Private Information in Sequential Common-Value Auctions*

Johannes Hörner† Julian Jamison‡

March 17, 2006

Abstract

We study an infinitely-repeated first-price auction with common values. Initially, bidders receive independent private signals about the objects' value, which itself does not change over time. Learning occurs only through observation of the bids. Under one-sided incomplete information, this information is eventually revealed and the seller extracts essentially the entire rent (for large discount factors). Both players' payoffs tend to zero as the discount factor tends to one. However, the uninformed bidder does relatively better than the informed bidder. We discuss the case of two-sided incomplete information, and argue that, under a Markovian refinement, the outcome is pooling: information is revealed only insofar as it does not affect prices. Bidders submit a common, low bid in the tradition of "collusion without conspiracy".

Keywords: repeated game with incomplete information; private information; ratchet effect; first-price auction; dynamic auctions

JEL codes: C72, D44

---

*We thank In-Koo Cho, Martin Cripps, Olivier Gossner, Mike Peters, Jeroen Swinkels, Joel Watson, Bob Weber, the managing editor, two anonymous referees, and seminar participants at Northwestern and several conferences (CETC, SED, NASMES, ESEM) for useful feedback, as well as Marcin Peski for research assistance.

†Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, 2001, Sheridan Rd., Evanston, IL 60208-2001: `j-horner@kellogg.northwestern.edu,` *and* HEC School of Management, Paris, *and* CEPR.

‡University of California, Berkeley and San Francisco: `julison@berkeley.edu`

1 Introduction

Imagine an auction in which several lots of 1961 Chateau Palmer, a red Bordeaux, are successively put up for sale.[1] The value of these lots is identical, but not commonly known. Bidders include both wine experts and less experienced buyers. All bids are revealed after each sale. How should buyers behave? What is the pattern of prices over time? How much is each bidder's proprietary information worth? How well does the seller perform?

To provide some insights into these questions, we analyze an infinitely-repeated game between two players. In each period, a single object is sold via a first-price sealed-bid auction. All objects have the same value. Each player receives a private signal once (at the beginning of time) that is correlated with the common value. All bids are observed, so it is a Bayesian game (i.e. with incomplete information only). Players maximize average discounted payoffs, and the analysis primarily focuses on the case of low discounting.

In a static framework, proprietary information plays two roles: it helps determine one's own value for the unit being sold, and it helps predict the other bidders' actions. The value of this information, and its effect on equilibrium bids and profits, has been subject to much analysis and is by now well-understood when the setting is symmetric, or when only one bidder has private information (Milgrom and Weber (1982); Engelbrecht-Wiggans, Milgrom, and Weber (1983)).

In the multi-period case, bids transmit information that is relevant for future play, allowing bidders to refine their estimate of the units' value, as well as to better predict their opponents' future bids. Thus, bidders face a trade-off: if a bidder's proprietary information indicates that the units are all of high quality, he would like to increase his chance of winning the object by bidding more aggressively, but doing so may prove costly later on, as other bidders may compete away the value of the information released through the bid.

In light of this intuition, we focus on several attributes of the model. In particular, we are interested in the bidders' equilibrium strategies *per se*; in the implied value of private information within such a model; and in the resulting implications for the seller's revenue.

We are not, however, interested in the rich collusive opportunities inherent to repeated interactions, which are present in our game just as in standard repeated games. Therefore, we restrict attention to equilibria in which strategies should depend only on beliefs (about the signals or types of all players) and not on payoff-irrelevant details of histories. In the case of one-sided incomplete information, this is particularly easy to achieve, by requiring the uninformed player to be "myopic" and maximize his flow payoff. In effect, we then consider the game between an informed, long-run bidder with proprietary information and an infinite sequence of uninformed, short-run bidders. Our results are most striking in this simple framework:

When only one bidder has proprietary information, all information is revealed in finite time, although the time required grows without bound as the discount factor tends to one. The informed player's average discounted total payoff then tends to one; further, in any undominated equilib-

[1] See Ashenfelter (1989) for more detail on this example and others.

rium, the uninformed player's average discounted total payoff converge to zero (as the discount factor tends to one), while the auctioneer's revenue converges to his first-best profit. In addition, the informed bidder's total payoff is small relative to that of the uninformed bidder.

That is, in any such equilibrium, the informed bidder discloses all of his proprietary information in finite time. The delay increases with the bidders' patience, but information revelation occurs "quickly" relative to the discount factor. Once information is revealed, bids equal the commonly known value of the object, and continuation payoffs are therefore zero (just as in Bertrand competition with symmetric costs). So, provided bidders are sufficiently patient, both bidders do poorly overall and the seller fares very well. While both bidders' payoffs converge to zero, the payoff of the informed bidder does so faster than that of the uninformed.

The intuition for this result starts with the fact that the uninformed bidder's actions have no informational content. Therefore, he has an incentive to break any ties in his favor. This leaves an informed bidder who received an optimistic signal no choice but to eventually bid strictly more than what he would have bid if his private signal was bad news, thereby eventually revealing some of his private information. The informed bidder's temporal trade-off is that by delaying this revelation, he is able to win with a lower bid, as pessimism makes the uninformed bidder more cautious in the future.

We show that none conclusion holds without the refinements used, myopia and weak dominance. We also show that our results remain valid when the horizon is long, but finite, even without discounting, generalizing thereby a result by Engelbrecht-Wiggans and Weber (1983) discussed below, or when there is more than one uninformed player.

We provide additional detail and intuition in the particularly simple case of binary signals. In line with the empirical literature, we find that equilibrium winning bids tend to decline until the information is fully revealed.

When there is two-sided uncertainty, we show by means of examples that similarly precise predictions require much stronger refinements. Because those refinements are less obvious, and their consequences more predictable, we restrict ourself to a discussion of the results, and refer the interested reader to the working paper for additional details. In that case as well, we can show:

When both bidders have proprietary information (and discounting is low), bidders' profits are at least as large as if they were both repeatedly bidding the lowest expected value that any type of either player could possibly have for the object.

That is, the average equilibrium bid is less than or equal to the most pessimistic expectation that any type of either player could have. The intuition for this result is very similar to the ratchet effect identified in games with one-sided incomplete information by Freixas, Guesnerie and Tirole (1985) and Hart and Tirole (1988).[2] Although the equilibrium payoffs are as if all

[2]Our result validates Laffont's intuition that "in this type of dynamic game a ratchet effect often happens as the bidders in the auction want to hide some of their private information to be in a more favorable position in the next stages." (Laffont 1997)

bids were "pooling", we note that almost all information may potentially still be disclosed in equilibrium. Sellers should therefore tend to prefer a one-time batch sale (replicating the static outcome) to a sequential auction if bidders are patient enough, or alternatively if the frequency of auctions is high enough. It is however possible for this ordering to be reversed, even in a symmetric version of the model. Loosely, this is because there is no winner's curse in such a pooling equilibrium: winning has no adverse implications for beliefs.

In order to facilitate comparison, we begin with the static case of the model. Here we completely characterize the equilibrium, which has not previously appeared. In this one-shot environment, bids fully reveal players' types. More precisely, successive types of a given player continuously randomize over contiguous (but non-overlapping) intervals. We solve explicitly for the expected revenue of the seller. Although our conclusion will be that the dynamic auction induces low revenue levels, it is interesting to observe that for some prior beliefs the auctioneer would still prefer a sequential to a static (batch-sale) auction; see Section 5.

Finally, we consider the special case of binary types (i.e. high or low signals only) in substantially greater detail. In this setting we give explicit formulations of the bidding functions and their (stochastic) evolution over time in the case of both one-sided and two-sided incomplete information. Furthermore, if the prior beliefs in the two-sided case are the same (i.e. symmetric across bidders), we characterize the equilibrium for all discount factors. The equilibrium is fully separating – meaning high types make themselves known immediately – if players are impatient and the prior is pessimistic; it is fully pooling if players are patient and the prior is pessimistic; and it is semi-pooling otherwise, meaning that high types separate with non-degenerate probability in a given period.

While it is outside the scope of this paper to fully develop the implications of our model for auction design, in Section 6 we briefly describe some of the different instruments that are available to the seller.

1.1 Literature

Related literature begins with an important early paper of Ortega-Reichert (1967), which describes the symmetric, pure-strategy equilibrium of a two-object, two-bidder, first-price sequential auction in a private value model. Bidders' signals for the two objects are independently distributed conditionally on a state variable whose distribution is commonly known. Because the second item's signal is only revealed after the first sale, there is no ratchet effect and the bidders' private information is revealed by their first bid.

Engelbrecht-Wiggans and Weber (1983) investigate what is a special case of our model.[3] They analyze a finite-object, common-value sequential auction between a perfectly informed bidder and an uninformed bidder in which the units' value – high or low – is perfectly correlated over time.

[3] We learned a great deal from their analysis.

As in our model, because the informed bidder cannot exploit his information without revealing it for the next stages, the uninformed bidder's payoff exceeds the informed bidder's payoff when the horizon is long enough. If it is likely that units are of high value, a one-stage batch sale generates a higher expected revenue than the sequential procedure.

Hausch (1986) re-examines Ortega-Reichert's model when the single private signal is received *before*, rather than after the first auction. For some parameters, players have an incentive to use mixed strategies to hide some of their information. As a consequence, simultaneous sales may be better for the seller.

Bikhchandani (1988) examines a finitely repeated, two-bidder *second-price* auction with common values. The players' signals are additive and identically and independently drawn from a common distribution. Players' signals are conditionally independent. One of the bidders has a privately known type. The strong type's signal is higher than the ordinary type's signal. Because winning is especially bad news for the uninformed player when his opponent is of the high type, the winner's curse is intensified for the uninformed player, forcing him to submit lower bids in equilibrium. This weakens the winner's curse for the informed bidder of the ordinary type and decreases his serious bids. The ordinary type has therefore an initial incentive to mimic the strong type. Similarly, Kwiek (2002) considers the effect of reputation on bidding in second-price auctions. Reputation affects the perception of a player's bidding behavior, but does not affect the value of the object to both bidders. Therefore, it is quite different from private information as understood in this paper.

In a recent paper, De Meyer and Saley (2002) analyze the limit of a finitely repeated first-price auction with one informed bidder and one uninformed bidder. There are two possible states of the world, and the players transact directly with each other (so it is a zero-sum game). They show that the price distribution asymptotically approximates a Brownian motion. Kyle (1985) examines the informational content of prices and the value of information to an insider in a model of sequential auctions in which buyers choose quantities and the price clears the market. To prevent the (unique) insider's actions from being perfectly informative, and to guarantee that an equilibrium in pure strategies exists, he introduces noise trading. In contrast, we are able (without adding noise) to directly characterize the mixed strategy equilibrium arising when incomplete information is one-sided.

