reserve Prices).

Given any number of bidders $n > 1$, there exists $\eta > 0$, such that, for all $\underline{u}/\bar{u} \in (1 - \eta, 1)$, $\bar{q}^o < q^r$.

Hence, for any prior $q \in [\bar{q}^o, q^r]$, reserve prices do not increase revenue.

Proof. We compare the two thresholds q^r and \bar{q}^o given by (35) and (36). We observe that q^r is increasing in \underline{u} and \bar{q}^o is independent of \underline{u} . We solve for

$$1 - (n + 1)^{-1/n} = \frac{\underline{u}}{(n + 1)\bar{u} - n\underline{u}}$$

or

$$\frac{\underline{u}}{\bar{u}} = \frac{(n + 1) \left((n + 1)^{\frac{1}{n}} - 1 \right)}{(n + 1) \left((n + 1)^{\frac{1}{n}} - 1 \right) + 1},$$

and the right-hand side is evidently smaller than 1 for all $n > 1$. □

The ratio of \underline{u} and \bar{u} appears as it represents the relative loss the seller incurs by relying on a competitive bid from the low valuation bidder \underline{u} rather than imposing a high reserve price equal to the high valuation \bar{u} . Clearly for $\underline{u} = 0$, a high reserve price is always optimal, but eventually as the competitive distance between the low valuation and high valuation decreases, a high reserve price forgoes too much revenue from the low valuation bidders.

We conclude by briefly commenting on the role of some of the restrictions of our model.

Private vs. Common Values: Throughout, we have assumed that bidders have private values. If values had a common element, we suspect that some information disclosure might be desirable. Recall that, from the literature on static auctions, revenue increases in the amount of information that is being disclosed, as it allows bidders to fine-tune their bids. The same should be true in sequential auctions. As information about other bidders' bids are disclosed, information percolates that might help bidders refine their estimate of the value of the good, and this might mitigate the detrimental effect of public information that has been discussed in this paper. Of course, the latter effect might completely inhibit learning: if the equilibrium is pooling, no information about other bidders' information will ever be transmitted, and learning cannot take place.

Persistent vs. Changing Values: Throughout, we have assumed that values never change. This is an obvious simplification. Its technical convenience is easy to grasp: if values could change over time, in the unobservable case, a loser who eventually wins could no longer be sure that his opponent has a high-valuation, and thus, that his opponent knows that he has a high-valuation as well. Perhaps his eventual win came about because the previous winner's valuation dropped. The impossibility of a pooling equilibrium remains valid, however, as high-valuation bidders have the same incentive to break ties in their favor. The impact of changing values on bidding dynamics is intuitively ambiguous. On one hand, it makes the winner less cautious about lowering his bid, because losing does not imply that the continuation payoff will be zero forever (values will not remain persistently high). On the other hand, the probability that the losers have a low valuation is bounded below (as their value might change from one period to the next), and this dampens the winner's incentive to lower his bid.

Binary Values: We assume throughout that each bidder has either a low or a high value for the object. The binary structure of the possible valuations was particularly helpful in the construction of the equilibrium in the case of unobservable bids (and it was necessary to get unique predictions in the other versions). More precisely, it allowed us to conclude that the bids would jump immediately to the high value after the past winning bidder lost for the first time, and would thereafter stay constant at the high value. With a finite number or a continuum of values, the analysis would become more intricate. With a finite number of values the resolution of uncertainty among two competing bidders with distinct values may require more than one period in which the identity of the winner switches. With two values, any winning bid above the low value identifies a high value bidder. With more than one value above the lowest value, the separation between adjacent values is unlikely to occur in a single event, as it would open the possibility to delay an aggressive bid by a finite number of periods to mimic a lower value type, and hence separation will be slower and with some cost in terms of the efficiency of the equilibrium allocation. As we already know from the two-period analysis of Landsberger, Rubinstein, Wolfstetter, and Zamir (2001), there is no hope in finding a closed-form expression for the strategies if values are drawn from an interval. Furthermore, because from one

period to the next, the winner decreases his bid by some finite amount, there would be no common knowledge of valuations once bid trajectories cross. We believe that ours is the first paper to explicitly solve for a Markov equilibrium in a game in which higher-order beliefs matter, and we hope that it will trigger further developments that will ultimately allow us to study such richer environments.

