



Cowles Foundation

for Research in Economics
at Yale University

Cowles Foundation Discussion Paper No. 1764

SHOULD AUCTIONS BE TRANSPARENT?

Dirk Bergemann and Johannes Hörner

August 2010

An author index to the working papers in the
Cowles Foundation Discussion Paper Series is located at:
<http://cowles.econ.yale.edu/>

This paper can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1657947>

Should Auctions be Transparent?*

Dirk Bergemann[†] Johannes Hörner[‡]

August 11, 2010

Abstract

We investigate the role of market transparency in repeated first-price auctions. We consider a setting with private and independent values across bidders. The values are assumed to be perfectly persistent over time.

We analyze the first-price auction under three distinct disclosure regimes regarding the bid and award history. Of particular interest is the minimal disclosure regime, in which each bidder only learns privately whether he won or lost the auction at the end of each round. In equilibrium, the winner of the initial auction lowers his bids over time, while losers keep their bids constant, in anticipation of the winner's lower future bids. This equilibrium is efficient, and all information is eventually revealed. Importantly, this disclosure regime does not give rise to pooling equilibria.

We contrast the minimal disclosure setting with the case in which all bids are public, and the case in which only the winner's bids are public. In these settings, an inefficient pooling equilibrium with low revenues always exists with a sufficiently large number of bidders.

KEYWORDS: First Price Auction, Repeated Auction, Private Bids, Information Revelation

JEL CLASSIFICATION: D44, D82, D83.

*This research is supported financially by NSF Grant SES 0851200 and NSF Grant SES 0920985. We would like to thank Vitor Farinha Luz for research assistance and seminar participants at the University of Mannheim and University of Munich, as well as at the GDR Jeux at Luminy for their constructive comments.

[†]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

Information revelation policies vary wildly 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. But the pinnacle of opaqueness surely belongs to Google for its sponsored search auction, in which even the algorithm based on which the winner is determined is unknown. To add to this confusion, some B2B auction platforms (for instance, FreeMarkets, or Covisint) offer their clients the possibility to choose how transparent, or opaque, they wish the auction format to be.

Undoubtedly, the choice of such feedback policies reflects a variety of considerations, such as security, privacy, risk of corruption, etc. 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) are publicly disclosed at the end of each round.

Of course, 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 which 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 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

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

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.

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 Organisation for Economic Co-Operation and Development (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.

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 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.

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 compares the results and Section 7 concludes. The appendix, Section 8, 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 V^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 would considerably complicate the analysis of the case in which bids are not observable.

Thus, the reward r_t^i in period t of player i with valuation V^i is equal to $V^i - b_t^i$ if he bids b_t^i and he wins the object, or is equal to 0 if he does not win the object. 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 information about actions, 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 a little bit tricky. In the observable case, players' beliefs (about others' values) are public after every history, and we can then 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 might have been, 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 losing bid, the loser has beliefs about the winner's beliefs, and they are certainly payoff-relevant from his point of view. 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 this is getting ahead of ourselves. First, let us define more precisely the information and the strategies available to the bidders 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, along with each private history into a distribution over bids.

In the observable case, bidder i knows the entire (ordered) sequence of bids in each period up

to t . That is, in that case, the public history up to period t is 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 period t is denoted $h_t \in H_t$ and it is equal to $((\mathbb{R}_+ \cup \{\underline{u}_+\})^n \times \{1, \dots, n\})^t$. (We set $H_0^i := \{\emptyset\}$, $H_0 := \{\emptyset\}$.)

In the winner-only observable case, the public history up to period 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 determines 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, then, is then summarized by a measurable map $\sigma^i : \Theta^i \rightarrow \Delta(\mathbb{R}_+ \cup \{\underline{u}_+\})$, and a Markov sequential equilibrium (hereafter, MSE) is a PBE in Markov strategies. In the observable case, these hierarchies of beliefs are trivial. As we shall see, they are equally simple in the case of winner-only observable bids. 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. The auction consists of $n + 1$ bidders, $i = 1, \dots, n + 1$, with two possible valuations, $V^i \in \{\underline{u}, \bar{u}\}$, 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 distribution of bids from each of his competitors, either with low or high valuation, is denoted by $F(b)$. Now, every bid in the support of the random bidding strategy must maximize the expected payoff

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

The indifference of the high-valuation bidder requires that the right-hand side be independent of b on its support, 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 over $(\underline{u}, \bar{u} - (1 - q)^n(\bar{u} - \underline{u}))$. Note that the low-valuation bidder receives zero net utility, while the high-valuation bidder receives a positive expected net utility 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. It is worth pointing out that, 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. It is easy to show, however, that the allocation from the first-price auction maximizes revenue among efficient mechanisms.²

3 Unobservable Bids

We begin our analysis with the case of unobservable bids. We shall first argue that the equilibrium cannot be pooling.

²Unlike other allocations that will come up, this allocation has the property that the incentive constraints are not uniformly strict.

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. But in the process of separation from a possible low-valuation bidder, 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 would naturally mean that the surplus would have to be shared, in particular with low-valuation bidders. But the benefit would be that the price would remain 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. But this opens the possibility for a high-valuation bidder to bid slightly more than the pooling bid \underline{u} , at a negligible cost, and to win the current auction for sure, leaving him with an increase in the current rent. As beliefs of the agents do not change, 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 sequential equilibrium does not exist with unobservable bids.

We should point out that 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 both the low- and high-type bidders are simultaneously indifferent over two different bids (i.e., over two different probabilities of winning: recall that 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). Second, note that this pooling bid must (at least in some period) be no larger than \underline{u} , for otherwise the low-type bidder would make negative profits. Consider any such period, and apply the argument given above.

Having established the impossibility of pooling in a Markov sequential 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 subsequent 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.³ 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, in the sense that the highest bid that a loser could conceivably make in a given period, i.e., the bid a loser would make if he lost

³We 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.

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 the bid \bar{b}_0 .

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.⁴ 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 (for as soon as he does, it becomes common knowledge that two bidders’ valuations are high.)

Note also that, in such an equilibrium, the process of belief updating remains relatively simple. To fix ideas, consider the case of 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

⁴Note, 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.

mostly interested in the qualitative findings might elect to skip it.

3.3 Deriving the Equilibrium Strategies

Fix the separating Markov sequential equilibrium to be described. Let F_t denote 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 that he faces 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 bid distribution of a loser who would have bid less than \underline{u} throughout.

The distinction between winner and loser is immaterial in the initial period, and thus $G_0 = F_0$. Given the aforementioned properties of the separating equilibrium we seek, it must be that $F_t(\underline{u}) = 1 - q$, where we recall that $1 - q$ is the prior probability that the valuation of the bidder is \underline{u} . 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 $G_t(\underline{u}_+) \geq 1 - q$.⁵

In general, a player's beliefs are pinned down by his entire private history. It turns out that, in the equilibrium we describe, 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 the general case of $n + 1$ bidders, and their proofs can be found in the Appendix.

⁵To 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.

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 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 .⁶ 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 $F_t(\beta)$. Now, the winner received the object in the preceding period with a bid b , and 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 $\beta_t(b)$. Clearly, 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 one more period. By contrast, if he loses the auction today, then it is common knowledge among the bidders that they both have a high valuation for the object, and hence, by Refinement **A**, all future bids will have to be equal to the high value \bar{u} and exhaust all the 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

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

⁶Implicitly, 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.

This allows to rewrite the value function of the winner as

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

which has the advantage that the unconditional distribution F_{t-1} of the preceding period $t-1$ no longer appears. Note also 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. It follows that the first term of the right-hand side must be constant over the support of F_t , which means that the equilibrium bid distribution of the loser is given by

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

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 have thus solved for the unconditional bid distribution of the loser, for all $t \geq 1$, as

$$F_t(b) = (1-q) \frac{\bar{u}-\underline{u}}{\bar{u}-b}. \quad (5)$$

Recall that we are looking for an equilibrium in which bids in periods $t \geq 1$ are deterministic functions of the initial bid (as long as the bidder's status as loser or winner persists). Since $F_t(b)$ is independent of t , this means that, from time $t \geq 1$, a high-valuation loser makes a constant bid (until he wins, if ever). The unconditional bid distribution of the loser is in fact identical to the equilibrium bid distribution (1) in the static auction derived earlier.

Let \bar{b} denote the highest bid in the support of a bidder's distribution.⁷ Then the bid \bar{b} must satisfy

$$(1-q)(\bar{u}-\underline{u})/(\bar{u}-\bar{b}) = 1,$$

⁷To be clear, the expression "support" refers to the support of the unconditional distribution. That is, it is the union over all bids that are made by the loser (resp. the winner) in the relevant period, over all histories of equilibrium bids that this player might have made so far.

and thus the net flow value of the winner is given by

$$\bar{u} - \bar{b} = (1 - q)(\bar{u} - \underline{u}). \quad (6)$$

This residual surplus is equal to the expected surplus that a bidder with a high valuation would receive in the first- or second-price static auction. Moreover, as the highest bid \bar{b} guarantees the winner to win the object in all future auctions with probability one, we have an early indication that the revenue of the repeated auction with unobservable bids may satisfy a flow revenue equivalence with a single static auction.