Hon-Snir, Monderer and Sela (1998) study a repeated first-price auction with independent private values, though they do not assume that players are fully rational (noting that "equilibrium analysis of repeated first-price auctions in the framework of repeated games with incomplete information is complex"). Instead they consider a broad class of learning schemes, and they are able to show that information is eventually fully revealed. Similarly, Oren and Rothkopf (1975) develop a behavioral model of how competitors react to solve for the bidding strategies.

Virág (2003) studies a first-price finitely repeated common value auction between two players, with one-sided incomplete information. He assumes that bids are not observed, only the identity of the winner. In addition, the uninformed bidder finds out the exact value of the object as soon as he wins it once. Therefore, information is complete once the uninformed bidder has won, and it follows that dynamics are very different from the ones in this paper. Abdulkadiroğlu and Chung (2003) assume that bidders collude and investigate the resultant mechanism design problem faced by the auctioneer.

Finally, a general analysis of repeated games with incomplete information is carried out by Aumann, Maschler and Stearns (1966-68).

The empirical literature on repeated auctions is extensive. Beginning with Ashenfelter (1989), many papers have documented the declining-price anomaly in sequential auctions of identical objects.[4] Van den Berg, van Ours and Pradhan (2001) study a repeated first-price auction of Dutch roses. Ashenfelter (1989) and Ginsburgh (1998) investigate wine auctions. Pezanis-Christou (2000) reports an analysis of price behavior in a sequence of fish auctions with retailers and wholesalers who are differentially informed (as in our model) and in which there is uncertain supply (equivalent to impatience). Collusion in repeated auctions, including in the California electricity markets and other areas, has also been subject to many studies. For instance, Baldwin, Marshall, and Richard (1997) report collusion at timber sales, while Pesendorfer (2000) suggests evidence of collusive behavior in first-price auctions for milk.

1.2 Outline of the paper

Section 2 presents the formal description of our model. Section 3 details the equilibrium for the static case ($\delta = 0$); this is not required in the remainder. Section 4 characterizes the equilibrium with one-sided incomplete information, and introduces an example for illustrative purposes. Section 5 considers two-sided incomplete information and includes the main result. It also extends the example, for arbitrary discount factors. Finally, Section 6 concludes. Proofs of the results are relegated to an appendix.

2 The Model

This paper considers an infinitely-repeated game between two risk-neutral bidders with quasilinear utility, player 1 and player 2. In every period $t = 0, 1, \ldots$ one indivisible unit is sold using a sealed-bid first-price auction. In case of a tie, each player wins with equal probability. We assume common values:

Bidders value each unit identically.

All units auctioned off have the same value, represented by the random variable V. That is, we assume perfect correlation over time:

The value of the units from one period to the next is unchanging.

[4]Note, however, that Raviv (2003) actually finds evidence for increasing prices in certain sequential auctions.

Before the game starts, each bidder i ($i \in \{1,2\}$) receives a signal S_i concerning the value of the object. The signal can take $m+1$ different values for player 1, and $n+1$ different values for player 2, that is, $S_1 \in M = \{0,1,\ldots,m\}$ and $S_2 \in N = \{0,1,\ldots,n\}$. Signals are also referred to as types. Players never observe the realization of V. Thus, conditional on signals $S_1 = j$ and $S_2 = k$, we may safely identify V with its expectation, and we denote this value by $v(j,k)$. A high (low) signal is statistical evidence for a high (low) value of the object. This is formalized by:

The valuation $v(j,k)$ is strictly increasing in each of its arguments. The variables S_1 and S_2 are themselves independent.[5]

We let $p(j,k)$ denote the ex ante probability that $S_1 = j$ and $S_2 = k$. The corresponding marginal distributions are $p_1(j) = \sum_{k=0}^{n} p(j,k)$ and $p_2(k) = \sum_{j=0}^{m} p(j,k)$. Finally, we denote by $E_k[v(j,k)]$ the expected value of a type j player 1,

$$E_k[v(j,k)] = \frac{\sum_{k=0}^{n} p(j,k)v(j,k)}{\sum_{k=0}^{n} p(j,k)},$$

and similarly for $E_j[v(j,k)]$.

Players maximize their payoff, which is the discounted sum of their profits in each auction, using a common discount factor $\delta \in [0,1)$. Therefore, their utility does not exhibit "diminishing marginal returns" for winning more units. We emphasize that players do not learn the value of a unit upon buying it. Learning, therefore, is restricted to inferring one's rival's information. Bids are observable. A bid is a *serious* bid if it is equal to the highest bid with strictly positive probability, and is a *losing* bid otherwise.

Players maximize their payoff, which is the discounted sum of their profits in each auction, using a common discount factor $\delta \in [0,1)$. Therefore, their utility does not exhibit "diminishing marginal returns" for winning more units. We emphasize that players do not learn the value of a unit upon buying it. Learning, therefore, is restricted to inferring one's rival's information. Bids are observable. A bid is a *serious* bid if it is equal to the highest bid with strictly positive probability, and is a *losing* bid otherwise.

The set of histories of length $t < \infty$, H_t, is defined to be the set of $t \geq 1$ pairs of elements of $\mathbb{R}_+$ (with H_0 being a singleton set containing the "empty history"). A strategy for player 1 (resp. player 2) is a measurable mapping $\sigma_1 : M \times \cup_t H_t \to \Sigma$ (resp. $\sigma_2 : N \times \cup_t H_t \to \Sigma$), where Σ is the set of distribution functions on $\mathbb{R}_+$. We say that player i's type k submits a bid b with positive probability if b is in the support of his distribution function. An infinite history is a countably infinite sequence of pairs of elements of $\mathbb{R}_+$; the set of these is denoted H. Given an

[5] This is, for instance, the environment of Maskin and Riley (2000). A more primitive assumption that leads to the same outcome is that S_1 and S_2 are independently pairwise affiliated with V, in the sense of Milgrom and Weber (1982). Note that we do not require independence for any of the results concerning one-sided incomplete information, even when it arises as a degenerate subcase of two-sided incompleteness.

infinite history $h \in H$, let $b_t^i(h)$ be player i's bid in period t. Player 1 type j's reward in period t along history h, given that player 2 is of type k, is then

$$r_{j,t}^1(h) = \left(v(j,k) - b_t^1(h)\right)\left(I_{W_t^1} + \frac{1}{2}I_{T_t}\right),$$

where I_A denotes the indicator function of the set A, and W_t^1 denotes the event that player 1 wins in period t, i.e.

$$W_t^1 := \left\{b_t^1(h) > b_t^2(h)\right\}.$$

Similarly, $T_t := \{b_t^1(h) = b_t^2(h)\}$ is the event that they tie. Player 2's reward is defined analogously. The prior distribution on types and a strategy profile $\sigma = (\sigma_1, \sigma_2)$ induce a probability distribution over infinite histories in the usual manner, and if we denote by $E_\sigma(\cdot)$ the associated expectation over types and histories, then player i type j's payoff is defined as $(1-\delta)E_\sigma\left[\sum_{t=0}^{\infty}\delta^t r_{j,t}^i(h)\,|j\right]$, his average discounted sum of rewards, conditional on his own type and given some strategy profile σ. Abusing standard terminology, we refer to continuation games as “subgames”.

The solution concept used is Perfect Bayesian Equilibrium (P.B.E.) with the standard restrictions as defined by requirements B(i)-(iv) in Fudenberg and Tirole (1991, pp. 331-2). As explained below, this solution concept puts essentially no restriction on the set of equilibrium outcomes, when the discount factor tends to one. We further wish to prune any equilibrium which depends on payoff irrelevant information, and therefore assume that continuation play only depends on (all players') current beliefs. That is, if there exist two histories after which every type of every player entertains the same beliefs about his opponent, then the actions that follow these histories are the same as well. We refer to this assumption as *Markov stationarity.*

While we describe the results of this game in the case of two-sided uncertainty in Section 5, the paper focuses for the most part on the case of one-sided uncertainty (Section 4). In this case, it is convenient to re-interpret the model and the solution concept as follows: Player 1 is a long-lived bidder who faces a sequence of short-run bidders, called Player 2. Player 1, the *informed* player, receives one of $m+1$ possible signals, while Player 2, the *uninformed* player, receive none. Indeed, Markov stationarity implies that Player 2 maximizes his (instantaneous) reward. Fix any history, and consider the decision of the uninformed player. His continuation payoff only depends on his posterior, but his posterior only depends on the informed player's action (and his prior). Therefore, Player 2's optimal action must maximize his reward. It is worth pointing out that the converse is not true: there are equilibria in which Player 2's action after every history, yet the equilibrium fails to be Markov stationary. Since strong predictions can be obtained under the weaker solution concept, we shall assume myopia:

In the case of one-sided uncertainty, Player 2's action maximizes his reward after every history.

Henceforth, we shall omit to explictly repeat this refinement whenever it is used - it is maintained throughout Section 4 up to but excluding Subsection 4.1. Another refinement, *weak* *dominance*, is then introduced and results are discussed with and without it. Section 5 returns to the case of two-sided uncertainty and discusses the results in that case.

3 Static Benchmark

Our main result implies that, in the infinitely-repeated auction, the average expected revenue of the auctioneer, as the discount factor tends to one, tends to the relatively low value λ (defined above). It is therefore important to derive the corresponding revenue for the static common-value auction with independent types, which is equivalent to the infinitely-repeated auction with $\delta = 0$. Existence of an equilibrium in this finite-type environment (using the Vickrey tie-breaking rule) is established by Maskin and Riley (2000, Proposition 2). In this section, we are able to derive explicit formulas for the equilibrium strategies and the expected revenue of the auctioneer, allowing a direct comparison with the sequential auction (large δ) to be made for any particular prior distributions. This section can be safely skipped by the reader solely interested in the repeated game.