We should emphasize that we believe that the thrust of the argument, the impossibility of a pooling equilibrium with unobservable bids, and the possibility of a pooling bid with observable bids would remain to hold with finitely many or a continuum of values. After all, if we take \underline{u} and \bar{u} as the lowest and the highest valuations (among a continuum), then a pooling equilibrium leads to a flow payoff $(\bar{u} - \underline{u}) / (n + 1)$ while the flow payoff from the efficient separating equilibrium is $(1 - q)^n / (\bar{u} - \underline{u})$, all evaluated from the point of view of the bidder with the highest valuations. Now, as Proposition 2 established for binary values, for every q , there exists n such that the separating equilibrium ceases to exist, while the pooling equilibrium survives.

7 Appendix

We generalize here the analysis given in the text to the case with $n + 1$ bidders. The loser's bid distribution in period $t \geq 1$ is denoted F_t . The distribution of the highest losing bid is given by $(F_t)^n$, and most of the analysis from the winner's point of view is identical to the case $n = 1$.

Proof of Lemma 1. The high-valuation winner's value function must satisfy the optimality equation

$$V_t(b) = \max_{\beta} \left\{ \frac{F_t(\beta)^n}{F_{t-1}(b)^n} ((1 - \delta)(\bar{u} - \beta) + \delta V_{t+1}(\beta)) \right\},$$

where b is the bid the winner made in period $t - 1$ (as before, we attempt to solve for an equilibrium in which the equilibrium bid is a summary statistic for the entire information of a player). Define $Y_t(b) := F_{t-1}(b)^n V_t(b)$ for all $t \geq 1$. Then $Y_t(b) = \max_{\beta} \{F_t(\beta)^n (1 - \delta)(\bar{u} - \beta) + \delta Y_{t+1}(\beta)\}$, from which it is clear that, since the right-hand side is independent of b , the winner is indifferent over all bids in the relevant interval, and Y_t is independent of b . It follows that, for all $t \geq 1$, and for some constant $\varphi_t \geq 0$, $F_t(b)^n = \frac{\varphi_t}{\bar{u} - b}$. Since our purpose is to construct an equilibrium in which only the low type bidder bids \underline{u} , we further have $(1 - q)^n = F_t(\underline{u})^n = \frac{\varphi_t}{\bar{u} - \underline{u}}$, from which we can solve for the constant φ_t , so that

$$F_t(b) = (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}. \quad (37)$$

This distribution being independent of t , the loser makes a bid that is independent of $t, \forall t \geq 1$. \square

Proof of Lemma 2. Let us define $X_t(b)$ as the continuation payoff of a player with a high valuation \bar{u} who lost in the first period and always bids b . Then

$$X_t(b) = \max_{\beta} \{ \Pr[i \text{ wins in period } t \text{ with } \beta] (1 - \delta)(\bar{u} - \beta) + \delta X_{t+1}(\beta) \}.$$
¹⁹

¹⁹As in the case of two players, this is an abuse of notation, because the continuation payoff is a function of the latest bid β only if this latest bid is the most informative one; however, this is necessarily the case if bidder i does not submit a bid strictly lower than the lowest bid for which he knows he will lose for sure, and this is without loss of generality. Note further that, for any higher bid β , his optimization problem is as if he had bid β throughout.