We record here the equilibrium strategy for the loser in case that there are $n + 1$ bidders, which is an immediate generalization of the formula above (see (38) in the Appendix). We insist here that this is a description of the equilibrium strategies on the equilibrium path. See section 3.4.2 for details about play off the equilibrium path.

Lemma 1 (The Loser's Bid Distribution).

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

$$F_t(b) = F(b) := (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}, \quad (7)$$

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.

Finally, note that (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

$$Y_t = \varphi = (1 - q)(\bar{u} - \underline{u}), \quad (8)$$

which is a constant value, independent of time and of the past bid b . With the solution to the loser's bid distribution and the unconditional payoffs, given by (7) and (8), we obtain the conditional

value of the winner, $V_t(b)$, by using the equation (3) as:

$$V_t(b) = \bar{u} - b. \quad (9)$$

The flow continuation value of the winner in period t is therefore equal to his value minus his successful bid in the past period. This simple representation of the continuation value follows directly from the constant bid distribution of the losers. Given that the bid b succeeded in the past period, the winner knows that by continuing to bid b forever, he can indeed win forever. Moreover, given the bidding strategy by the losers, the winner is indifferent between continuing to bid b or lowering his bid to β . As we will show next, the winner's bid will have to be decreasing, rather than constant, to maintain the indifference of the losers. Nonetheless, using the same intuition, or formally the recursion of the value function given by (2), it follows that the continuation value of the winner is a martingale, i.e.

$$\mathbb{E}[V_{t+1}] = V_t(b) = \bar{u} - b. \quad (10)$$

We shall come back to the martingale property of the continuation value shortly to obtain a complete characterization of the intertemporal properties of the equilibrium bids and equilibrium values.

3.3.2 The Winner's Bidding Strategy

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

$$W_t(b) = \max_{\beta} \left\{ \frac{G_t(\beta) - G_{t-1}(b)}{1 - G_{t-1}(b)} (1 - \delta)(\bar{u} - \beta) + \delta \frac{1 - G_t(\beta)}{1 - G_{t-1}(b)} W_{t+1}(\beta) \right\}, t \geq 1. \quad (11)$$

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 flow payoff of $\bar{u} - \beta$ (after which

bids jump to \bar{u}) or it can be insufficiently low, in which case the loser remains in his loser's status, at least 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 event of the past period, in which, by construction, the loser lost with a bid b , so that the winner's bid must have been at least as high. As before in the analysis of the winner's problem, it is useful to restate this equation with the help of an auxiliary function. We denote by $X_t(b)$ the expected continuation value from losing with a bid b in period $t - 1$, or

$$X_t(b) := (1 - G_{t-1}(b)) W_t(b). \quad (12)$$

With this definition, we rewrite (11) to get

$$X_t(b)/(1 - \delta) = \max_{\beta} \{(G_t(\beta) - G_{t-1}(b))(\bar{u} - \beta) + \delta X_{t+1}(\beta)/(1 - \delta)\}, \quad t \geq 1. \quad (13)$$

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

$$G'_t(\beta)(\bar{u} - \beta) - (G_t(\beta) - G_{t-1}(b)) + \frac{\delta}{1 - \delta} X'_{t+1}(\beta) = 0, \quad t \geq 1. \quad (14)$$

Note also that, from the envelope theorem applied to (13),

$$X'_t(b)/(1 - \delta) = -G'_{t-1}(b)(\bar{u} - \beta). \quad (15)$$

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 (14),

$$G'_t(b)(\bar{u} - b) - (G_t(b) - G_{t-1}(b)) + \frac{\delta}{1 - \delta} X'_{t+1}(b) = 0. \quad (16)$$

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 (15), and using the fact that $\beta = b$, we can eliminate $X'_{t+1}(\beta)$ from (14), to obtain

$$(1 - \delta)G'_t(b)(\bar{u} - b) = G_t(b) - G_{t-1}(b), \quad t \geq 2. \quad (17)$$

Because the support of F_t and G_t must coincide, we also have that $G_t(\bar{u}) = 1$. Thus, we have an ordinary differential equation and a boundary condition that allow us to solve for G_t , provided that 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 (14), is given by

$$G'_1(\beta)(\bar{u} - \beta) - (G_1(\beta) - G_0(b)) + \frac{\delta}{1 - \delta} X'_2(\beta) = 0. \quad (18)$$

Now, while we cannot assert anymore that $\beta = b$ in $t = 1$, we can appeal to the hypothesis of monotone bidding strategies to relate the bid b in $t = 0$ to the bid β in $t = 1$ by observing that

$$G_0(b) = F_0(b) = F_1(\beta) = (1 - q)(\bar{u} - \underline{u}) / (\bar{u} - \beta), \quad (19)$$

where the final equality had been established in (5). Thus, we can write (18) as

$$G'_1(\beta)(\bar{u} - \beta) - G_1(\beta) + \frac{(1 - q)(\bar{u} - \underline{u})}{\bar{u} - \beta} + \frac{\delta}{1 - \delta} X'_2(\beta) = 0.$$

Equation (15) for $t = 2$ can be used here as well to eliminate $X_2'(\beta)$. We obtain that the differential distribution G_1 must satisfy the differential equation

$$(1 - \delta)G_1'(b)(\bar{u} - b) = G_1(b) - \frac{(1 - q)(\bar{u} - \underline{u})}{\bar{u} - b}, \quad (20)$$

along with the boundary condition $G_1(\bar{u}) = 1$.

Thus, the equations (17) and (20), 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 (17) and (20) is the special case for two bidders of the formula given in the following lemma, where as mentioned, by convention, we interpret $G_t(\underline{u}) - (1 - q)$ as 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 $t \geq 1$,

$$G_t(b) = \frac{1}{\delta^t} F(b) + F(b)^{\frac{1}{1-\delta}} \sum_{\tau=0}^t \frac{1 - \delta^{\tau-t}}{\tau!} \left(\ln F(b)^{-\frac{1}{1-\delta}} \right)^\tau, \quad (21)$$

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

It follows from this formula that the bids of the winner are decreasing over time from $t \geq 1$ onward. A closed-form expression for the distribution can be derived from (21), which involves the incomplete gamma function. As it yields no further insight, we omit it here.

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

$$\max_b \left\{ F_0(b)(\bar{u} - b) + \frac{\delta}{1 - \delta} Y_1(b) + \frac{\delta}{1 - \delta} X_1(b) \right\}. \quad (22)$$

The bid b in the initial period determines the flow payoff $(\bar{u} - b)$ and the likelihood of receiving it. In addition, it determines the continuation value conditional on being the winner $Y_1(b)$ or the loser $X_1(b)$, where we maintain the notation that we introduced in (3) and (12) that already

accounts for the likelihood of each event. We also recall that the continuation value from winning $Y_1(b)$ is independent of b , and equal to $(1 - q)(\bar{u} - \underline{u})$.

It remains to determine the continuation value $X_1(b)$ from losing in the initial period $t = 0$. From the envelope theorem (15), we know that $X_1'(b) = -G_0'(b)(\bar{u} - \beta)$, and from (19), we have that $(\bar{u} - \beta) = (1 - q)(\bar{u} - \underline{u})/G_0(b)$. Combining, we get that

$$X_1'(b)/(1 - \delta) = -(1 - q)(\bar{u} - \underline{u}) \frac{G_0'(b)}{G_0(b)}.$$

Note that $X_1' < 0$, i.e., X_1 is decreasing in b . Using that $F_0 = G_0$, and taking the indefinite integral gives

$$X_1(b)/(1 - \delta) = -(1 - q)(\bar{u} - \underline{u}) \ln F_0(b) + C_0,$$

for some constant C_0 . If in the initial period bidders are indifferent over all bids in some interval, the profit computed in (22) must be independent of b over this interval. Thus, inserting our formulas for X_1, Y_1 into (22), it must be the case that

$$F_0(b)(\bar{u} - b) - \delta(1 - q)(\bar{u} - \underline{u}) \ln F_0(b) = K_0, \tag{23}$$

for some constant K_0 . This implicitly defines F_0 . Because $F_0(\underline{u}) = 1 - q$, it follows that

$$K_0 = (1 - q)(\bar{u} - \underline{u})(1 - \delta \ln(1 - q)).$$

Note, in particular, that the highest bid in the support is given by $\bar{u} - K_0$ (plug $F_0(b) = 1$ into (23)). The defining equation of the initial bid distribution can now be solved using the Lambert function.⁸ Formally, we have the following lemma.

Lemma 3 (The Initial Bid Distribution).

⁸The 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 $W_{-1}(0_-) = -\infty$.