Without loss of generality, we normalize $v(0,0)$ to 0. Because there are only two players, the supports of the bidding distributions of two different types of the same player intersect in one point if these types are consecutive. For each player and every pair of consecutive types of this player, let $\alpha(\kappa)$ be this bid, where arguments are picked in any way such that $\alpha(1) \leq \cdots \leq \alpha(m+n)$. Let S_i^1 denote the support of the bid distribution of player 1's type i, and analogously S_j^2 for player 2's type j. If κ is the index that corresponds to the intersection of the supports of player 1's (resp. player 2's) type i and $i+1$ (resp. type j and $j+1$), let $m(\kappa) = i+1$ and $n(\kappa) = \max\left[j \in N \mid \alpha(\kappa) \in S_j^2\right]$ (resp. $n(\kappa) = j+1$ and $m(\kappa) = \max\left[i \in M \mid \alpha(\kappa) \in S_i^1\right]$). Let $\alpha(0) = 0$, $m(0) = n(0) = 0$, and denote the highest bid in either player's support by $\alpha(m+n+1)$. Let $p(i)$ (resp. $q(j)$) be the probability that player 1 (resp. player 2) is of type i (resp. of type j), and let $F_{m(\kappa)}$ (resp. $G_{n(\kappa)}$) be the bid distribution of player 1's type $m(\kappa)$ (resp. player 2's type $n(\kappa)$) on $[\alpha(\kappa), \alpha(\kappa+1)]$. Let $P(i) = \sum_{l=0}^{i} p(l)$, $Q(j) = \sum_{l=0}^{j} q(l)$, and define recursively $s(\cdot)$ by $s(0) = 0$ and, for all $1 \leq \kappa \leq m+n+1$, for $\kappa - 1$ the largest integer for which $s(\kappa-1)$ has been defined so far, if $\min\left[x \mid x = P(i) > s(\kappa-1) \text{ for some } i\right] \neq \min\left[x \mid x = Q(j) > s(\kappa-1) \text{ for some } j\right]$, then $s(\kappa)$ equals the lowest of these minima. Otherwise, $s(\kappa) = s(\kappa+1)$ equals their common value. Let $v(\kappa) := v(m(\kappa), n(\kappa))$. We show in the appendix that $\alpha(\cdot)$ satisfies the recursion $\alpha(0) = 0$, and, for all $1 \leq \kappa \leq m+n+1$:

$$\alpha(\kappa) = \sum_{l=0}^{\kappa-1} \left(s(l+1) - s(l)\right) v(l).$$

We prove (in the appendix) that:

Theorem 1 *If $\delta = 0$, the expected revenue R for the seller is given by*

$$R = \sum_{l=0}^{m+n} (1 - s(l))^2 (v(l+1) - v(l)),$$

and the distribution functions used by the players are given by, for $b \in [\alpha(\kappa), \alpha(\kappa+1)]$,

$$\begin{aligned} p(m(\kappa)) F_{m(\kappa)}(b) &= s(\kappa+1) \frac{v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)}{v(m(\kappa), n(\kappa)) - b} - \sum_{j=0}^{m(\kappa)} p(i) \text{ and} \\ q(n(\kappa)) G_{n(\kappa)}(b) &= s(\kappa+1) \frac{v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)}{v(m(\kappa), n(\kappa)) - b} - \sum_{j=0}^{n(\kappa)} q(j). \end{aligned}$$

4 One-sided Incomplete Information

This section derives properties common to all equilibria in the case in which the type space is a singleton for exactly one of the players. So suppose that $M = \{0, 1, \ldots, m\}$, with $m \geq 1$, while $n = 0$. Accordingly, we refer to player 1 as the informed player, or player I, and to player 2 as the uninformed player, or player U, and we write $v(j)$ or v_j instead of $v(j, 0)$, where $(j, 0) \in M \times N$, as well as p_t^j for the probability assigned to type j by the uninformed player, in period t (given some history h_t clear from the context). Without loss of generality, we assume that all types in M are assigned strictly positive probability. Also, let $v_t^U = \sum p_t^k v_k$ denote the expected value of a unit. In addition, let $h_t = (h_t^I, h_t^U)$, where h_t^I is the history of bids by the informed player along h_t, and h_t^U is the history of bids by the uninformed bidder. Given a history h_t, let $\underline{k}_t$ denote the lowest type of the informed player assigned positive probability. The lowest element in the support of (player I's) type k bid distribution is denoted α_t^k; the lowest element in the support of player I's bid distribution among all types assigned positive probability, α_t^I; the lowest element in the support of player U's bid distribution, α_t^U.

Similarly, the highest element in the support of (player I's) type k bid distribution is denoted β_t^k; the highest element in the support of player I's bid distribution among all types assigned positive probability, β_t^I; the highest element in the support of player U's bid distribution, β_t^U. Finally, the continuation payoff of Player I's type k after history h_t is denoted $\pi_{h(t)}^k$ or π_t^k for short, and the continuation payoff of Player U is denoted $\pi_{h(t)}^k$ or π_t^U for short.

As discussed in Section 2, Player 2 is assumed to maximize his (instantaneous) reward. As will be clear from the first lemma, this guarantees that the uninformed player always bids at least as much as $v_{\underline{k}_t}$, the lowest value conceivable given h_t. Therefore, bidding strictly more than his signal is a dominated action for Player I's type $\underline{k}_t$ in the supergame. Such dominated actions are ruled out whenever Refinement **R** is invoked.

Refinement R (weak dominance): For all histories h_t, $\beta_t^{\underline{k}_t} \leq v_{\underline{k}_t}$.

This refinement is standard in the literature: it is used, for instance, by Maskin and Riley (2003), and in a closely related context, by Engelbrecht-Wiggans and Weber (1983), who note that it is implied by trembling-hand perfectness. We characterize below the equilibria with or without this refinement.

By assumption, the uninformed bidder plays a myopic best-response to his beliefs. This implies that the uninformed bidder will "break" any tie occurring with positive probability in his favor, provided only that doing so yields a positive payoff. We avoid introducing additional notation for infinitesimal bids by adopting the convention that, whenever incomplete information is one-sided, ties are won by the uninformed player.[6]

Given an equilibrium, let H_t^* be the subset of $t-$histories H_t (including the null history) that have positive probability under the equilibrium strategies. For $i > 0$, let

$$T_i := \inf \left\{ t \in \mathbb{N}_0 ; p_{t+1}^i \in \{0, 1\}, \forall h_{t+1} \in H_{t+1}^* \right\},$$

that is, T_i is the length of the longest finite history having positive probability under the equilibrium strategies in which uncertainty about type i persists until period t (included). The main result of this section is:

Theorem 2 *$\exists \delta' < 1$, $\forall \delta \in (\delta', 1)$:*

(i) $T := \max_{j \geq 1} T_j < \infty$ if $\delta > \bar{\delta}$ for some $\bar{\delta} < 1$; $\lim_{\delta \to 1} T = \infty$; $\lim_{\delta \to 1} \pi_0^k = 0$, all $k \in M$;

*(ii) Under **R**, $\lim_{\delta \to 1} \delta^T = 1$; $\lim_{\delta \to 1} \pi_0^U = 0$; $\lim_{\delta \to 1} \pi_0^U / (1 - \delta) = \infty$; $\lim_{\delta \to 1} \pi_0^k / \pi_0^U = 0$, all $k \in M$.*

That is, uncertainty about types can only persist for finitely many periods, although this number of periods increases without bound as the informed player become more patient (under R1, the patience of the uninformed player plays no role). At the same time, the informed player's payoff tends to zero. Although $\lim_{\delta \to 1} T = \infty$, $\delta^T \to 1$, under **R**. That is, although patient higher types may mimic lower types for a long time, this time is still short relative to their discount factor. In addition, $\lim_{\delta \to 1} \pi_0^U = 0$ and $\lim_{\delta \to 1} \pi_0^k / \pi_0^U = 0$: although each player's average payoff tends to zero when the discount factor tends to 1 (which is unsurprising given the previous conclusion), the uninformed bidder fares better than the informed bidder. Indeed, the proof shows that, in equilibrium, the informed bidder may win at most once before at least some information is revealed (and bids jump upward), while the uninformed bidder is likely to

[6] We follow Engelbrecht-Wiggans and Weber (1983) in this convention.

win many units cheaply. Along any equilibrium path, once the informed bidder fully reveals his private information, both bidders submit bids equal to the object's value.

The proof is divided in several steps.

Lemma 1 *$\forall h_t \in H_t^*$, $\alpha_t^U \geq \alpha_t^I$ and $\alpha_t^U \geq \beta_t^{k_t}$, so that $\pi_t^{k_t} = 0$.*

Proof: The proof is by induction on the number of types K remaining in the support of the uninformed player's beliefs.
Consider any history $h_t \in H_t^*$ such that $K = 1$, so that the informed player is believed to be some type k with probability one. We first argue that $\pi_t^k = 0$. Consider any period $t' \geq t$. Either $\alpha_{t'}^I \geq v_{t'}^U$ and k's reward in period t' is nonpositive, or $\alpha_{t'}^I < v_{t'}^U$. In the latter case, the uninformed can secure a strictly positive reward by bidding $(\alpha_{t'}^I + v_{t'}^U)/2$, for instance. This implies that the uninformed does not submit a losing bid in that period, so that $\alpha_{t'}^U \geq \alpha_{t'}^I$, implying that k's reward must be 0 in this case as well. Therefore, $\pi_t^k = 0$, which implies in particular that $\alpha_t^U \geq v_t^U$ (since type k could profitably bid $(v_t^U + \alpha_t^U)/2$ otherwise).
We now claim that $\alpha_t^U \geq \beta_t^k$. If $\beta_t^k \leq v_t^U$, then the result follows from $\alpha_t^U \geq v_t^U$. If $\beta_t^k \geq v_t^U$, then β_t^k must be a losing bid, for k's reward (and thus his continuation payoff) is strictly negative otherwise. The claim follows, as well as $\alpha_t^U \geq \alpha_t^I$, since $\beta_t^k \geq \alpha_t^k = \alpha_t^I$.
Assume that the result holds whenever the number of types in the support of the uninformed player's beliefs is $K - 1$ or less, and consider any history $h_t \in H_t^*$ such that the support of the uninformed player's belief has exactly K elements. Let $k = k_t$. We first argue that $\pi_t^{k_t} = 0$. Consider any period $t' \geq t$. Either $\alpha_{t'}^I \geq v_k$ and k's reward in period t' is nonpositive, or $\alpha_{t'}^I < v_k$. In the latter case, the uninformed can secure a strictly positive reward by bidding $(\alpha_{t'}^I + v_k)/2$, for instance. This implies that the uninformed does not submit a losing bid in that period, so that $\alpha_{t'}^U \geq \alpha_{t'}^I$, implying that $k_{t'}$'s reward must be 0 in this case as well. If $k_{t'} = k$, his reward is therefore 0, while if $k_{t'} > k$ then by the induction hypothesis, $\pi_{t'}^{k_{t'}} = 0$, so that $\pi_{t'}^k = 0$, and therefore, type k's reward is 0 in this case as well. Therefore, $\pi_t^k = 0$, which implies in particular that $\alpha_t^U \geq v_k$ (since type k could profitably bid $(v_k + \alpha_t^U)/2$ otherwise). The remaining conclusions follow as in the case $K = 1$. ■

As mentioned above, this lemma implies that $\alpha_t^U \geq v_{\underline{k}_t}$ for any history $h_t \in H_t^*$. Therefore, if beliefs are degenerate on some type k, the uninformed player's payoff must be nonpositive, implying that the informed player must bid at least v_k. However, since the lemma also establishes that the uninformed player must win all units if his beliefs are degenerate, he cannot bid more than v_k. Thus, both players submit the bid v_k with probability one in this case.