We do not condition on player i having always lost before, or the winner having always won (*i.e.*, it could be that the game is “over” and i might not know about it). The choice only matters if those two events obtain, but so maximizing the conditional or unconditional payoff is equivalent. Let

$$p_t(b, \beta) := \Pr[i \text{ wins in period } t \text{ with } \beta \text{ and bid } b \text{ in all previous periods}],$$

given that he always bid b before, and lost in the initial period. First-order conditions give

$$\frac{d}{d\beta} p_t(b, \beta)(\bar{u} - \beta) - p_t(b, \beta) + \frac{\delta}{1 - \delta} X'_{t+1}(\beta) = 0,$$

while the envelope theorem states that

$$\frac{X'_t(b)}{1 - \delta} = \frac{d}{db} p_t(b, \beta)(\bar{u} - \beta). \quad (38)$$

Combining gives, for $t \geq 2$ (remember that $b = \beta$ then)

$$\left(\frac{d}{d\beta} p_t(b, b) + \delta \frac{d}{db} p_{t+1}(b, b) \right) (\bar{u} - b) - p_t(b, b) = 0. \quad (39)$$

We must now solve for the probabilities $p_t(b, \beta)$. Fix some player i who lost in the initial period with a bid b . Let G_t denote the unconditional bid distribution of the winner in period t (we do not condition on the fact that player i lost in the initial period with a particular bid b). Also, given t , define the function β by $G_t(b) = G_{t-1}(\beta(b))$. That is, if the initial winner bids b in period t , he must have bid $\beta(b)$ in the previous period. Suppose now that player i bid b in all periods up to $t - 1$. What are his odds of winning for the first time in t with a bid $b - \varepsilon$, for small $\varepsilon > 0$? First, the bid $b - \varepsilon$ must be the highest bid among the losing bids, which occurs with probability $F(b)^{n-1}$. Second, the winner’s previous bid must have been in the interval $[b, \beta(b - \varepsilon)]$: if it were lower, i would have

won before; if it were higher, i would not win in period t . So the probability he wins is

$$F(b - \varepsilon)^{n-1}(G_{t-1}(\beta(b - \varepsilon)) - G_{t-1}(b)) = F(b - \varepsilon)^{n-1}(G_t(b - \varepsilon) - G_{t-1}(b)). \quad (40)$$

Let us now consider instead the case in which he increases his bid to $b + \varepsilon$ in period t . What is the probability that he then wins in that period? The probability is then

$$F(b + \varepsilon)^{n-1} \int_{b+\varepsilon}^{\beta(b+\varepsilon)} g_{t-1}(x)dx + \int_b^{b+\varepsilon} F(x)^{n-1} g_{t-1}(x)dx.$$

Indeed, either the winner bid x in the range $[b, b + \varepsilon]$ in period $t - 1$, and for i to win in t , he must have won in $t - 1$ (*i.e.* the others bid below x , which imply that they are also outbid by i in t), or he bid x in the range $[b + \varepsilon, \beta(b + \varepsilon)]$, and all that is needed then is that the other initial losers bid no more than $b + \varepsilon$. This probability can be rewritten as

$$F(b + \varepsilon)^{n-1}(G_t(b + \varepsilon) - G_{t-1}(b + \varepsilon)) + \int_b^{b+\varepsilon} F(x)^{n-1} g_{t-1}(x)dx. \quad (41)$$

As expected, these probabilities and their derivatives with respect to ε coincide for $\varepsilon = 0$, so

$$p_t(b, b) = F(b)^{n-1}(G_t(b) - G_{t-1}(b)), \quad (42)$$

$$\frac{d}{d\beta} p_t(b, b) = F(b)^{n-1} g_t(b) + (n - 1) f(b) F(b)^{n-2} (G_t(b) - G_{t-1}(b)). \quad (43)$$

We also get, from (40) or (41),

$$\frac{d}{db} p_t(b, b) = -F(b)^{n-1} g_{t-1}(b). \quad (44)$$

We have $\left(\frac{d}{d\beta} p_t(b, b) + \delta \frac{d}{db} p_{t+1}(b, b) \right) (\bar{u} - b) - p_t(b, b) = 0$, and plugging in the value for p_t and its

derivatives from (42)–(44) gives

$$\begin{aligned} (F(b)^{n-1}g_t(b) + (n-1)f(b)F(b)^{n-2}(G_t(b) - G_{t-1}(b)) - \delta F(b)^{n-1}g_t(b))(\bar{u} - b) \\ - F(b)^{n-1}(G_t(b) - G_{t-1}(b)) = 0. \end{aligned}$$