The bid distribution at $t = 0$ is given by

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

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

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

In particular, as $\delta \rightarrow 0$, the bid distribution in the initial period converges to the static auction. More generally, the bid distribution of the static auction first-order stochastically dominates this initial distribution (cf. (23)). Bids jump up from the initial to the second period, for the (high-valuation) winner and the losers' alike.

3.4 The Separating Equilibrium: Summary

3.4.1 Main Features of the Equilibrium

Lemmata (1)–(3) provide explicit solutions for the equilibrium strategies. It is time to summarize the main qualitative findings that were either mentioned in passing, or that immediately follow from these solutions.

We begin with the observation that the bids in the initial period are lower than in the static first-price auction, as established by Lemma 3. 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. This would suggest that winning the initial auction is like winning the static auction, but with an additional prize provided by a more favorable continuation value. Now, if the continuation value of winning or losing would be independent of the current bid, then clearly each bidder would bid more aggressively initially as it would look like the static auction, but with a prize strictly larger than the flow payoff \bar{u} . But the analysis of the bidding problem in the initial period, given by (22), demonstrates that continuation value depend on the information provided by the initial bid in an important way. In fact, we showed that the payoff contribution from winning, $Y_1(b)$, is constant in the initial bid b . Now, since the probability of winning is increasing with a larger

bid, this implies immediately 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 corresponding first price auction as the incentives to make high 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.




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 $t = 6$)

The reason why the winner decreases his bid over time is obvious: he cautiously explores how low he can get while still winning. At some point (except again, if his bid was at the very top, a zero probability event), 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 particular reason to raise his.

Although he has lost so far in every period, which should push him towards higher bids, he also 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 than if he increased his bid. The equilibrium exactly balances these forces, so that a constant bid is optimal.

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).

We just observed that the price of the winner (conditionally on continued winning) is decreasing over time. But as the current winner may eventually lose against a competitor, subsequent bids may either go up or down, and hence we are also interested in the evolution of the winning bid, independent of the identity of the bidder. The martingale property of the value function of the winner, established earlier in (10), yields the critical insight, and establishes that the winning bid is a supermartingale. We recall that in equilibrium, the current winner loses at most once, and thereafter, the winning bids jumps upward to \bar{u} , and the object is allocated randomly among the highest bidders. Now, in the random event of the first loss, the realized utility for the past winner is zero. We now observe that the realized utility would also be zero if we were to replace the actual bid with the higher (fictional) bid \bar{u} and now reassign the object to the past winner, and thereafter would assign the object with probability one to the past winner at price (and bid) \bar{u} . Under this modified bid process, the value function of the winner is unchanged and the flow of realized utilities is also exactly as in the equilibrium. But under this modified bid process, the past winner receives the object with probability one forever, and as his value is constant, and the

value function is a martingale, so must be the modified bid process. But as the true winning bid is always below the modified bid process, it follows that the winning bid process is a supermartingale. Now, the revenue of the seller is exactly the stream of winning bids, and hence his value function is a supermartingale as well, i.e. it is decreasing in expectation. As the social value of the allocation is constant following the initial period, hence a martingale, it follows that the value function of the loser must pick up the residual movement, and hence that it is a submartingale, i.e. the prospects of the losers are increasing over time in expectation.

Let us mention a few comparative statics here, formally established in the 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. This is illustrated in Figure 2, which graphs the same functions as Figure 1, but for a much higher discount factor ($\delta = 99/100$). As is clear, bid functions look very similar to each other (and to the static first-price bid distribution). The horizontal line corresponds to the bid distribution in period 54; therefore, it is still the case that bids eventually go to \underline{u} , and they do so quite fast.




Figure 2: Bid distribution for $n = 4$ in periods $t = 1, \dots, 6$ and 54, $q = 1/3$, $\delta = 99/100$, $\bar{u} = 1 = 1 - \underline{u}$ (bottom $t = 1$, top $t = 54$)

Figure 3 shows the same bidding curves of the winner with 4 rather than 2 bidders (the dashed lines are the corresponding curves for $n = 2$ from Figure 2). We see that the bids are higher (bidding is more aggressive) with more bidders, and that this persists over time, as the winner is more careful, here as well, in his bid discovery process.




Figure 3: Bid distribution for $n = 4$ in periods $t = 1, \dots, 6$, $q = 1/3$, $\delta = 9/10$, $\bar{u} = 1 = 1 - \underline{u}$ (bottom $t = 1$, top $t = 6$)

We summarize these observations in the following lemma.

Lemma 4. *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 \bar{u} as $n \rightarrow \infty$.*
4. *Expected revenue tends to the revenue of the static auction as $\delta \rightarrow 1$.*

What if δ models frequency, rather than patience? If we set δ equal to $e^{-r\Delta}$, so that the discount factor corresponds to the discounting over an interval of time Δ , given some constant interest rate $r > 0$, then, for fixed t , as $\Delta \rightarrow 0$, the distribution G_t tends to a well-defined limit.

Define $T := \frac{1}{r} \ln \frac{1}{1-q}$. It is shown in Appendix (see section 8.1.1) that

$$\lim_{\Delta \rightarrow 0} G_t(b) = \begin{cases} e^{rt}(1-q) \left(\frac{\bar{u}-u}{\bar{u}-b}\right)^{\frac{1}{n}} & \text{if } b \in (\underline{u}, \bar{u} - (\bar{u} - \underline{u})(1-q)^n e^{nrt}) \text{ and } t < T, \\ 1 & \text{if } b > \underline{u} \text{ otherwise.} \end{cases}$$

3.4.2 Completing the Equilibrium Specification

As mentioned, the equilibrium strategies have only been described on path so far. What happens after histories that are off the equilibrium path?

Note 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.⁹ Therefore, we may focus on later deviations.

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

⁹For 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.

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.

We conclude this section by commenting on uniqueness. 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 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. At the most we were able to show that the high-valuation bidder cannot be willing to bid \underline{u} indefinitely.

Furthermore, it is not clear either that, within the class of strategy profiles in which the high-valuation bidders always bid at least \underline{u}_+ , our equilibrium is unique. This requires monotonicity in the strategies: if the winner's second-period bid is decreasing in the initial bid, then it might make sense for the loser's second-period bid to be decreasing in his initial bid as well, and conversely. At least, we were not able to rule out such strategy profiles. What is clear, however, is that the current separating Markov equilibrium could be used as a punishment scheme to support strategy profiles that are not Markov, such as bid rotation.

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 the equilibrium 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

$$(1 - \delta)(1 - q)^n(\bar{u} - \underline{u})/(n + 1) + \delta((1 - q)^n + (1 - \delta)nq(1 - q)^{n-1})(\bar{u} - \underline{u}). \quad (24)$$

To understand (24), 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 will allow 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. We obtain the following result.

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 sequential

equilibrium does not exist with observable bids.

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

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

Clearly, 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 subsequently 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 sequential 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.¹⁰ 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 prob-

¹⁰To prevent any confusion, we avoid the notation F, G introduced in the unobservable case.

ability 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 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.¹¹ Finally, let \bar{b}_t denote the highest bid in the bidders' support in period $t \leq T$ (conditional on a history in which all informed bidders bid \underline{u} throughout, a statement we shall omit from now on). We observe that while this is a private-values setting, this game and its solution bear some similarities with the common-values game analyzed in Hörner and Jamison (2008). Therefore, our analysis will be complete, but concise.

Because the informed bidder will always bid 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 $T < \infty$. Hence, 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} . 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

$$1 - q_{t+1} = \frac{1 - q_t}{F_t^I(\underline{u})}, \quad (25)$$

and so

$$1 - q_0 = \prod_{t=0, \dots, T} F_t^I(\underline{u}).$$

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

¹¹If $T = +\infty$, 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, $q_{T+1} = 0$.

from the Markov assumption), we have, for all $t \leq T$,

$$F_t^U(b)F_t^I(b)^{n-1}(\bar{u} - b) = \bar{u} - \bar{b}_t = F_t^I(b)^n(\bar{u} - b), \quad (26)$$

and so

$$F_t^U(b) = F_t^I(b) =: F_t(b) \text{ for all } b \text{ and all } t, \quad (27)$$

with the convention that $F_t^U(\underline{u})$ is the probability assigned by the uninformed bidder to \underline{u}_+ .

Finally, because an informed bidder is indifferent between bidding just above \underline{u} in period 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 $t < T$,

$$F_t^n(\underline{u}) = \delta F_t^{n-1}(\underline{u})F_{t+1}^n(\underline{u}), \text{ or } F_t(\underline{u}) = \delta F_{t+1}^n(\underline{u}). \quad (28)$$

The equality must be replaced by an inequality in period T , i.e., since $F_{T+1}(\underline{u}) = 1$, we must have $F_T(\underline{u}) \geq \delta$.