As a second remark, this lemma also implies that any serious bid rules out type $\underline{k}_t$, so that, in the continuation game, type $k \geq \underline{k}_t$ does not submit more than $k - \underline{k}_t$ additional serious bids. This implies that his payoff is bounded above:

$$\pi_t^k \leq (1 - \delta) \sum_{i=\underline{k}_t}^{k-1} (v_k - v_i) .$$

This implies that $\lim_{\delta\to 1}\pi_0^k = 0$, all $k \in M$, as asserted in Theorem 2.

Lemma 2 $\forall h_t \in H_t^*$, $\beta_t := \beta_t^U = \beta_t^I \leq v_t^U$.

Proof: Suppose first that $\beta_t^U > \beta_t^I$, contrary to the claim. Then the uninformed player would strictly gain by bidding $\left(\beta_t^U + \beta_t^I\right)/2$ rather than β_t^U (using R1). If $\beta_t^U < \beta_t^I$ instead, then consider the lowest type k whose bid support includes β_t^I. His continuation payoff is 0, by Lemma 1. Therefore, he would gain by bidding $\left(\beta_t^U + \beta_t^I\right)/2$ rather than β_t^I.

Therefore, by bidding β_t, the uninformed player wins with probability one (recall that he wins ties). His reward must be positive, so that $\beta_t \leq v_t^U$. ■

Observe that Lemma 2 implies that the highest type $\bar{k}_t$ in the support of the uninformed player's belief can secure a reward of $(1-\delta)(v_{\bar{k}_t} - v_t^U)$ which is strictly positive as long as beliefs are nondegenerate on $\bar{k}_t$. This gives an upper bound on the number of consecutive periods (say, $\hat{T}^{\bar{k}_t}$) in which type $\bar{k}_t$ is willing to submit losing bids. Given Lemma 1, this implies that the support of the uninformed player's beliefs must shrink in finitely periods. In particular, since the lowest type $\underline{k}_0$ only submits losing bids, there exists some sequence of losing bids such that, after $\hat{T}^{\bar{k}_0} + 1$ periods, type $\bar{k}_0$ is ruled out by the uninformed, and the probability assigned to type $\underline{k}_0$ has not decreased. The argument can therefore be repeated from this point on, yielding an upper bound on the number of periods in which the highest remaining type in the support of the uninformed player's beliefs is willing to submit losing bids. By induction, there exists a horizon $\hat{T}$ and a sequence (of length $\hat{T}$) of losing bids such that, conditional on this sequence, the uninformed player assigns probability one to type $\underline{k}_0$. Therefore, each type but the lowest one can secure a strictly positive payoff: $\pi_0^k \geq \delta^{\hat{T}}\left(v_k - v_{\underline{k}_0}\right)$.

This argument holds more generally: given any history $h_t \in H_t^*$ such that beliefs are not degenerate on $\bar{k}_t$, there exists an integer $\hat{T}_t$ and a sequence of losing bids of length $\hat{T}_t$, such that, given any history following h_t in which the informed player's bids are given by this sequence, the uninformed player assigns probability one to type $\underline{k}_t$. Therefore, $\pi_t^k > 0$ for all $k > \underline{k}_t$. There is an important caveat to this discussion. While $\hat{T}_t$ is finite, it depends on the probability assigned to the lowest type that is still in the beliefs' support after history $h_t \in H_t^*$ (more precisely, on the spread $v_{\bar{k}_t} - v_t^U$). Therefore, it does not follow from this that T (the upper bound over *all* histories until all uncertainty is resolved) is finite, because there may well be histories after which the belief assigned to such types is arbitrarily small.

Nevertheless, an important consequence of our discussion is:

Lemma 3 $\forall h_t \in H_t^*$, $\alpha_t^U = \beta_t^{\underline{k}_t}$.

Proof: Suppose not, i.e., given Lemma 1, suppose that $\alpha_t^U > \beta_t^{\underline{k}_t}$. This implies that at least one type of the informed bidder must submit with positive probability a bid in the interval $\left(\beta_t^{\underline{k}_t}, \alpha_t^U\right]$. Pick such a bid b and consider the lowest type $\hat{k}$ submitting this bid. Since $b > \beta_t^{\underline{k}_t}$, we have $\hat{k} > \underline{k}_t$. Since b is a losing bid, and $\hat{k}$ is the lowest type assigned positive probability in the continuation game, it follows that type $\hat{k}$'s payoff from submitting such a bid is zero, a contradiction given the argument in the text. ■

As we know from Lemma 1 that the informed player wins only finitely many times, the undiscounted game between him and a sequence of short-run players is well-defined (results so far have not used strict discounting). However, the same line of reasoning shows that this game admits no equilibrium, as long as the prior is nondegenerate. To see this, consider any history on the equilibrium path specifying an infinite sequence of losing bids by the informed player, along which the probability assigned to the lowest type is non-decreasing. This is possible, since we have seen that this type only submits losing bids. We claim that there exists $\hat{T}$ such that, along that sequence, this probability is one for all $t > \hat{T}$. Indeed, the highest type -who can secure a strictly positive payoff by submitting β_0- must submit a serious bid with probability one in some period $\hat{T}_1$, for otherwise he would be willing to lose forever, resulting in a payoff of 0. From time $\hat{T}_1 + 1$ on, the same argument applies to the second-highest type still assigned positive probability, etc.

This implies that the payoff of the second lowest type, type 1, must be at least $v_1 - v_0$, since he can submit losing bids until period $\hat{T}$ and win at least once by bidding slightly more than v_0 in the following period. At the same time, in equilibrium, he must submit a serious bid with positive probability before period $\hat{T}$. Yet in any such period t, the uninformed player must bid strictly more than α_U^t with positive probability, implying that the expected payoff from the serious bid is strictly less than $v_1 - v_0$ (the continuation payoff being necessarily zero).[7]

The next lemma proves that the infinitely repeated game actually ends (in the sense that all information is revealed and actions become static), endogenously, in [boundedly] finite time. Roughly, this is because the informed player (unless he is the lowest type) must enjoy his information rents at some point, but to do so he must make a serious bid and thus give up some of that same information. Due to the discount factor, he is only willing to wait for some well-defined period before all of this occurs. Slightly paradoxically, the specific argument we employ requires δ large enough in order to yield that information is revealed in finite time, but of course for small δ we expect it to occur even more rapidly, since the incentives are even stronger for the informed bidder not to wait; indeed, as we saw above, for $\delta = 0$ the game is essentially over immediately.

Lemma 4 *There exists $\bar{\delta} < 1$, $\forall \delta > \bar{\delta}$, $T < \infty$.*

Proof: As the priors are nondegenerate, $v_0^U < \bar{v} := v_{\bar{k}}$, where $\bar{k} := \bar{k}_0$ is the highest type. The uninformed player is never willing to bid more than his expectation, so $\beta_0 < \bar{v}$ and hence $\pi^{\bar{k}} > 0$. This (along with $\delta < 1$) implies that at least type $\bar{k}$ makes a serious (i.e. potentially winning) bid by some time T_1 with probability one. Meanwhile, type 0 never makes a serious

[7] The same argument applies if payoffs are evaluated according to the overtaking criterion. The limit of means criterion is consistent with our characterization, but also with other equilibrium outcomes.

bid. Thus, the support of player U's beliefs is necessarily strictly smaller at time $T_1 + 1$ than at time 0, no matter what sequence of bids has been observed.

So consider the beginning of period $T_1 + 1$: suppose first that only losing (more precisely, nonserious) bids have been observed so far from the informed player. By the reasoning above, there is a new highest type $\bar{k}' := \bar{k}_{T_1+1} < \bar{k}$. Since type 0 only makes losing bids, there is at least some subsequence of such bids (by the informed) so that the new expected value $v^U_{T_1+1}$ is bounded away from $v_{\bar{k}'}$ by the original prior on type 0.[8] Since type $\bar{k}'$ could have made precisely those bids if he so chose, this places a lower bound on $\pi^{\bar{k}'}_{T_1+1} > 0$ and hence a uniform upper bound on the number of subsequent periods in which he is willing to continue to submit losing bids, just as above. Call this T_2.

If, on the other hand, we consider a history along which the informed submits a serious bid by period T_1 (WLOG we shall assume it is in period T_1 itself), we cannot uniformly bound the posteriors at that point. In particular it is possible that after certain separating bids we have $v^U_{T_1+1}$ arbitrarily close to $v_{\bar{k}}$ (where $\bar{k}$ might really be some $\bar{k}_{T_1+1} < \bar{k}_0$, but it makes no difference to the argument). In this case the reasoning above does not go through directly, but note that we can still place a uniform lower bound on the highest types' payoff as follows.

Instead of separating in period T_1, $\bar{k}$ could have pooled (submitted a losing bid) and then placed the highest equilibrium bid in each of the subsequent two periods. Now [conditional on observing losing bids] the highest bid must be decreasing over time.[9] Hence $\beta_{T_1+1} \leqq \beta_{T_1}$ and (after seeing a bid of β_{T_1+1}) $\beta_{T_1+2} \leqq v^U_{T_1+2} \leqq v_{\bar{k}-1} < v_{\bar{k}}$, where the last inequality follows because type $\bar{k}$ was expected to separate in period T_1. Therefore $\pi^{\bar{k}}_{T_1} > \delta(v_{\bar{k}} - \beta_{T_1}) + \delta^2(v_{\bar{k}} - v_{\bar{k}-1})$, whereas by submitting the highest bid in period T_1 (which the highest type, by single crossing, must be willing to do) type $\bar{k}$ achieves $\pi^{\bar{k}}_{T_1} = (v_{\bar{k}} - \beta_{T_1}) + \delta C$, where C is his continuation value. Thus for δ close enough to 1, we get a [strictly positive] lower bound on C that is uniform since it depends only on the initial parameters of the model. As before, this implies an upper bound T'_2 on the number of additional periods in which the highest type is willing to pool.

It is now clear that we can continue this process: after $T_1 + \max\{T_2, T'_2\}$ periods the support of beliefs has strictly decreased twice, and so on. In each case we follow one or the other of the lines of reasoning above, depending on whether or not separation has occurred within the 'subgame'. Finally, then, we have shown that beliefs are necessarily degenerate by some fixed time T. ■

The final lemma is the key of the result. The central step is to show that, in a precise sense, either the spread of serious bids $\beta_t - \alpha^U_t$ is small, or information is revealed quickly. Indeed, as the uninformed player must be indifferent between bidding β_t and α_t, it must be that the value

[8] In fact this must be true 'on average', in a well-defined sense, but we shall not need this stronger statement.