We can then eliminate the higher powers of F and obtain

$$\left((1 - \delta)g_t(b) + (n-1)\frac{f(b)}{F(b)}(G_t(b) - G_{t-1}(b)) \right) (\bar{u} - b) - (G_t(b) - G_{t-1}(b)) = 0, \quad (45)$$

a difference-differential equation to be solved for G_t given G_{t-1} . We further eliminate $f(b)/F(b)$ by observing that from (37), $\frac{f(b)}{F(b)} = \frac{1}{n}\frac{1}{\bar{u}-b}$. Plugging into (45),

$$\left((1 - \delta)g_t(b) + \frac{n-1}{n}\frac{1}{\bar{u}-b}(G_t(b) - G_{t-1}(b)) \right) (\bar{u} - b) - (G_t(b) - G_{t-1}(b)) = 0, \quad (46)$$

or

$$\left((1 - \delta)g_t(b)(\bar{u} - b) - \frac{1}{n}(G_t(b) - G_{t-1}(b)) \right) = 0. \quad (47)$$

The derivation of G_1 follows then exactly the same steps as in the text, giving

$$(1 - \delta)g_1(b)(\bar{u} - b) = \frac{1}{n} \left(G_1(b) - \left(\frac{\gamma}{\bar{u} - b} \right)^{\frac{1}{n}} \right),$$

where $\gamma := (1 - q)^n(\bar{u} - \underline{u})$, g_1 is the derivative of G_1 and $G_1(\bar{u} - \gamma) = 1$. For $t \geq 2$,

$$(1 - \delta)g_t(b)(\bar{u} - b) = \frac{1}{n} (G_t(b) - G_{t-1}(b)),$$

with $G_t(\bar{u} - \gamma) = 1$. Set $y := (\bar{u} - b)/\gamma$, $H_t(y) := G_t(b)$, so $(1 - \delta)nyh_t(y) + H_t(y) - H_{t-1}(y) = 0, t \geq 2$,

and $(1 - \delta)nyh_1(y) + H_1(y) - y^{-\frac{1}{n}} = 0$, where $H_t(1) = 1, t \geq 1$. The solution is

$$H_t(y) = \frac{y^{-1/n}}{\delta^t} + y^{-1/(1-\delta)n} \sum_{\tau=0}^t \frac{1 - \delta^{\tau-t}}{\tau!} \left(\frac{\ln y}{(1-\delta)n} \right)^\tau, \quad (48)$$

that is, in terms of the distribution G_t ,

$$G_t(b) = \frac{1}{\delta^t(1-q)} \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{\frac{1}{n}} + \left[(1-q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{\frac{1}{n}} \right]^{\frac{1}{1-\delta}} \sum_{\tau=0}^t \frac{1 - \delta^{\tau-t}}{\tau!} \left(\ln \left((1-q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{\frac{1}{n}} \right)^{-\frac{1}{1-\delta}} \right)^\tau,$$

which establishes the lemma. \square

Proof of Theorem 1. It remains to determine F_0 , the initial bid distribution. The payoff of bidding b for a player with a high-valuation is given by $F_0(b)^n(\bar{u}-b) + \delta Y_1(b)/(1-\delta) + \delta X_1(b)/(1-\delta)$, where $Y_1(b)$ is the (unconditional) continuation payoff after winning with an initial bid of b , evaluated from the second period onward, and $X_1(b)$ is the (unconditional) continuation payoff from losing after an initial bid of b , evaluated from the second period onward. As in the case $n = 1$, $Y_1(b)/(1-\delta) = (1-q)^n(\bar{u}-\underline{u})$ is a constant. Further, from the envelope theorem,

$$X'_1(b)/(1-\delta) = \frac{d}{db} p_1(b, \beta)(\bar{u} - \beta) = -nF_0(b)^{n-1} f_0(b)(\bar{u} - \beta),$$

where f_0 is the density of the distribution F_0 . From $F_0(b) = F_1(\beta) = (1-q) \left(\frac{\bar{u}-\underline{u}}{\bar{u}-\beta} \right)^n$, we can solve for $\bar{u} - \beta$ in terms of F_0 . Plugging this into the previous formula, we get

$$X'_1(b)/(1-\delta) = -n \frac{f_0(b)}{F_0(b)} (1-q)^n (\bar{u} - \underline{u}).$$