Equations (25), (26) and (28) allow us to solve for the equilibrium. In particular, (25) and (28) yield

$$1 - q_0 = \prod_{t=0, \dots, T} F_t(\underline{u}) = \delta^{\frac{n^{T+1} - (n-1)^{T-n}}{(n-1)^2}} F_T(\underline{u})^{\frac{n^{T+1} - 1}{n-1}} \in I_T \quad (29)$$

where the intervals I_T , subsets of the unit interval, are defined as

$$I_T := \left[\delta^{\frac{n^{T+2} - (n-1)^{(T+1)-n}}{(n-1)^2}}, \delta^{\frac{n^{T+1} - (n-1)^{T-n}}{(n-1)^2}} \right].^{12}$$

Because the intervals $\{I_T\}_{T \in \mathbb{N}_0}$ form a partition of the unit interval $[0, 1)$, there exists a unique solution $T(q_0)$ to (29). Thus the game of one-sided incomplete information ends in finite time with the discovery or not of a second high-valuation bidder. Given $T = T(q_0)$, we obtain $F_T(\underline{u})$ from (29). From (28), we then get $F_t(\underline{u})$ for all $t \leq T$, and (26), applied to an arbitrary b and to $b = \underline{u}$ then gives the distribution F_t (and the value of \bar{b}_t) for all t .

¹²For $n = 1$ (two bidders), Equation (29) uses the convention $(n^{T+1} - 1)/(n - 1) = T + 1$ and $(n^{T+1} + T - n(T + 1))/(n - 1)^2 = T(T + 1)/2$.

Given a prior belief $q < 1$, it is easy to see 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 5 (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 of the original 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

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

as $\delta \rightarrow 1$. Writing out the payoff V^U of the uninformed player is a little messier, and so is omitted here (see Appendix for more details). For our purposes, 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 only comes from the possibility that all his opponents have low valuations. This is good news 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

$$\frac{1}{n+1} \geq 1 - \delta + \delta V^U(q)/(\bar{u} - \underline{u}). \quad (30)$$

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

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 sequential 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 then get

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

Note that the left-hand side of (30) tends to the lowest payoff that a high-valuation bidder can guarantee, provided low-valuation bidders do not bid more than \underline{u} . Indeed, a high-valuation bidder can always secure $(1 - q)^n(\bar{u} - \underline{u})$ by always bidding \underline{u}_+ . Therefore, the pooling equilibrium exists whenever it yields an individually rational payoff to the high-valuation bidder. As we shall see, the same observation 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 $q \in (\underline{q}^o, \bar{q}^o)$. In that case, the equilibrium must involve semi-pooling (when $\delta \rightarrow 1$). 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




Figure 4: Markov equilibria in the observable case, as a function of q .

some point, beliefs are such that a separating equilibrium exists. In 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. On the other hand, we have been unable to rule out the existence of such equilibria in the region in which pooling equilibria exist (and we suspect such equilibria might exist). We summarize this discussion in the following theorem. See also Figure 4.

Theorem 1. *A (Markov sequential) equilibrium always exists. Furthermore,*

1. *If $q \in [0, \underline{q}^o]$, the unique equilibrium is separating.*
2. *If $q > \underline{q}^o$, no separating equilibrium exists. Furthermore, if*
 - (a) *$q \in (\underline{q}^o, \bar{q}^o]$, all equilibria are semi-pooling;*
 - (b) *If $q > \bar{q}^o$, 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.¹³ Again, in order to solve the game in which bidders have symmetric information, we must start with the case in which exactly one bidder is known to have a high valuation, while the others are not, because such informational structures arise in continuation games 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. In consequence, 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 the incomplete information was resolved. In the current environment where only the winning bid is observable, the updating process actually depends on the realized bid of the uninformed bidder. His bid, as long as it is the winning bid, provides an upper bound for losing bids. In consequence, the level of his bid will determine the rate at which the updating occurs. In contrast to the observable case, it is therefore convenient to describe the strategies in terms of the uninformed bidder's posterior belief about the informed bidder's valuation, which all bidders can derive from the observable bids of this bidder.

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.

¹³While we assume so, it is not necessary that the winner's identity be disclosed: if it is not disclosed, 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.

By winning, he not only gets the reward in that period, but he also increases the probability that the game “goes on” (if another bidder ever outbids him, then it will be known that there are two high-valuation bidders, and so their bids will be \bar{u} from then on). It follows that the uninformed bidder bids at least \underline{u}_+ .

Any informed bidder with a high valuation 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 winning with such a bid would reveal that he has a high 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 bidding distribution is therefore also indexed by the belief q . Write F_q^U for this distribution.

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 pure strategy could not be optimal. Clearly, he cannot be the only player submitting bids in this range, and the high-valuation informed bidder must randomize over this interval 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 ,

$$V^U(q) = (1 - q)^n(\bar{u} - \underline{u}) = \bar{u} - \bar{b}_q. \quad (31)$$

Note that, by Bayes’ rule, the probability that the uninformed bidder assigns to any of his op-

ponents 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,

$$\begin{aligned} (1 - q)^n(\bar{u} - \underline{u}) &= F_q^I(b)^n \left((1 - \delta)(\bar{u} - b) + \delta V^U \left(1 - \frac{1 - q}{F_q^I(b)} \right) \right) \\ &= F_q^I(b)^n (1 - \delta)(\bar{u} - b) + \delta (1 - q)^n (\bar{u} - \underline{u}), \end{aligned}$$

where the second equality uses the first equality of (31) to eliminate V^U . It follows that

$$F_q^I(b) = (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}. \quad (32)$$

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

$$V^I(q) = (1 - \delta)(\bar{u} - \bar{b}_q) = (1 - \delta)(1 - q)^n(\bar{u} - \underline{u}), \quad (33)$$

where we have used (31) 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 the object. If he loses, then he only gets a positive continuation payoff if the informed bidder wins. That is, the informed bidder must bid some $\beta > b$, and all other informed bidders must bid less than β . Therefore,

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

where the second equality uses (32) and (33). Plugging into (33) 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 (31)), to get

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

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, since winning carries the additional benefit of prolonging the game. This additional benefit becomes all the more important as the discount factor increases, and as $\delta \rightarrow 1$, the uninformed bidder is very likely to make a bid that is near the upper end of the bid support.

Equations (32) and (34) fully characterize the equilibrium of the game with one-sided incomplete information, with payoffs given by (31) and (33). Note that, after bidding b , the uninformed bidder updates his beliefs to $q' = 1 - (1-q)/F_q^I(b) = 1 - ((\bar{u}-b)/(\bar{u}-\underline{u}))^{1/n}$, so that the highest bid he makes in the following period is

$$\bar{b}_{q'} = \bar{u} - (1 - q'^n)(\bar{u} - \underline{u}) = b.$$

That is, the uninformed bidder's bids are non-increasing over time: while he bids more aggressively than the informed bidders, the uninformed bidder cautiously decreases his bids over time nevertheless, as he becomes more optimistic that his opponents have low valuations. We summarize this discussion in the following lemma.

Lemma 6 (Bidding with One-Sided Information and Winner-Only Observable Bids).

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

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.*

The bidding game with one-sided information in the winner-only observable bid environment shares a number of features with the bidding game after the initial period in the unobservable bid environment. The similarities in the strategies clearly can be traced back to the similarities in the information held by the participating agents. 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 bidding behavior either. The losing bidders in either game form know that they face a bidder with a high valuation. In addition, in the winner-only observable environment they know the past bids of the winning bidder. In either game form, the losing bidders have the same opportunities, namely to win exactly once against the current winning bidder. In consequence, the equilibrium bid distribution of the losing bidders as a function of the posterior belief q , is identical, as expressed by (7) and (32), respectively. By contrast, the behavior of the winning bidder is stationary (in the posterior q) in the winner-only observable environment. As the winning bid is observable, the level of the winning bid carries no additional information, and thus the winner is not concerned about the informational content of his bid.

At this point, it is also 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 in the current period. Indeed, the bidding strategy of an informed bidder in the observable environment is less aggressive. In particular, from (25), we see that 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 (26) and (32), 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 (27) with (34). The defensive posture of the bidders in the observable environment also translates into lower revenues for the seller and higher net values for the bidders.

Proposition 4 (Revenue Comparison with One-Sided Incomplete Information).

1. *The equilibrium values of the informed and uninformed bidders are larger in the observable environment than in the winner observable environment.*
2. *The equilibrium revenues of the seller are smaller in the observable environment than in the winner observable environment.*

We emphasize that the differential distribution of the surplus in these two environments occurs even as the equilibrium allocation is efficient in both environments. In the observable environment, the informed bidder faces a larger exposure risk and hence bids less aggressively. It is the very possibility of exposure that also supports a pooling equilibrium for a larger set of parameters in the observable than in the winner-observable environment, as we will see in the next result.

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

$$(1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q) = (1 - \delta)(\bar{u} - \underline{u}) + \delta(1 - q)^n(\bar{u} - \underline{u}), \quad (35)$$

where the equality uses (31). 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 sequential equilibrium does exist with winner-only observable bids.