[9] Otherwise whichever type k was the lowest type willing to submit the highest bid in the following period (and was thus expecting a zero continuation value) would strictly prefer to submit the highest bid in the current period instead.

of the unit conditional on winning with β_t is sufficiently large, relative to the value conditional on winning with α_t^U, to compensate for the difference in bids. In fact, a partial converse to this holds as well. As a consequence, it must be that $\lim_{\delta\to 1} T = \infty$, for otherwise, given some history h_t, the second lowest type of the informed player could secure a payoff arbitrarily close to $(1-\delta)(v_k - v_{\underline{k}_t})$ simply by waiting. Yet the value of T gives a lower bound (independent of δ) on the probability with which this type must submit a serious bid in one of the periods before, and in turn, this gives a lower bound on the spread of serious bids in this period, and thus an upper bound, strictly below $(1-\delta)(v_k - v_{\underline{k}_t})$, on the payoff that this type can get by submitting such a bid. For δ close enough to one, this gives the desired contradiction.

Once we have deduced this relationship between the speed of information revelation and the bid spread, we finally (for the first and only time) invoke refinement **R**, which pins down the lower end α of serious bids. This allows us to take the final steps in the proof, and in particular to show that T is 'soon' relative to δ, which in turn determines all long-run payoffs.

Lemma 5 *Under **R**,* $\lim_{\delta\to 1} \delta^T = 1$; $\lim_{\delta\to 1} \pi_0^U = 0$; $\lim_{\delta\to 1} \pi_0^U/(1-\delta) = \infty$; $\lim_{\delta\to 1} \pi_0^k/\pi_0^U = 0$, *all* $k \in M$.

Proof: Recall that α_t^U is the lowest serious bid of the uninformed player, and β_t his highest bid (given Lemma 2). As always p_t^k is U's belief in period t (after seeing only losing bids so far) about I being of type k; we introduce $\hat{p}_t^k$ to denote U's probability of facing a type k given that he won with a bid of α_t. Thus $\hat{p}_t^k$ is an appropriately weighted average of the p_{t+1}^k values that follow the different nonserious bids in period t. Analogously, we use γ_t^k for the [unconditional] probability that the informed player is of type k and that he submits a serious bid (i.e. strictly above α_t^U) in period t, and thus separates, and we use γ_t for the sum of these across k. The following claim formalizes the intuition that if the uninformed is willing to make a spread of bids, then he must be compensated for doing so via a correspondingly greater expected value for the object.

Claim 1: $\beta_t - \alpha_t^U > \varepsilon$ implies $\gamma_t > \varepsilon/\bar{v}$.

Proof First note that from Bayes Rule we have $\hat{p}_t^k = (p_t^k - \gamma_t^k)/(1-\gamma_t)$. Meanwhile, U must be indifferent between bidding β_t and α_t^U, so $\sum_k p_t^k v_k - \beta_t = (1-\gamma_t)\left(\sum_k \hat{p}_t^k v_k - \alpha_t^U\right) = \sum_k (p_t^k - \gamma_t^k) v_k - (1-\gamma_t)\alpha_t^U$ (using the above), and hence $\sum_k \gamma_t^k v_k = \beta_t - (1-\gamma_t)\alpha_t^U$. Thus $\gamma_t \bar{v} = \sum_k \gamma_t^k \bar{v} > \sum_k \gamma_t^k v_k = \beta_t - (1-\gamma_t)\alpha_t^U \geqq \beta_t - \alpha_t$ and the result follows. This is of course not necessarily a tight bound.▼

Applying this single-period reasoning to the game as a whole, we see that T cannot remain bounded as δ increases. Otherwise, two mutually contradictory events would necessarily ensue: the higher types of the informed player would be separating with noticeable (i.e. bounded away from zero) probability in at least some periods (since T is bounded); and their payoff in any period in which they're willing to separate would be arbitrarily close to what they could obtain from being thought to be the lowest type for sure (which occurs in period $T+1$, but δ is going to

1). This latter, combined with what we just observed, implies that there must be arbitrarily little separation in any given period, giving us the contradiction. This is formalized in the following result.

Claim 2: $\lim_{\delta\to 1} T = \infty$ and $\lim_{\delta\to 1} \pi_0^k = 0$, all $k \in M$.

Proof We begin with an immediate partial converse to the preceding claim. Namely, if the spread $\beta_t - \alpha_t^U$ is small, then the uninformed would strictly prefer to bid β_t, paying only slightly more but winning against an additional $\gamma_t^{\bar{k}}$ of mass with value $\bar{v} > v_t^U \geqq \beta_t$, unless $\gamma_t^{\bar{k}}$ is concomitantly small, or $p^{\bar{k}}$ is very large. But now suppose that $\lim_{\delta\to 1} T < \infty$: since type 1 (or more generally type $\underline{k}_t + 1$) can always pool until period $T + 1$, and at that point secure a payoff of $v_1 - v_0$, his unnormalized (but discounted) payoff must tend toward exactly $v_1 - v_0$ as $\delta \to 1$. [We know that $\alpha_t^U \geq v_0$ so he cannot hope to achieve anything greater.] Furthermore, since he wins at most once, he must expect to win almost for sure at a price of essentially v_0, and of course very quickly relative to δ. Similarly, type 2 can always mimic type 1, gaining $v_2 - v_0$, and then (out of equilibrium) bid just above v_1, gaining an additional $v_2 - v_1$. Continuing, we find that the highest type $\bar{k}$ must also win at the cheapest possible price (i.e. v_{k_t}) every time he places a serious bid in equilibrium. But since (by single crossing) he must be willing to bid β_t in every period $t < T$, this implies that $\beta_t \to v_{\underline{k}_t}$ as $\delta \to 1$, and hence that (along an equilibrium trajectory followed by type $\bar{k}$) $\beta_t - \alpha_t^U \to 0$ and so $\gamma_t \to 0$. This is a contradiction, as by definition he separates entirely by time T, which was supposedly bounded independently of δ. Therefore $\lim_{\delta\to 1} T = \infty$, from which it is obvious that the [normalized] payoff of the informed is vanishing.▼

The next link in the chain of reasoning is to put an upper limit on the amount of time (relative to δ) before 'almost all' of the mass of higher types has dissipated, implying that the spread in bids is small (using the first lemma above) and therefore that it is cheap to win (using **R**). This will allow us to conclude that (as players become more and more patient) $\delta^T \to 1$ despite $T \to \infty$. We start with the following claim concerning dissipation of mass.

Claim 3: $\forall h_t \in H_t^*$ there is a uniform upper bound (depending only on $p_t^{\underline{k}_t}$) on the number S of consecutive periods t' in which $\gamma_{t'}$ can exceed any given $\varepsilon > 0$.

Proof On first glance, since γ_t is an unconditional probability of separation, it might seem that $\sum_{t'=t}^{t+S} \gamma_t \leqq 1 - p_t^{\underline{k}_t}$ for any S, using the fact that the lowest type never separates. But upon reflection, although this will lead us in the right direction, it is not true: if there is a miniscule probability of the lowest type, then it is perfectly possible for the unconditional chance of separation to be extremely high, and yet (conditional on pooling) the new unconditional chance of separation in the subsequent period to be equally high. This cannot continue ad infinitum, however. In particular, suppose that $\gamma_{t'} \geqq \varepsilon$ in periods $t' = t, t+1, ..., t+S$. Then we must have $(1-\varepsilon)^S(1 - p_t^{\underline{k}_t}) \geqq \varepsilon$, i.e. there must be sufficient mass of the higher types remaining for the unconditional probability of separation to feasibly be ε. Thus $S \leqq \ln\left(\varepsilon/(1 - p_t^{\underline{k}_t})\right) / \ln(1-\varepsilon)$ serves as our bound; note that as ε increases the maximum S decreases, as expected. We have ignored the question of whether all higher types separate with equal probabilities, but it is irrelevant since we consider total mass only. Finally, note that although there may be multiple sequences of losing (pooling) bids, the argument holds along any individual one of them and the same S is binding in each case.▼

Unfortunately, $S(\varepsilon)$ does (naturally) depend on $p_t^{\underline{k}_t}$ at any particular time t, and we cannot put a uniform lower bound on this latter number across all equilibrium histories (different losing bids may lead to different posteriors). But we do know (via an application of the pigeonhole principle) that there is at least one sequence of losing bids for which the posterior on the lowest type is no lower than the prior. Furthermore, in equilibrium the choice of losing bid cannot affect the future distribution used by U (as otherwise all types k of the informed, conditional on bidding nonseriously, would make whatever losing bid caused the most favorable future outcomes). Specifically, it can't affect either $\alpha_{t'}$ or $\beta_{t'}$ for $t' > t$, and thus if $\beta_t - \alpha_t$ exceeds ε along some such sequence it necessarily does so along all such sequences. In the end, then, we are able to obtain bounds $S_\mu(\varepsilon)$ for each 'subgame' (defined by a belief support μ) that are uniform in the original priors. There are a finite number of types and thence also of possible supports μ, so we may take an overall $S^*(\varepsilon) = \max_\mu S_\mu(\varepsilon)$.

We may now proceed to prove part (ii) of the theorem: Under **R**, $\lim_{\delta\to1}\delta^T = 1$; $\lim_{\delta\to1}\pi_0^U = 0$; $\lim_{\delta\to1}\pi_0^U/(1-\delta) = \infty$; $\lim_{\delta\to1}\pi_0^k/\pi_0^U = 0$, all $k \in M$. Fix some small $\varepsilon > 0$. From the lemmas above, we know that there is a maximum number of periods $S^*(\varepsilon)$ (bounded uniformly in δ) for which γ_t can exceed $\varepsilon/\bar{v}$ and hence for which $\beta_t - \alpha_t$ can exceed ε. We now finally apply our weak dominance refinement to pin down α_t^U; as $\alpha_t^U = \beta_t^{\underline{k}_t}$, and $\alpha_t^U \geq v_{\underline{k}_t}$ (as argued below Lemma 1), by **R**, since $\beta_t^{\underline{k}_t} \leq v_{\underline{k}_t}$, it must be that in fact $\alpha_t^U = v_{\underline{k}_t}$. So after some history $h_t \in H_t^*$, any type k can be assured of waiting at most $S^*(\varepsilon)$ periods before being able to win with probability one at a price of no more than $v_{\underline{k}_t} + \varepsilon$. Taking $\delta \to 1$ and noting that this holds for any $\varepsilon > 0$, we see that the limit of the unnormalized (but discounted) payoff for type k must be arbitrarily near (and therefore equal to) $\sum_{k'=0}^{k-1}(v_k - v_{k'})$.[10]

By definition, at least one type is willing to pool for T periods. Since we know that each type k wins at most k units, and the preceding paragraph implies that they win at least k units, they must indeed win exactly k times, and each time receive almost the full $v_k - v_{k'}$ (where 'almost' means closer and closer as δ approaches 1). Hence it is impossible that $\lim_{\delta\to1}\delta^T < 1$, as then at least one type for at least one unit would not receive the full $v_k - v_{k'}$. After T, beliefs are degenerate and instantaneous payoffs are zero for both I and U, so $\lim_{\delta\to1}\delta^T = 1$ implies immediately that $\lim_{\delta\to1}\pi_0^U = 0$. However, since $\lim_{\delta\to1}T = \infty$ and $\lim_{\delta\to1}\beta_t = \alpha_t$ pointwise in t (which is clear from the description of I's unnormalized payoffs), the uninformed wins many units (in expectation) at a price boundedly below his expected value. As a consequence, his own unnormalized payoff is unbounded – unlike for the informed player. This completes the proof. ■

[10] Technically, the argument holds directly only for $k = 1$, but just as in Lemma x above, it is easy to extend it from there to all k.