Integrating then yields $X_1(b)/(1-\delta) = -n(1-q)^n(\bar{u}-\underline{u}) \ln F_0(b) + C_0$, for some constant C_0 . Because the payoff of bidding b must be independent of b over the support, we thus obtain that (substituting

for X_1 and Y_1)

$$F_0(b)^n(\bar{u} - b) - n\delta(1 - q)^n(\bar{u} - \underline{u}) \ln F_0(b) = K_0, \quad (49)$$

for some constant K_0 that is independent of b . By using the fact that $F_0(\underline{u}_+) = 1 - q$, we get K_0 by plugging $b = \underline{u}_+$, and so $K_0 = (\bar{u} - \underline{u})(1 - q)^n(1 - n\delta \ln(1 - q))$. By plugging $b = \bar{b}$, where \bar{b} denotes the upper extremity of the support of F_0 , we get that $\bar{b} = \bar{u} - K_0$. This and (49) uniquely characterize F_0 , and the closed-form solution follows from standard properties of the Lambert function.

So far, the equilibrium strategies have only been described on path. What happens after histories that are off the equilibrium path? Note first that there are no non-trivial deviations in the initial period: the high-valuation bidder has nothing to gain from bidding more than the highest bid in the support of his opponents' distribution, nor from bidding less. For the latter, observe that bidding \underline{u}_+ strictly dominates bidding \underline{u} : high-valuation opponents bid more anyway; so it only makes a difference in the event that they all have low valuations, but in that event, it is better for him to break the tie in his favor. Therefore, we may focus on deviations after the initial period. Consider the winner first. Given that the losers make constant bids, the most informative bid b that the winner has made from period 1 to period t , given some arbitrary private history of his, is the lowest bid that he has made so far. Therefore, his beliefs are as if he had made this bid b in period t , and since the winner's equilibrium strategy is onto (*i.e.* for every bid b in the support, there is a history for which this bid b is the equilibrium bid in period t), we may specify that he then behaves as if he had followed all along the equilibrium strategy that leads to bid b in period t (in fact, this specification is implied by the Markov assumption). The situation for a high-valuation loser i is similar: because the winner's bids follow some decreasing trajectory, what matters is, given the private history of bidder i , what is the highest lower bound b that he assigns to the winner's bid in period t ? Then bidder i 's problem is identical to the one he would face had he followed the equilibrium strategy leading to a bid of b in period t , and the Markov assumption then implies that, whatever private history he actually has, he behaves from that period onward according to this equilibrium strategy (*i.e.*, he submits the constant bid b from then on).

Finally, as the reader might recall, the optimality equations for the winner and loser implicitly assumed that bids were chosen in a range ensuring that this bid would be more informative than the previous ones. For the winner, this means that the new bid β is no higher than the previous bid b that he submitted; plainly, given that the losers make constant bids, higher bids are suboptimal (because any bid $\beta \geq b$ is a winning bid anyway). For the loser, this means that the new bid β is not strictly lower than the maximal bid that ensures that, given the previous bid b he submitted, and the winner's equilibrium strategy, the loser is guaranteed to lose again. Plainly again, making any lower bid cannot constitute a profitable deviation. By construction then, there are no profitable deviations. This completes the description of the separating equilibrium. \square