In fact, by comparing the two possible payoff streams above, we can assert that pooling is an equilibrium whenever

$$\frac{1}{n+1} \geq 1 - \delta + \delta(1-q)^n, \quad (36)$$

i.e. the belief q should exceed some threshold q^w . Note that, as $\delta \rightarrow 1$,

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

which yields the same condition as in the case of observable bids.

Let us now examine a candidate separating equilibrium. This is an equilibrium in which, in the initial period, all high-valuation bidders submit bids that are at least \underline{u}_+ . The distribution F_q on $[\underline{u}, \bar{b}_q]$ that characterizes this equilibrium must satisfy

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

where the two terms of the left-hand side correspond to the respective payoffs from winning and losing with a bid $b > \underline{u}$, and the right-hand side is this payoff computed for the bid $b = \bar{b}$. Simplifying (using (31) and (33)) 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

$$\frac{\bar{u}-b}{\bar{u}-\underline{u}} \left(\frac{F_q(b)}{1-q} \right)^n = 1 + \delta \ln \left(\frac{F_q(b)}{1-q} \right)^n, \quad (37)$$

using the boundary condition $F_q(\underline{u}) = 1 - q$ that characterizes a separating equilibrium. Equation (37) uniquely characterizes the distribution F . An explicit formula can be given in terms of the (branch -1 of the) Lambert function W , namely

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

From (37), we can also solve for the upper extremity of the support of the bid distribution,

$$\bar{b}_q = \bar{u} - (1 - q)^n (1 - \delta n \ln(1 - q)) (\bar{u} - \underline{u}).$$

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. We also observe that the bidding supports coincide for $t \geq 1$ as well. Therefore, his payoff is the same in both cases, and given that the equilibrium is efficient, the revenue is the same in both separating equilibria.

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.¹⁴ We should however emphasize that the equilibrium play across the two environments is quite distinct. In the unobservable environment, the bidders randomize only in the initial period and then pursue a deterministic policy as a function of their initial randomized choice, whereas in the winner-only observable environment, the bidders continue to randomize in all periods until the uncertainty has been completely resolved.

Explicitly, the payoff $V(q)$ of the high-valuation bidder, given belief q , is given by

$$V(q) = (1 - \delta)(\bar{u} - \bar{b}) + \delta(1 - q)^n(\bar{u} - \underline{u}) = (1 - \delta(1 - \delta) \ln(1 - q)^n)(1 - q)^n(\bar{u} - \underline{u}).$$

¹⁴This follows by direct computation. The unconditional distribution in the winner-only observable case in a given period is a convolution over those distributions in previous periods, because earlier bids affect the belief q , and so it is the later distribution that is used. The details are omitted and available upon request.




Figure 5: Markov equilibria in the winner-only observable case, as a function of q .

Note that the separating equilibrium exists for all parameter values. Indeed, if the other high-valuation bidders separate, it is strictly better for a high-valuation bidder to submit the bid \underline{u}_+ rather than \underline{u} . In either case, the bidder will only win 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 will learn that all other bidders have low valuation, and so he will always win 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 interval $(\underline{u}, \bar{b}_q)$. The details are provided in Appendix. Such equilibria must necessarily converge (in finite time or not) to pooling, which means that the belief must be above the threshold guaranteeing existence of pooling equilibria (characterized by (36)). The following theorem summarizes these findings, as does Figure 5.

Theorem 2. *A (Markov sequential) 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.*

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.

6 Information Disclosure and Equilibrium

We are now in a position to compare the properties of the bidding equilibria across the different disclosure regimes. The structure of the equilibria in the distinct disclosure regimes are determined by the conflicting forces of ratcheting versus information revelation and learning. The private information of each bidder is completely persistent and hence each bidder attempts to use (and hence reveal) his private information optimally over the course of the bidding game. In particular, as part of the equilibrium strategy, the bidder may try to hide and disguise his private information for some time in form of a pooling (or semi-pooling) strategy; this is the logic of the “ratchet effect.” 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 each 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. For example, if pooling were to be a candidate strategy for the bidders, but the bids are unobservable, then the coarse information provided through the allocation decision would only give probabilistic information about the actual bidding strategy employed by the bidders. In fact, in the unobservable bid environment, the information is so coarse that a

pooling equilibrium cannot be supported at all. On the other hand, we saw that as the value of private information becomes more important, both in the sense of a larger discount factor and more competition, then the ratchet effect is eventually going to be strong enough to support a pooling equilibrium with either observable or winner-observable bids. The following corollary then summarizes the impact of the disclosure regime for the emergence of the ratchet effect, based on the earlier analysis.

Corollary (Observable vs. Unobservable Bid Environment).

1. *There does not exist a pooling equilibrium in the unobservable bid environment.*
2. *There exists a pooling equilibrium in the observable and the winner-observable environment for sufficiently large δ and n .*

In other words, the ratchet effect in the bidding environment is weakened by the gains offered by more aggressive strategy. If the bids are unobservable, then the more aggressive strategy, in the form of higher bids cannot even be distinguished from a more defensive strategy and hence undermines the very existence of a pooling equilibrium. The ratchet effect is also weakened if the immediate gains from revealing the private information become relatively more important, i.e. as there is more discounting and/or less competition. In this case, a bidder with a high valuation intends to convey his information through an aggressive bid, and similarly seeks to benefit from the information conveyed through low bids. An environment with less information now suppresses the informational efficiency and hence the competitiveness associated with aggressive bids. In fact, as we compare the revenues from separating equilibria across the different disclosure regimes, we find that additional bid information increases the revenues accruing to the seller.

Corollary (Separating Equilibria).

1. *The revenues in the separating equilibria across the three environments compares as follows; it is largest in the observable case and lower and equal in the other two cases.*
2. *As $\delta \rightarrow 1$, the average revenue in the separating equilibrium in every environment converges from below to the revenue of the static first price auction.*

However, the informational advantage provided by the observable bid environment vanishes as the discount factor becomes large. In other words, the revenue in the unobservable environment also converges against the theoretic optimum provided by the static first price auction under the initial prior distribution of the valuations. The revenue gap between the theoretical optimum and the equilibrium in the environments with restricted bidding information can be attributed to the limited information transmission in these environments. To see this, consider the position of a high valuation bidder who considers an aggressive bid and associated continuation payoffs. In either the unobservable or the winner-observable bid environment, even a winning bid today will not offer him a complete insight into the strength of his competitors. After all, he will not observe their realized bid. More importantly, as he deliberates on the strength of his own bid, he realizes that winning with a lower bid, besides reducing the current bidding cost, provides him with more information about the true valuation of his competitors in the subsequent periods. This naturally depresses his bid relative to the static considerations. In consequence, a high valuation buyer bids distinctly from a low valuation buyer, but not as aggressively as he would in the absence of intertemporal considerations. In abstract terms, this implies that the incentive constraints which separates high from low type is satisfied, but as strict rather than weak inequality, and this gap in turn creates the revenue gap in restricted information environments.

We are thus lead to compare the different disclosure regimes in terms of their impact on the equilibrium strategies. In terms of the efficiency of the allocation, we find that the unobservable environment does unambiguously better in that it does not permit the possibility of a pooling equilibrium with the associated inefficiency in the allocation. In terms of revenues, we found that the revenues in the separating equilibrium are increasing with more available bidding information. But the gains arising from additional information are limited with respect to the possible losses which arise with the pooling equilibrium.

Corollary (Separating vs. Pooling Equilibria).

In every environment, a separating equilibrium achieves a strictly higher revenue than the pooling equilibrium (if it exists).

We should emphasize that the above result explicitly allows the comparison of a separating equilibrium in a given disclosure regime with a pooling equilibrium in a different disclosure regime,

for a constant payoff environment across the regimes. We also observe that the above result can be strengthened in the observable and the winner-observable environment to state that the separating equilibrium achieves a higher revenue than all equilibria subject to Refinement **A**. In the unobservable environment, we do not have a complete characterization of the equilibrium set, although we suspect that the separating equilibrium is the unique equilibrium subject to Refinement **A**.

7 Concluding Remarks

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 unobservable 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. The later policy was shown to share equilibrium properties with both the minimal and the maximal disclosure policy.

We established in Proposition 1 that the policy of minimal information disclosure eliminates the possibility of a pooling equilibrium with low revenues. 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 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.

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.

Similar technical difficulties arise if the number of bidders fluctuates over time, or, more importantly, if there were more than two valuations. 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 to study such richer environments.

Trust in the Auctioneer: As mentioned in the introduction, 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.