4.1 The role of the refinement(s)

For all the results of the previous section, we have assumed that Player 2 is myopic (that is, that he maximizes his instantaneous reward). In addition, some results have been derived under refinement **R**, weak dominance. How necessary are these refinements? As we will argue, none of the conclusions of Theorem 2(i) is valid in the absence of myopia, and none of the conclusions of 2(ii) is valid in the absence of **R**, even if myopia is strengthened to Markov stationarity (that is, even if attention is restricted to strategy profiles that only depend on the current uninformed player's beliefs).

If myopia is dropped, the folk theorem is valid, even if **R** is imposed: *any strictly positive payoff pair* (π^U, π^I) *such that* $\pi^U + \pi^I < v_0^U$ *is an equilibrium payoff pair provided the discount factor is sufficiently close to one.* In the equilibrium outcome described in Theorem 2, payoffs tend to zero as δ tends to one. Therefore, this equilibrium can be used as a threat to enforce other equilibrium behaviors. The construction follows closely Ausubel and Deneckere (1989) and is omitted. This can be done with or without information revelation.

The second refinement (weak dominance) plays a more subtle role. Roughly speaking, the logic of the proof of Theorem 2(ii) is as follows. As long as uncertainty persists, the upper extremity of the bid supports must decrease at the geometric rate of discounting. By **R**, the lower extremity cannot exceed the signal of the lowest type assigned positive probability. This yields a lower bound on the length of the uninformed bidder's bid support. In turn, this yields a lower bound on the probability with which the informed bidder reveals his private information in a given period, and thus, an upper bound on the time at which all uncertainty must be resolved.

If **R** is dropped, a counterexample satisfying myopia (and even Markov stationarity) can be found for each of the claims of Theorem 2(ii). Without **R**, there is a continuum of equilibria. Each of them is characterized by the value, in each period, of the lower extremity of the (serious) bid support, which can be chosen anywhere between v_{k_t} and v_t^U. In some, information revelation occurs at an arbitrarily small rate (so that $\delta^T \to 0$, rather than $\delta^T \to 1$). Similarly, the unnormalized payoff of the uninformed player, $\pi_0^U/(1-\delta)$, can be chosen anywhere between 0 and infinity, while the normalized payoff, π_0^U, need not tend to 0. In the simple example in which the objects' value is either 1 with probability p, or 0 (see below), it can be shown that there exists a Markov stationary equilibrium in which the uninformed player's payoff equals π, when players are sufficiently patient, if and only if $\pi \in (0, p - (1 - e^{-p}))$ (observe that the upper bound from the folk theorem, p, cannot be attained with myopia). The proof is available from the authors.

4.2 The robustness of the results

What if the horizon is finite? If the horizon is long enough relative to discounting, nothing changes. More precisely, any equilibrium satisfying **R** remains an equilibrium for the finitely repeated game that lasts at least T periods (where T is defined as above). Results are however different when the horizon is short relative to discounting. Take for instance the case of no discounting (the results remain valid for δ close enough to one). Engelbrecht-Wiggans and Weber (1983) examine the case of two types. On the one hand, as they observe, refinement **R** is necessary to avoid a continuum of equilibria. This conclusion remains valid with finitely many types. For instance, it is an equilibrium outcome for both players to bid v_0^U (with probability one) in every period but the last, and behave as in the static auction in the last period, in which the highest bid is precisely v_0^U. The payoff of the uninformed player is then 0, independently of the length of the game, a result that contrasts sharply with their characterization in the two-type case (with **R**). On the other hand, myopia is not necessary in the two-type case, and probably not in the finitely many type case either, as the equilibrium characterization may be obtained by backward induction alone.[11] Proving this appears however difficult, and we will therefore assume both myopia and **R** in what follows.

Consider therefore the undiscounted, N-finitely repeated game, and assume refinement **R**. Lemmas 1 and 2 remain valid. In fact, we have the following generalization of Engelbrecht-Wiggans and Weber (proof available from the authors). The notation follows the infinite horizon case with discounting.

Theorem 3 *Under **R**,*

(i)

$$\lim_{N\to\infty}\frac{\pi_0^U}{N} = v_0^U - v_0,\ \lim_{N\to\infty}\pi_0^k = \sum_{i=0}^{i=k-1}(v_k - v_i)\,.$$

For all history $h_t \in H_t^*$:

(ii) the support of serious bids of type $k > k_t$ is independent of k and, given k_t (i.e. the support of beliefs at time t), is independent of t.

(iii)

$$\alpha_t^k \leq \alpha^U \Leftrightarrow N - t \geq k - k_t.$$

This implies that, in equilibrium, the lowest type submits a losing bid with positive probability (in fact, with probability one) in every period, while the second lowest one does so with positive probability in every period (conditional on having done so in all previous periods) up to (but not including) the last period, the third lowest one does so in every period up to the penultimate period, etc. This is a new feature relative to the results of Engelbrecht-Wiggans and Weber.

There are two main differences with the discounted case. First, while in the discounted case, the discount factor "endogenously" pins down how long uncertainty persists, i.e. the value of T, here uncertainty about the informed player's type may persist until the last period. Second,

[11] That is, as long as equilibrium uniqueness can be established at each stage of the backward induction.

this implies that the expected fraction of periods in which the uninformed player wins the object cheaply tends to one, so that his average payoff tends to $v_0^U - v_0$ rather than 0.

The basic idea behind the proof is that, although uncertainty may be resolved before the end of time[12], it cannot be the case that it necessarily will be. Otherwise, in that final period T of uncertainty, any type that was supposed to separate would strictly prefer to wait and do so more cheaply in the following period, without incurring any delay cost. Meanwhile, it can't be that three (or more) types are still in the belief support of the uninformed in the true final period N. This would imply that the highest of those types is expecting to win (at most) once more, when he would prefer to win in the previous period (more cheaply), not yielding full information about his exact type, and then again in the final period (even if more expensively at that point). Since type k of the informed can't win more than k times (exactly analogously to the discounted case: he gives away information with each serious bid, causing all future bids to jump), this implies that he expects to win precisely k times. Of course, the uninformed wins the rest of the units, and prices are relatively low while uncertainty persists (which is potentially throughout), so this pins down the payoffs.

What if there is more than one uninformed player? As before, a folk theorem obtains if uninformed players are not myopic. If they are myopic, Lemma 1 and 2 remain valid, and in addition each uninformed player's payoff is zero, since competition in each period dissipates rewards. Therefore, all statements in Theorem 2 relative to this payoff must be accordingly modified. In addition, observe that this zero payoff condition gives an additional boundary condition which pins down the lower end of the uninformed players' bid support (as a function of his posterior belief conditional on observing a losing bid). It turns out that this suffices to further characterize the equilibrium outcome. For example, in the case in which the object is either worth 1, with probability p, or worth 0, it is no longer true that $\delta^T \to 1$, but instead $\delta^T \to 1 - p$. The other conclusions in Theorem 2 remain valid. The level of the (payoff-irrelevant) losing bid(s) of the informed player is undetermined, as the competition between uninformed players is sufficient to pin down their lowest serious bid. This means that refinement 2 is neither automatically satisfied nor ruled out by multiple uninformed players.

4.3 A simple example

To better understand the logic behind this result, it is useful to concentrate on a special case. We consider in what follows the case $m = 1$, i.e. in which the informed player can be either a high type or a low type only, and we normalize values to $v(0) = 0$, $v(1) = 1$.

If the informed bidder gets a low signal, he bids 0 in every period (by $\mathbf{R}$). Therefore, as soon as a strictly positive bid is observed, all information is revealed, and bids are 1 from then on. It remains to determine strategies for histories such that all bids by the informed player have

[12] In fact with very high probability it will be.

been zero so far, and all notations that follow are conditioned on such a history. We let p_t be the probability that the value is 1 in period t. By definition $p_0 = p$, which we assume in $[0, 1)$ to avoid trivialities. Let β_t denote the highest bid in the support of either player and U_t is the bid distribution in period t used by the uninformed bidder. Let H_t denote the bid distribution in period t used by the informed bidder with a high signal. We claim that:

1. if $1 - p_t < \delta$, then the informed player with high signal must bid 0 with strictly positive probability. If he does not, then given his equilibrium bid, his continuation payoff is zero (because his signal will be known) and his payoff in period t must therefore be his payoff in the static auction given beliefs p_t, that is $(1 - p_t)(1 - \delta)$. On the other hand, by bidding 0 instead, he will be able to win the unit in period $t + 1$ at negligible cost, as the uninformed bidder will wrongly believe that he faces an informed player with low signal. His payoff from doing so is therefore $\delta(1 - \delta)$, which must be lower than his equilibrium payoff, an immediate contradiction.

2. If $1 - p_t \geq \delta$, then the informed bidder with high signal cannot bid 0 with positive probability: by doing so, he could hope for no more than $\delta(1 - \delta)$, while he gets $(1 - p_t)(1 - \delta)$ by bidding p_t immediately.

3. The equilibrium sequence $\{p_t\}$ is weakly decreasing and gets below $1 - \delta$ in finite time. This follows from observations 1 and 2, Bayes' rule, the obvious fact that an informed bidder with low signal bids zero in equilibrium, and the fact that the informed bidder with high signal is only willing to bid zero with positive probability finitely many times: by waiting t periods before submitting a strictly positive bid, the high-signalled informed bidder cannot achieve a payoff larger than $\delta^t(1 - \delta)$, while he can get $(1 - p)(1 - \delta)$ by bidding p immediately.

Let T denote the first period t in which $1 - p_t \geq \delta$. As long as the informed bidder has bid 0 and until period T, the informed bidder submits either a bid of 0 (with positive probability), or continuously randomizes his bid over some support $[0, \beta_t]$. Meanwhile, the uninformed bidder submits either a bid of 0 (with discrete probability), or continuously randomizes over the same support $[0, \beta_t]$. If the informed bidder has bid 0 in every period up to T, then his strategy assigns zero probability to a bid of 0 in period T (he only continuously randomize over $[0, \beta_T]$), and his private information is thereby revealed. As soon as his private information is revealed, bids remain constant.