References

- ALBAEK, S., P. MOLLGARD, AND P. OVERGAARD (1997): “Government Assisted Oligopoly Coordination,” *Journal of Industrial Economics*, 45, 429–443.
- BLUME, A. (2003): “Bertrand Without Fudge,” *Economics Letters*, 78, 167–168.
- CANTILLON, E., AND M. PESENDORFER (2006): “Auctioning Bus Routes: The London Experience,” in *Combinatorial Auctions*, ed. by Y. Shoham, P. Cramton, and R. Steinberg, pp. 573–593. MIT Press, Cambridge.
- CASON, T., K. KANNAN, AND R. SIEBERT (2009): “An Experimental Study of Information Revelation Policies in Sequential Auctions,” Discussion paper, Purdue University.
- CHWE, M. (1989): “The Discrete Bid First Auction,” *Economics Letters*, 31, 303–306.
- CRAMTON, P., AND J. SCHWARTZ (2002): “Collusive Bidding in the FCC Spectrum Auctions,” *B.E. Journal of Economics Analysis and Policy: Contributions*, 1.
- DIZDAR, D., A. GERSHKOV, AND B. MOLDOVANU (2011): “Revenue Maximization in the Dynamic Knapsack Problem,” *Theoretical Economics*, 6, 157–184.
- ENGELBRECHT-WIGGANS, R. (1983): “A Sequential Auction Involving Asymmetrically Informed Bidders,” *International Journal of Game Theory*, 12, 123–127.
- FEVRIER, P. (2003): “He Who Must Not Be Named,” *Review of Economic Design*, 8, 99–119.
- HÖRNER, J., AND J. JAMISON (2008): “Sequential Common-Value Auctions with Asymmetrically Informed Bidders,” *Review of Economic Studies*, 75, 475–498.
- JACKSON, M., J. SWINKELS, L. SIMON, AND W. ZAME (2002): “Communication and Equilibrium in Discontinuous Games of Incomplete Information,” *Econometrica*, 70, 1711–1740.

- KAYA, A., AND Q. LIU (2015): “Transparency and Price Formation,” *Theoretical Economics*, 10, 341–383.
- KLEIJNEN, J., AND F. SCHAIK (2007): “Sealed-Bid Auction of Dutch Mussels: Statistical Analysis,” Discussion Paper 83, Tilburg University and University of Amsterdam.
- LANDSBERGER, M., J. RUBINSTEIN, E. WOLFSTETTER, AND S. ZAMIR (2001): “First-Price Auctions When the Ranking of Valuations is Common Knowledge,” *Review of Economic Design*, 6, 461–480.
- MASKIN, E., AND J. RILEY (2003): “Uniqueness of Equilibrium in Sealed High-Bid Auctions,” *Games and Economic Behavior*, 45, 395–409.
- MERTENS, J., AND S. ZAMIR (1985): “Formalization of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 14, 1–29.
- MILGROM, P., AND R. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089–1122.
- OECD (2008): “Fighting Cartels in Public Procurement,” *OECD Policy Brief*, pp. 1–8.
- PAI, M., AND R. VOHRA (1979): “Optimal Dynamic Auctions and Simple Index Rules,” *Mathematics of Operations Research*, 38, 682–697.
- SELTEN, R., AND J. BUCHTA (1999): “Experimental Sealed-Bid First-Price Auctions with Directly Observed Bid Functions,” in *Games and Human Behavior: Essays in Honor of Amnon Rapoport*, ed. by D. Budescu, I. Erev, and R. Zwick, pp. 79–102. Lawrence Erlbaum Associates, Mahwah, New Jersey.
- THOMAS, C. (2010): “Information Revelation and Buyer Profits in Repeated Procurement Competition,” *Journal of Industrial Economics*, 58, 79–105.

TU, Z. (2007): “Why Do We Use the Dutch Auction to Sell Flowers: Information Disclosure in Sequential Auctions,” Discussion paper, University of Pittsburgh.

WILSON, T. (2000): “B2B Sellers Fight Back on Pricing,” *Internet Week*, p. December 12.

YAO, Z. (2007): “Sequential First-Price Auctions with Multi-Unit Demand,” Discussion paper, UCLA.