8 Appendix

8.1 Unobservable Bids

We generalize here the analysis given in the text to the case with $n + 1$ bidders. In the equilibrium we construct, the loser's bid distribution in period $t \geq 1$ is denoted F_t . The distribution of the highest losing bid is therefore 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

$$F_t(b) = (1 - q) \left(\frac{\bar{u} - \underline{u}}{\bar{u} - b} \right)^{1/n}. \quad (38)$$

This distribution being independent of t (and of δ), it follows that the loser makes a bid that is independent of t , for all $t \geq 1$. ■

Proof of Lemma 2. The analysis here departs somewhat from the case $n = 1$. 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

$$X_t(b) = \max_{\beta} \{ \Pr[i \text{ wins in period } t \text{ with } \beta](1 - \delta)(\bar{u} - \beta) + \delta X_{t+1}(\beta) \} .^{15}$$

We do not condition on player i having always lost before, or the winner having always won (i.e., it could well be that the game is “over” and i might not know about it). Of course 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

$$\frac{X'_t(b)}{1 - \delta} = \frac{d}{db} p_t(b, \beta)(\bar{u} - \beta). \quad (39)$$

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

$$\left(\frac{d}{d\beta} p_t(b, b) + \delta \frac{d}{db} p_{t+1}(b, b) \right) (\bar{u} - b) - p_t(b, b) = 0. \quad (40)$$

¹⁵As 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 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 $b - \varepsilon$, where $\varepsilon > 0$ is small? First, the bid $b - \varepsilon$ must be the highest bid among the constant losing bids, which occurs with probability $F(b)^{n-1}$. Second, the previous bid by the winner 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

$$F(b - \varepsilon)^{n-1}(G_{t-1}(\beta(b - \varepsilon)) - G_{t-1}(b)) = F(b - \varepsilon)^{n-1}(G_t(b - \varepsilon) - G_{t-1}(b)). \quad (41)$$

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

$$F(b + \varepsilon)^{n-1}(G_t(b + \varepsilon) - G_{t-1}(b + \varepsilon)) + \int_b^{b+\varepsilon} F(x)^{n-1} g_{t-1}(x) dx. \quad (42)$$

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

$$p_t(b, b) = F(b)^{n-1}(G_t(b) - G_{t-1}(b)), \quad (43)$$

and

$$\frac{d}{d\beta}p_t(b, b) = F(b)^{n-1}g_t(b) + (n-1)f(b)F(b)^{n-2}(G_t(b) - G_{t-1}(b)). \quad (44)$$

We also get, from (41) or (42),

$$\frac{d}{db}p_t(b, b) = -F(b)^{n-1}g_{t-1}(b). \quad (45)$$

We must now solve the resulting differential equation. 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 (43)–(45) 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

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

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 (38),

$$f(b) = \frac{1}{n}(1 - q) \left(\frac{\bar{u} - u}{\bar{u} - b} \right)^{\frac{1}{n}} \frac{1}{\bar{u} - b}, \quad (47)$$

and hence

$$\frac{f(b)}{F(b)} = \frac{1}{n} \frac{1}{\bar{u} - b}.$$

Plugging into (46),

$$\left((1 - \delta)g_t(b) + \frac{n-1}{n} \frac{1}{\bar{u} - b} (G_t(b) - G_{t-1}(b)) \right) (\bar{u} - b) - (G_t(b) - G_{t-1}(b)) = 0, \quad (48)$$

or

$$\left((1 - \delta) g_t(b) (\bar{u} - b) - \frac{1}{n} (G_t(b) - G_{t-1}(b)) \right) = 0. \quad (49)$$

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 $G_t(\bar{u} - \gamma) = 1$. Let us define $y := (\bar{u} - b)/\gamma$, and for all t , $H_t(y) = G_t(b)$, so that, for $t \geq 2$,

$$(1 - \delta) n y h_t(y) + H_t(y) - H_{t-1}(y) = 0,$$

and also

$$(1 - \delta) n y h_1(y) + H_1(y) - y^{-\frac{1}{n}} = 0,$$

where $H_t(1) = 1$, for $t \geq 1$.

The solution is

$$H_t(y) = \frac{y^{-1/n}}{\delta^t} + y^{-1/(1-\delta)n} \sum_{\tau=0}^t \frac{1 - \delta^{\tau-t}}{\tau!} \left(\frac{\ln y}{(1-\delta)n} \right)^\tau, \quad (50)$$

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. ■

Proof of Lemma 3. It remains to determine F_0 , the bid distribution used in the initial period. To

do so, observe that 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)

$$F_0(b)^n(\bar{u} - b) - n\delta(1 - q)^n(\bar{u} - \underline{u}) \ln F_0(b) = K_0, \tag{51}$$

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

$$K_0 = (\bar{u} - \underline{u})(1 - q)^n(1 - n\delta \ln(1 - q)). \tag{52}$$

By plugging $b = \bar{b}$, where \bar{b} denotes the upper extremity of the support of F_0 , we also get that $\bar{b} = \bar{u} - K_0$. Equations (52) and (51) uniquely characterize F_0 , and the closed-form solution follows from standard properties of the Lambert function. ■

8.1.1 Comparative Statics of Subsection 3.4.1

Comparative statics of G_t are more easily computed in terms of the distribution H_t , given by (50). In particular, the second term can be re-written as

$$\frac{\Gamma\left(t+1, \ln y^{\frac{1}{(1-\delta)^n}}\right) - \delta^{-t} y^{-1/n} \Gamma\left(t+1, \ln y^{\frac{\delta}{(1-\delta)^n}}\right)}{\Gamma(t+1)},$$

where $\Gamma(t)$ is the ordinary gamma function, and $\Gamma(t, s)$ is the incomplete gamma function $\Gamma(t, s) = \int_s^\infty u^{t-1} e^{-u} du$.

All the comparative statics follow. Because $\Gamma(t, s) \rightarrow 0$ as $s \rightarrow \infty$, it follows that the second term tends to zero as $\delta \rightarrow 1$. Therefore,

$$\lim_{\delta \rightarrow 1} G_t(b) = F(b),$$

which is equal to the distribution of the losers. Comparative statics with respect to n follow since n only affects G_t through the $F(b) = y^{-1/n}$ term.

Consider now the case in which we interpret $\delta \rightarrow 1$ as the result of an increase in the frequency at which successive auctions take place, and fix some time $t = k\Delta$, where Δ is the (small) delay between two successive auctions, so that $\delta = e^{-r\Delta} \simeq 1 - r\Delta$, where $r > 0$ is the fixed discount rate. So we have

$$\frac{\Gamma\left(\frac{t}{\Delta} + 1, \frac{\ln y}{nr\Delta}\right) - y^{-1/n} (1 - r\Delta)^{-\frac{t}{\Delta}} \Gamma\left(\frac{t}{\Delta} + 1, \frac{1 - \Delta r}{\Delta} \frac{\ln y}{rn}\right)}{\Gamma\left(\frac{t}{\Delta} + 1\right)}.$$

Let us now define $\psi := (\ln y)/rnt$. Because

$$\lim_{\Delta \rightarrow 0} (1 - r\Delta)^{-\frac{t}{\Delta}} = e^{rt},$$

it follows that

$$H_t(y) \rightarrow e^{rt} y^{-1/n} + \lim_{x \rightarrow \infty} \frac{\Gamma(x+1, \psi x) - y^{-1/n} \left(1 - \frac{tx}{x}\right)^{-x} \Gamma(x+1, \psi(x-rt))}{\Gamma(x+1)},$$

where $x := t/\Delta$. It follows from standard expansions that the normalized incomplete gamma function $Q(x+1, \psi x) := \Gamma(x+1, \psi x)/\Gamma(x+1)$ satisfies

$$\lim_{x \rightarrow \infty} Q(x+1, \psi x) = \begin{cases} 1 & \text{if } \psi < 1, \\ 1/2 & \text{if } \psi = 1, \\ 0 & \text{if } \psi > 1. \end{cases}$$

and the same holds for $Q(x+1, \psi(x-\kappa))$, for fixed κ, ψ . (see Tricomi (1950)). We thus obtain

$$\lim_{\Delta \rightarrow 0} H_t(y) = \begin{cases} 1 & \text{if } \psi < 1, \\ e^{rt} y^{-\frac{1}{n}} & \text{if } \psi > 1. \end{cases}$$

So we obtain that the limit distribution G_t is as follows. Define $T := \frac{1}{r} \ln \frac{1}{1-q}$.

$$\lim_{\Delta \rightarrow 0} G_t(b) = \begin{cases} e^{rt} (1-q) \left(\frac{\bar{u}-\underline{u}}{\bar{u}-b}\right)^{\frac{1}{n}} & \text{if } b \in (\underline{u}, \bar{u} - (\bar{u} - \underline{u})(1-q)^n e^{nrt}) \text{ and } t < T, \\ 1 & \text{if } b > \underline{u} \text{ otherwise.} \end{cases}$$