We can solve for the equilibrium strategies. Because the informed bidder with high signal is willing to bid 0 before, we must have, for $t < T$,

$$1 - \beta_t = \delta\left(1 - \beta_{t+1}\right), \text{ or } 1 - \beta_t = \delta^{T-t}(1 - \beta_T)$$

In addition, because of Bayes' rule,

$$1 - p_{t+1} = \frac{1 - p_t}{1 - \beta_t}.$$

In period T, as the high-signalled informed player only submits strictly positive bids, it follows that $p_T = \beta_T$. Therefore, the probability $1 - p_T$ is given by:

$$1 - p_T = \frac{1 - p_0}{\left(1 - \beta_{T-1}\right)\left(1 - \beta_{T-2}\right)\cdots\left(1 - \beta_0\right)} = \frac{1 - p_0}{\delta\left(1 - \beta_T\right)\delta^2\left(1 - \beta_T\right)\cdots\delta^T\left(1 - \beta_T\right)},$$

and thus:

$$1 - p_T = \left(1 - p\right)^{1/(T+1)}\delta^{-T/2}.$$

The time T must satisfy $1 > 1 - p_T \geq \delta$. Therefore, T satisfies:

$$\delta^{(T+1)(T+2)/2} \leq 1 - p < \delta^{T(T+1)/2}.$$

As the family of sets $\left\{\left[\delta^{(n+1)(n+2)/2}, \delta^{n(n+1)/2}\right)\right\}_{n\in\mathbb{N}}$ partitions the unit interval $[0,1)$, this establishes the existence and uniqueness of:

$$T = \min\left\{t \in \mathbb{N}; 1 - p \geq \delta^{(T+1)(T+2)/2}\right\}.$$

Observe that, in accordance with the previous theorem, $\lim_{\delta\to1} T = \infty$, but $\lim_{\delta\to1}\delta^T = 1$. Payoffs and bid distributions are derived in the appendix. Figure 2 illustrates the equilibrium strategies.

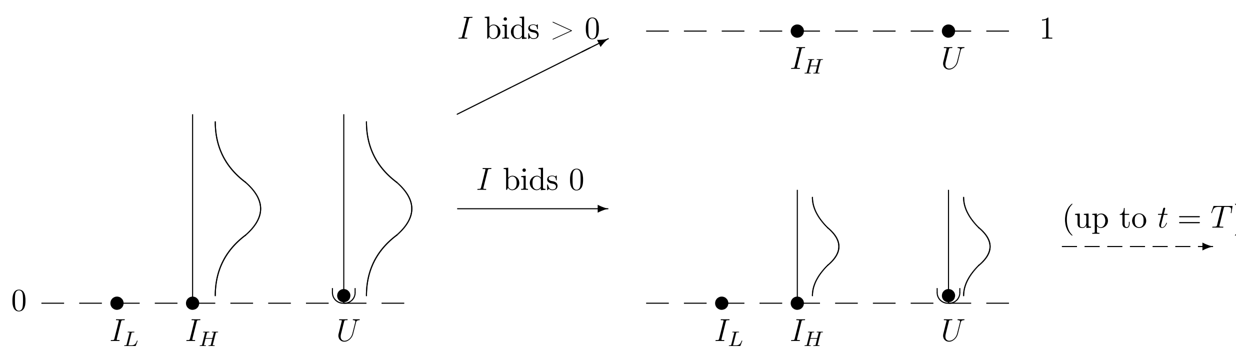

I_L,I_H: low- and high-signalled informed bidder ; U: uninformed bidder

Figure 2: (Random) equilibrium bid trajectory

Since Ashenfelter (1989), several papers have studied or documented the "declining price anomaly". The price trajectory can be determined in this example. The expected maximum bid (conditional, as usual, on the informed bidder having bid 0 up to $t-1$) is

$$E_t = \left(1 - \delta^{T-t}\left(1 - p_T\right)\right)^2,$$

which is decreasing in t. The *unconditional* expectation F_t of the winning bid in period $t \geq 1$, $t \leq T$, is given by:

$$F_t = 1 - \delta^T (1 - p_T) \frac{\delta^{-t} - 1}{\delta^{-1} - 1} \left(1 - \left(1 - \delta^{T-t} (1 - p_T)\right)^2\right),$$

which is decreasing in t as well. Of course, for $t > T$, it is equal to the prior, p, and is larger than the corresponding expectation for all $t \leq T$. Details can be found in the appendix.

5 Two-sided Incomplete Information

Under one-sided incomplete information, patient players achieve a very low payoff. Therefore, the incentive to completely reveal one's private information, thereby becoming *de facto* uninformed, should be low when incomplete information is two-sided. By induction, one may therefore expect equilibrium bids to be uninformative, so that, by Markov stationarity, a player should repeatedly submit the same bid over time. Clearly, it cannot be an equilibrium for the two players to submit repeatedly two bids that differ across players, as one of the players' payoff would be zero. Thus, one could reasonably suspect the equilibria of the game with two-sided incomplete information to be "pooling" equilibria, in which both players submit the same bid b^* independently of their type, where b^* is low enough to be individually rational even for the lowest type. That is, since conditional on winning in such circumstances there is no winner's curse, the bid should not exceed λ, where λ is defined as

$$\lambda := \min \langle E_k[v(0,k)], E_j[v(j,0)] \rangle .$$

Thus λ is the smallest expected value that any type of any player has for the object.[13]

Our main theorem, below, states that the equilibrium outcome is indeed very different in the case of two-sided incomplete information than it is in the case of one-sided incompleteness. Existence is straightforward to prove, while the characterization of what equilibrium outcomes can occur – although more involved – follows the basic reasoning above. Highest types will not want to separate because of the drastic nature of the one-sided outcome that would ensue. If other types either fully or partially separate (where the latter is construed to mean that a subset of types separates, i.e. the posterior belief support of the opponent is strictly smaller), then an inductive argument tells us that some sort of pooling must occur, and if δ is large enough, then all of these possible pooling outcomes must be very similar; otherwise all types of the player would have a strict preference for one over another. All of this holds whether or not a player expects his opponent to separate or pool in the current period, since of course he has no way

[13] Technically, it should be $\lambda = \min \langle E_k[v(j^*,k)], E_j[v(j,k^*)] \rangle$, where $j^* = \min \langle j : p_1(j) > 0 \rangle$ and $k^* = \min \langle k : p_2(k) > 0 \rangle$, but we can normalize both of these to be type 0 without loss of generality.

to affect that at this point. Note that the revenue implications of our result are most easily discussed in the context of the binary case introduced in the previous section, so we hold off on specifics in that regard until the continuation of the example.

We require two types of refinements here: the first is Markovian in the sense that it forces behavior to be measurable with respect to the coarsest summary statistic, and second are two technical assumptions that apply in certain specific circumstances. We also maintain refinement **R** from the one-sided case, applicable only when beliefs are degenerate for one player.[14]

Abusing notation slightly, let $\sigma_i^k(s)$ denote the weight given to pure strategy s by type k of player i in a mixed profile σ. Also, let $EP(s, \sigma_{-i})$ denote the total expected discounted probability of winning when using strategy s against σ_{-i}.

Markov stationarity: Actions chosen depend only on current beliefs. Further, for some $C > 0$, if s and s' are in the support of player i in an equilibrium σ, and $|EP(s, \sigma_{-i}) - EP(s', \sigma_{-i})| < \epsilon$, then $\left|\frac{\sigma_i^k(s)}{\sigma_i^k(s)+\sigma_i^k(s')} - \frac{\sigma_i^{k'}(s)}{\sigma_i^{k'}(s)+\sigma_i^{k'}(s')}\right| < C\epsilon$ for all k, k' such that $\sigma_i^k(s)+\sigma_i^k(s') > 0$ and $\sigma_i^{k'}(s)+\sigma_i^{k'}(s') > 0$.

The second half of the statement above rules out the possibility that beliefs are used to separate actions that are 'payoff-equivalent'. It says that if two strategies yield almost the same result in terms of probabilities (and hence also payments), then they must be used in similar proportions by any types who use either of them at all. Without this assumption, beliefs themselves can be used solely to mark time and are thus not necessarily as coarse as possible. In the discussion below, we illustrate this possibility by means of counterexamples.

Continuity: If along some equilibrium path, the beliefs of both players have well-defined limits, then so do the supports of bidding strategies.

Convexity: Suppose that after some history h_t, player i's type k is assigned probability one. If both b and $b' > b$ are equilibrium bids for player i in period t, then any bid $b'' \in (b, b')$ also leads to degenerate beliefs on k.[15]

Examples of equilibria that do not satisfy these latter assumptions (thereby illustrating their 'necessity' for our result) are available from the authors. Note that our equilibrium in the case of one-sided uncertainty trivially satisfies the technical refinements, and can be easily seen to satisfy the full Markov condition given that no such s and $s' \neq s$ exist for the informed player.

Finally, we extend our tie-breaking rule in the one-sided case – namely that all ties are won by the uninformed. Note that in that case it was merely a convenience in order to avoid the introduction of a smallest unit of account. In particular, without that rule, the uninformed

[14] In fact, we shall simply assume without further ado that the payoffs in the one-sided case act as terminal payoffs in the two-sided case if beliefs have become degenerate on the highest type, which never happens in equilibrium anyway.

[15] Thus a related but generally stronger type of restriction that would work equally well concerns monotonicity of belief updating: if some bid b' leads to degenerate (i.e. unchanged) beliefs on type k, then any bid $b'' < b'$ should not lead to beliefs on higher types of player i.

would bid ε above the lowest serious bid α^I of the informed, the lowest type $\underline{k}_t$ of the informed doesn't want to win at a price above his valuation anyway, and higher types of the informed are indifferent between losing for sure versus winning and revealing [some of] their information (so they will continue to make losing bids $b \leqq \alpha^I$ with the appropriate probability.

So the natural analogue in the two-sided case is that if neither belief is degenerate (or both are), then each player has a 50% chance of winning in the case of a tie, while if only one player has degenerate beliefs (about his opponent), then that player loses ties (since he is 'informed'). This is fully incentive-compatible and merely avoids the inconvenience of introducing a unit of account. Therefore we adopt such a convention, which will be necessary in subgames that mimic the one-sided case.

Then we have the following:

Theorem 4 $\forall \varepsilon > 0 \;\; \exists \delta' < 1$*:* $\forall \delta \in (\delta', 1)$*, if* $p_1(m), p_2(n) < 1$ *then*

i) there exists an equilibrium: in particular, for any bid $b^* \in [0, \lambda]$ *(where* λ *is defined as above), the outcome in which all types of both players bid* b^* *in every period is supported by a P.B.E. (satisfying* ***R****, Markov, Continuity, and Convexity)*[16]*;*

ii) furthermore, in any P.B.E. satisfying ***R****, Markov, Continuity, and Convexity, there is a* $b^* \in [0, \lambda]$ *such that all types of both players receive equilibrium payoffs within* ε *of the payoffs corresponding to the outcome that involves pooling forever at* b^**.*

Proof. See the appendix.