8.2 Observable Bids

This subsection provides some details on semi-pooling equilibria in the environment with observable bids. This requires, first of all, a better understanding of the payoff V^U of the uninformed bidder in the game of one-sided incomplete information. Let throughout $r := 1 - q$. Recall that

$$F_s(\underline{u}) = \delta F_{s+1}(\underline{u}) = \dots = F_T(\underline{u}) n^{T-s} \delta^{\sum_{j=0}^{T-s-1} n^j} = F_T(\underline{u}) n^{T-s} \delta^{\frac{n^{T-s}-1}{n-1}}.$$

So

$$\times_{t=s+1}^T F_t(\underline{u}) = F_T(\underline{u}) \sum_{t=s+1}^T n^{T-t} \delta^{\sum_{t=s+1}^T n^{\frac{T-t-1}{n-1}}} = F_T(\underline{u}) \frac{n^{T-s-1}}{n-1} \delta^{\frac{n^{T-s-1}}{(n-1)^2} - \frac{T-s}{n-1}}.$$

The payoff V^U can be derived under the assumption that the uninformed bidder bids \underline{u}_+ in all periods. In that case, he obtains

$$(1 - \delta) \sum_{s=0}^T \delta^s \times_{j=0}^s F_j(\underline{u})^n (\bar{u} - \underline{u}) + \delta^{T+1} r^n (\bar{u} - \underline{u}).$$

Because $\times_{j=0}^s F_j(\underline{u}) = r / \times_{t=s+1}^T F_t(\underline{u})$, we get

$$V^U = r^n \left[1 + (1 - \delta) \sum_{s=0}^{T-1} \left(F_T(\underline{u})^{-\frac{n}{n-1}(n^{T-s}-1)} \delta^{s - \frac{n}{(n-1)^2}(n^{T-s}-1) + \frac{n}{n-1}(T-s)} - 1 \right) \right].$$

For future purposes, recall that, because $\times_{t=0}^T F_t(\underline{u}) = r$, $F_T(\underline{u})$ is increasing in r , and so V^U/r^n is decreasing in r over each interval of values of r for which T is constant. But since V^U is continuous at values of r for which $F_T(\underline{u}) = \delta$, it follows that V^U/r^n is decreasing in r (i.e. increasing in q) for all values of r . Rather surprisingly, it can further be shown that $(V^U - r^n)/(1 - \delta)$ is bounded above, uniformly in δ , but we shall not need this in the sequel.

We may now turn to the analysis of semi-pooling equilibria in the game of symmetric information. In such an equilibrium, at least in some initial phase, each high-valuation bidder mixes between submitting the bid \underline{u} , and continuously randomizing over some interval $(\underline{u}, \bar{b}_q)$.

In particular, each high-valuation bidder must be indifferent between submitting the bid \underline{u} and \underline{u}_+ . Suppose that this is the case in period t , and that high-valuation bidders are willing to submit a bid above \underline{u} in period $t+1$ (assuming no player has done so in the past). Then it must hold that

$$\begin{aligned} & F_t(\underline{u})^n \left(\frac{1 - \delta}{n + 1} (\bar{u} - \underline{u}) + \delta F_{t+1}(\underline{u})^n ((1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q_{t+2})) \right) \\ & + \delta n F_t(\underline{u})^{n-1} (1 - F_t(\underline{u})) V^I(q_{t+1}) = F_t(\underline{u})^n ((1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q_{t+1})). \end{aligned} \quad (53)$$

Here, q_{t+1} and q_{t+2} (and q_t , used below) denote the beliefs of the players in the corresponding period, given that information is still symmetric, i.e. all players have always bid \underline{u} so far. To understand this equality, consider the right-hand side first. This is the payoff from bidding \underline{u}_+ . In that case, the high-valuation bidder (say, player i) wins something only in the event that all other bidders bid \underline{u} (if even one other bidder bids more, bidder i will lose now and because there will

be common knowledge that there are two high-valuation bidders, continuation payoffs are zero as well in this event. If all other bidders bid \underline{u} , which occurs with probability $F_t(\underline{u})^n$, he gets $1 - \delta$ in the current period, and the payoff $V^U(q_{t+1})$ as the uninformed player in a game of one-sided incomplete information in which informed bidders have high valuations with probability q_{t+1} . The left-hand side is the payoff from bidding \underline{u} in period t . If all other bidders do so as well (first term), then player i wins in the current period with probability $1/(n+1)$, and his continuation payoff can be computed under the assumption that he bids \underline{u}_+ in the following period, in which case his continuation payoff is given by the same formula as the right-hand side that we have just reviewed, with a shift in the relevant period. In case exactly one bidder bids above \underline{u} , player i also gets some continuation payoff (second term), namely, he gets the payoff of an informed, high-valuation bidder in the game with one-sided incomplete information.

We can use the fact that $F_t(\underline{u}) = (1 - q_t)/(1 - q_{t+1})$ (and the corresponding formula for $t+1$) to eliminate the distributions. Dividing through by $(1 - \delta)/(\delta(\bar{u} - \underline{u}))$ and re-arranging then gives

$$\begin{aligned} & \left(\frac{1 - q_{t+1}}{1 - q_{t+2}} \right)^n (1 - v^U(q_{t+2})) + n \frac{q_t - q_{t+1}}{1 - q_t} \frac{v^I(q_{t+1})}{1 - \delta} + \frac{(1 - q_{t+1})^n}{1 - \delta} \left(\frac{v^U(q_{t+2})}{(1 - q_{t+2})^n} - \frac{v^U(q_{t+1})}{(1 - q_{t+1})^n} \right) \\ = & \frac{n}{n+1} \frac{1}{\delta}, \end{aligned} \quad (54)$$

where $v^U(q) := V^U(q)/(\bar{u} - \underline{u})$, $v^I(q) := V^I(q)/(\bar{u} - \underline{u})$. Because, as we have seen, $V^U(q)/(1 - q)^n$ is increasing in q , yet $q_{t+2} \leq q_{t+1}$, the last term of the left-hand side is negative, so that it must be the case that

$$\left(\frac{1 - q_{t+1}}{1 - q_{t+2}} \right)^n (1 - v^U(q_{t+2})) + n \frac{q_t - q_{t+1}}{1 - q_t} \frac{v^I(q_{t+1})}{1 - \delta} \geq \frac{n}{n+1} \frac{1}{\delta}.$$

Further, we also know that $v^U(q_{t+2})/(1 - q_{t+2})^n \geq 1$, as well as $v^I(q_{t+1})/(1 - \delta) \leq (1 - q_{t+1})^n$, so *a fortiori* we must have

$$1 - (1 - q_{t+1})^n + n \frac{q_t - q_{t+1}}{1 - q_t} (1 - q_{t+1})^n \geq \frac{n}{n+1} \frac{1}{\delta}.$$

The left-hand side is decreasing in q_{t+1} .¹⁶ So it must be the case that, plugging in $q_{t+1} = 0$ in the last inequality,

$$n \frac{q_t}{1 - q_t} \geq \frac{n}{n+1} \frac{1}{\delta}, \text{ or } q_t \geq \frac{1}{1 + (n+1)\delta},$$

and so we need at the very least that $q \geq \frac{1}{1+(n+1)\delta}$, which is the boundary value for the existence of separating equilibria. Obviously, if $q \leq \frac{1}{1+(n+1)\delta}$, it was without loss that we assumed that bidding more than \underline{u}_+ was optimal for the high-valuation bidder, as we know that pooling is not an equilibrium outcome for such values of the prior belief. This establishes that there cannot be semi-pooling equilibria if (or once) $q \leq \frac{1}{1+(n+1)\delta}$.

Suppose now that semi-pooling ends in pooling, i.e., there exists a period $t + 1$ after which pooling on the bid \underline{u} is optimal, provided no bidder has submitted a higher bid ever. Then we must have, in period t ,

$$F_t(\underline{u})^n \frac{\bar{u} - \underline{u}}{n+1} + \delta n F_t(\underline{u})^{n-1} (1 - F_t(\underline{u})) V^I(q_{t+1}) = F_t(\underline{u})^n ((1 - \delta)(\bar{u} - \underline{u}) + \delta V^U(q_{t+1})),$$

by the same logic as above. Re-arranging, this is equivalent to

$$F_t(\underline{u}) \left((1 - \delta) \frac{n}{n+1} + \delta v^U(q_{t+1}) - \frac{\delta}{n+1} \right) = \delta n (1 - F_t(\underline{u})) v^I(q_{t+1}).$$

However, because pooling must be an equilibrium from period $t + 1$ onward, (30) provides an upper bound on $v^U(q_{t+1})$. Substituting it into the left-hand side, we obtain that the left-hand side is non-positive, and therefore the equality can only hold if $F_t(\underline{u}) = 1$, contradicting the fact that period $t + 1$ was the first period in which there was pooling.

Therefore, semi-pooling can never end into pooling, and so high-valuation bidders must be willing to separate in every period, as it was assumed, for period $t + 1$, in (53).

¹⁶The derivative with respect to q_{t+1} is $n(n+1)(1 - q_{t+1})^{n-1}(q_{t+1} - q_t)/(1 - q_t) \leq 0$.

Existence Recall that, from (54), a semi-pooling equilibrium must satisfy

$$\begin{aligned} & \left(\frac{1 - q_{t+1}}{1 - q_{t+2}} \right)^n (1 - v^U(q_{t+2})) + n \frac{q_t - q_{t+1}}{1 - q_t} \frac{v^I(q_{t+1})}{1 - \delta} + \frac{(1 - q_{t+1})^n}{1 - \delta} \left(\frac{v^U(q_{t+2})}{(1 - q_{t+2})^n} - \frac{v^U(q_{t+1})}{(1 - q_{t+1})^n} \right) \\ &= \frac{n}{n + 1} \frac{1}{\delta}, \end{aligned}$$

Note that the left-hand side is increasing in q_t . Also, the terms involving q_{t+2} are

$$\left(\frac{1 - q_{t+1}}{1 - q_{t+2}} \right)^n \left(1 - v^U(q_{t+2}) + \frac{v^U(q_{t+2})}{1 - \delta} \right) = (1 - q_{t+1})^n \left(\frac{1}{(1 - q_{t+2})^n} + \frac{\delta v^U(q_{t+2})}{(1 - \delta)(1 - q_{t+2})^n} \right),$$

and since both terms are increasing in q_{t+2} , the left-hand side of (54) is also increasing in q_{t+2} .

Note now that, when $q_{t+1} = q_{t+2}$, the left-hand side reduces to

$$1 - v^U(q_{t+2}) + n \frac{q_t - q_{t+1}}{1 - q_t} \frac{v^I(q_{t+1})}{1 - \delta}.$$

Note that the first two terms are less than the right-hand side, $n/(\delta(n + 1))$, if and only if $q_{t+2} \leq \bar{q}$. It is then clear that, if $q_{t+1} = q_{t+2} < \bar{q}$, there exists a unique solution q_t to (54), and this solution is such that $q_t > q_{t+1} = q_{t+2}$; in fact, it is bounded above the common value $q_{t+1} = q_{t+2}$. Consider now any pair $q_{t+2} < q_{t+1}$, with $q_{t+1} \in (\underline{q}, \bar{q})$. Fixing q_{t+1} , since q_t must increase as q_{t+2} decreases (because the right-hand side of (54) is increasing in both of them), and since q_t is already larger than q_{t+1} for the choice $q_{t+2} = q_{t+1}$, it follows that there always exists $q_t > q_{t+1}$ solving (54), and that, proceeding inductively, such a sequence $q_t, q_{t-1}, q_{t-2}, \dots$ cannot converge to a value strictly below \bar{q} (either it converges to \bar{q} or it eventually exceeds \bar{q}). This, alongside with continuity of the solution q_t of (54), establishes existence of the semi-pooling equilibrium for initial values of q in (\underline{q}, \bar{q}) .

8.3 Winner-Only Observable Bids

Proof of Proposition 4. (1.) We start with the informed bidder. In the winner-only observable environment, the largest bid in the support is

$$\bar{b}_{q_0} = \bar{u} - (1 - q_0)^n(\bar{u} - \underline{u}). \quad (55)$$

The informed bidder can therefore win with probability 1 with the bid (55) in a single period. The net utility is therefore

$$\bar{u} - \bar{b}_{q_0} = (1 - q_0)^n(\bar{u} - \underline{u}).$$

In the observable environment, the largest bid in the support is

$$\bar{b}_{q_0} = \bar{u} - \left(\frac{1 - q_0}{1 - q_1}\right)^n (\bar{u} - \underline{u}). \quad (56)$$

The informed bidder can therefore win with probability 1 with the bid (56) in a single period. The net utility is therefore

$$\bar{u} - \bar{b}_{q_0} = \left(\frac{1 - q_0}{1 - q_1}\right)^n (\bar{u} - \underline{u}).$$

Since $1 - q_1 < 1$, it follows that $\frac{1 - q_0}{1 - q_1} < (1 - q_0)$.

We then consider the uninformed bidder. In the winner observable environment, the largest bid in the support is

$$\bar{b}_{q_0} = \bar{u} - (1 - q_0)^n(\bar{u} - \underline{u}).$$

The uninformed bidder can therefore win with probability 1 with the bid (55) in every single period. The net utility is therefore

$$\bar{u} - \bar{b}_{q_0} = (1 - q_0)^n(\bar{u} - \underline{u}). \quad (57)$$

In the observable environment we can describe the equilibrium value of the uninformed agent as follows:

$$(\bar{u} - \underline{u}) \left(\sum_{t=0}^T \delta^t \left(\prod_{s=0}^t (F_s(\underline{u}))^n \right) + \frac{\delta^{T+1}}{1 - \delta} \prod_{s=0}^{T+1} (F_s(\underline{u}))^n \right).$$

The equilibrium value is obtained by the following strategy. The uninformed bidder bids \underline{u}^+ in every period t . The uninformed bidder wins in period t with a positive net utility $(\bar{u} - \underline{u})$ if and only if the informed bidders bid \underline{u} to hide their value. Now using the updating rule

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

or

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

it follows that the ex ante winning probability in period t is

$$\prod_{s=0}^t (F_s(\underline{u}))^n = \left(\frac{1 - q_0}{1 - q_{t+1}} \right)^n, \quad (58)$$

and from $t \geq T$, this yields:

$$\prod_{s=0}^t (F_s(\underline{u}))^n = \frac{1 - q_0}{1 - q_{t+1}} = (1 - q_0)^n, \quad (59)$$

since for $t \geq T$, $1 - q_{t+1} = 1$. But now we can compare (57) with (58) and (59) and find that the ex-ante probability of receiving the net utility $(\bar{u} - \underline{u})$ is everywhere (weakly) higher in the observable environment than in the bidder-only observable environment, as

$$\frac{1 - q_0}{1 - q_{t+1}} \geq 1 - q_0,$$

and the weak inequality holds as equality for all $t \geq T$.

(2.) The equilibrium allocation is efficient in both environments, and hence generate the same social value. The revenue comparison across the two environments then follows directly from the comparison of the residual values received by the informed and uninformed bidders. ■

Finally, we provide some details on semi-pooling equilibria with winner-only observable bids. In the equilibrium the high-valuation bidders assign positive probability to \underline{u} , but also continuously randomize over some interval $(\underline{u}, \bar{b}_q)$. Clearly, if q is very low given δ , the high-valuation bidder

is not willing to make the bid \underline{u} (as he believes that bidding \underline{u}_+ guarantees him almost the upper bound on his equilibrium payoff, namely $\bar{u} - \underline{u}$). So either the candidate semi-pooling converges to the pooling equilibrium, or there is a last period after which, if the winning bid is once again \underline{u} , the equilibrium played is the separating one.

However, semi-pooling cannot end up in separation. Suppose so for the sake of contradiction, and consider the payoff from bidding \underline{u}_+ , relative to \underline{u} , in the last period in which the high-valuation bidder assigns positive probability to the bid \underline{u} . We might as well condition on the other bidders bidding \underline{u} , for otherwise the continuation payoff is the same in both cases. In such a candidate equilibrium, indifference between these two bids requires

$$(1 - \delta) \frac{\bar{u} - \underline{u}}{n + 1} + \delta V \left(1 - \frac{1 - q}{F_q(\underline{u})} \right) = (1 - \delta)(\bar{u} - \underline{u}) + \delta V^U \left(1 - \frac{1 - q}{F_q(\underline{u})} \right), \quad (60)$$

where the left-hand side (respectively right-hand side) is the payoff from bidding \underline{u} (respectively \underline{u}_+). Rearranging gives

$$-\delta^2 \left(\frac{1 - q}{F_q(\underline{u})} \right)^n \ln \left(\frac{1 - q}{F_q(\underline{u})} \right)^n = \frac{n}{n + 1}.$$

However, the left-hand side is at most $\delta^2 e^{-1} < 1/2 \leq n/(n + 1)$, which gives the desired contradiction. A candidate semi-pooling equilibrium must thus converge to pooling, and the corresponding belief must be such that a pooling equilibrium exists. It is easy to see that we can solve (60) for $F_q(\underline{u})$ if q is above this threshold. In that case, in the last period before pooling (if there is such a period), (60) gives

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

which admits a solution $F_q(\underline{u}) \in [1 - q, 1]$ if and only if q is above the threshold that defines semi-pooling equilibria, and characterized by (36). This shows that semi-pooling equilibria can then be constructed in this range of parameters (obviously, more complicated semi-pooling equilibria involving many periods of semi-pooling can be constructed as well).

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.
- 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.
- 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.
- ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT (2008): “Fighting Cartels in Public Procurement,” *OECD Policy Brief*, pp. 1–8.
- 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.
- TRICOMI, F. (1950): “Asymptotische Eigenschaften der Unvollstaendigen Gammafunktion,” *Mathematische Zeitschrift*, 53, 136–148.
- 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.