So the theorem states that as long as neither player is believed with probability one to be his highest possible type, then the equilibrium outcome must, at least in terms of payoffs (and therefore also revenue for the seller), be essentially pooling at some level no larger than λ. If beliefs do not satisfy the condition of the theorem, i.e. they are one-sided degenerate on either type m or type n, then the equilibrium play must be as in the truly one-sided model described in Section 4 (in particular, not pooling and essentially fully revenue-extracting). But as the theorem states, if beliefs are one-sided degenerate on some lower type k of player 2 (say), then there are instead pooling equilibria at bids less than or equal to $\lambda = v(0, k)$. This shows clearly the distinction between knowledge (certainty) and mere belief. Finally, we point out that the theorem does apply in the degenerate belief case that $p_1(j) = 1$ *and* $p_2(k) = 1$, as long as both $j < m$ and $k < n$.

[16] Of course, as per its definition, **R** (like Convexity) only applies when beliefs on one side have become degenerate.

5.1 A simple example, continued

We extend the binary (i.e. $m = n = 1$) example developed in Section 4 to two-sided incomplete information. We will suppose that the parameters are symmetric, i.e. $p(j,k) = p(k,j)$, $v(j,k) = v(k,j)$, for all $j \in M$, $k \in N$. Let $p = \Pr\{S_i = 1\}$ denote the probability that a player observes a high signal. We continue assuming that types are independently distributed (so that p is indeed unconditional on one's own signal) and normalize $v(1,1) = 2$, $v(0,0) = 0$ and set $v(0,1) = 1$.[17] To simplify the exposition, we will only describe symmetric Markov stationary equilibria[18] in which bids are at least as high as what the object is commonly known to be worth (an assumption we refer to as *no underbidding*). This latter condition may be interpreted as a boundary condition that drives the payoff of one of the low types down to zero; clearly it is beneficial for the seller.[19] To get a better understanding of the role of discounting, we describe equilibrium strategies for all $\delta \in [0,1]$.

We call a mixed action profile *separating* if the low types of each player bid 0 (with probability one) and the high types continuously randomize over some interval $[0,\beta]$, with $\beta > 0$; as *semi-pooling* if low types bid a common bid p', while high types randomize on $[p',\beta]$, $\beta > p'$, with an atom at p'; and as *pooling* if both types of both players make a common bid p'. We say that *separation* occurs if player i's posterior, upon observing some particular bid (which we may then call a *separating* bid) by his opponent, is degenerate upon one specific type. Any other event is referred to as *pooling* (even when the posterior differs from the prior). As mentioned previously, we restrict attention to equilibria in which strategies are symmetric. If a high type separates (and thus becomes the uninformed player in the continuation subgame), we follow the equilibrium described in Section 4, except of course that all bids are shifted upward by 1. Note that this implies that all payoffs are the same as in the usual one-sided game.

If the discount factor is low ($\delta < \frac{1}{2}$) and the probability of the high signal is low ($p < 1 - 2\delta$), then mimicking a low type (i.e. bidding 0) is not profitable for a high type: by bidding p, he can win immediately at low cost, whereas mimicking the low type only pays off later, and delay is costly. The outcome in this case (in the first period) is thus identical to the outcome of the static game. However, if p is large ($p \geqq 1 - 2\delta$), winning immediately is too costly. Therefore, the high-signalled player mixes between a pooling bid (losing with higher probability but inducing his opponent to be more pessimistic, which is good for later profits) and a separating bid (reaping lower but immediate profits). If both players choose to submit the pooling bid, the high-signalled player faces a similar trade-off in the following period, and identical reasoning applies until either a separating bid is observed, or the common belief drops below the threshold $(1 - 2\delta)$ in which

[17] We have actually verified that the results admit natural extensions to the case without independence and with arbitrary $v(0,1) \in (v(0,0), v(1,1))$.

[18] By a (Markov stationary) symmetric equilibrium, we mean an equilibrium in strategies that specify the same mixed actions whenever beliefs are symmetric.

[19] Note in particular that it selects the highest (that is, at λ) fully pooling outcome in Theorem 3.

case play is as described before.

Equilibrium behavior for large discount factors ($\delta \geq 1/2$) is more intriguing. No matter how large the discount factor, there exists a pooling equilibrium only if p is not too large.[20] Indeed, if p is large relative to the discount factor, then being uninformed in a one-sided continuation game is a relatively profitable situation, as it implies winning often an object whose expected value is high (this follows from the limiting properties of the uninformed player's payoff π^U relative to p and δ). Submitting a separating bid is therefore attractive, even if it involves eventually bidding the true value of the object. Hence, if p is large relative to δ, the trade-off between the endless benefits from a pooling bid and the larger, finite-time rewards from a separating bid induces high-signalled players to semi-pool. In fact, if both realized bids are pooling, high-signalled players keep semi-pooling, and posterior beliefs asymptotically approach some limiting value. If the initial prior itself is below or equal to this threshold, players fully pool in all periods. Of course, when the discount factor tends to 1, this threshold also tends to 1 (so non-pooling outcomes disappear), in agreement with Theorem 3.

We summarize the previous discussion in the following Theorem.

Theorem 5 *If $m = n = 1$ and players are symmetric, there is a unique equilibrium distribution over outcomes (i.e. infinite histories) for all p and all $\delta < 1$. In equilibrium, while beliefs remain nondegenerate: for $\delta < 1/2$, actions are separating if $1 - p \geqq 2\delta$ and semi-pooling otherwise; for $\delta \geqq 1/2$, actions are pooling if $\pi^U(p) \leqq \delta - 1/2$ and semi-pooling otherwise.*

Proof. See the appendix.

This theorem is illustrated in Figure 3. The auctioneer's revenue in the static auction is $2p^2$, while it is p in the pooling equilibrium.[21] Therefore, the expected revenue in the static auction exceeds the expected revenue in the dynamic auction (for δ large enough) provided that the likelihood of the high signal is large enough.[22] The reason why pooling may generate larger revenues than the static equilibrium has been already described: since players learn nothing from winning, the winner's curse disappears, and the pooling bid is therefore larger than the lower extremity of the bidding support in the static equilibrium.

[20] Observe that this is consistent with Theorem 3: for fixed p, there exists $\bar{\delta}$ large enough such that the outcome is pooling if $\delta > \bar{\delta}$.

[21] The outcome in the separating equilibrium is the same as in the static auction, but the revenue is higher: it is $(1-\delta)S + \delta V$, where $S = 2p^2$ is the expected static revenue and $V = 2p$ is the expected value of the object.

[22] Keep in mind, however, that we have been assuming no underbidding in this section; therefore, what is derived here is an upper bound to the seller's payoff in a pooling equilibrium.

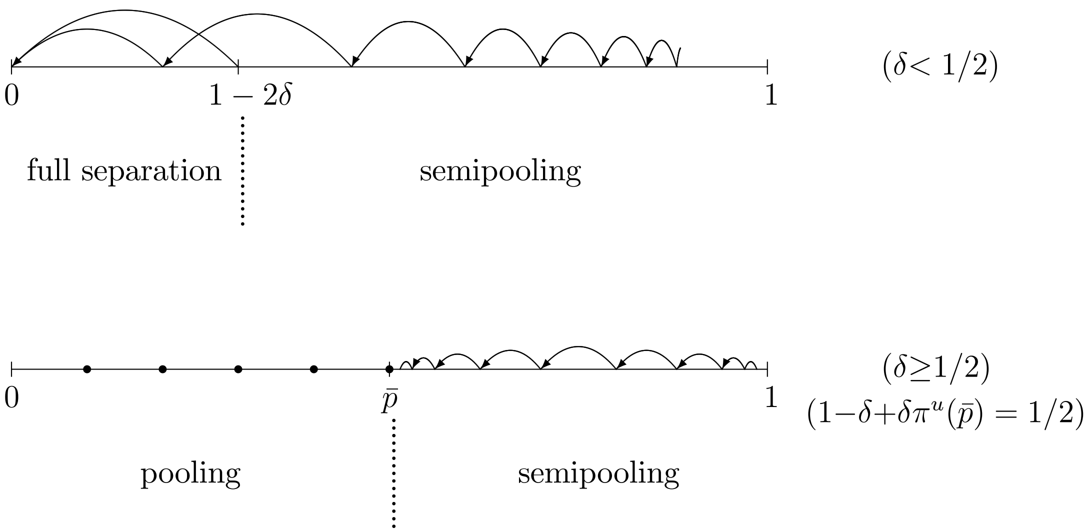

Figure 3: Phase portrait for $p_1 = p_2$

We have also considered the case of asymmetric priors: i.e., the case of $p_2 > p_1$. We do not go into the details here (they are available from the authors upon request), but it is perhaps worth pointing out that, although bidding is typically fairly complex, there is continuity with both the symmetric case and the static case (i.e. as δ goes to 0, equilibrium play converges to the static equilibrium). For $\delta \geq 1/2$, the equilibrium involves pooling below a boundary connecting the diagonal (i.e. the symmetric case) to either the edge that corresponds to $p_1 = 0$ (for $\delta < 2/3$) or to the edge corresponding to $p_2 = 1$ (for $\delta \geq 2/3$). The implications for revenue are similar to the symmetric case: the seller's payoff in the pooling region, p_1 (which is certainly lower than the overall expected value of $p_1 + p_2$), may or may not be lower than the revenue in the static game, $p_1^2 + p_2^2$. Intuitively, the more asymmetric that the two players are (either in terms of type-spaces or priors), the more confidently can we say that expected revenue is lower in the dynamic version of the auction.

5.2 Discussion of the refinements

There are two potential problems with the basic reasoning at the beginning of this section. First, measurability with respect to beliefs does not adequately capture the idea behind Markovian stationarity: namely, measurability with respect to a coarsest sufficient variable. The following example illustrates the problem.

Example 1: Player $i = 1, 2$ has two equiprobable signals: $v_H > v_L \geq 0$ and it is common knowledge that the object is worth at least 1. In the first T periods, the support of player 1's bid is $\{1/3, 1/2\}$, while player 2 bids $b_t \geq 1$ independently of his type. In period $t \leq T$, player 1's type v_H bids 1/2 with probability $m_t/ (m_t + 1)$, where m_t is the $(t + 1)$-th prime number, while type v_L bids 1/3 with probability $m_t/ (m_t + 1)$. After period T, along any equilibrium path, both players repeatedly bid 0. If any other bid is submitted by player $i = 1, 2$, his opponent assigns henceforth probability 1 player i being of type v_H, and the equilibrium with one-sided incomplete information follows. For each T and each bounded sequence $\{b_t\}$, we can choose δ large enough such that the payoff from following the equilibrium strategies exceeds the payoff from submitting any other bid.

For any $t' \leq T$, it is easy to see that the set of equilibrium histories $\bigcup_t H_{t \leq t'}$ can be mapped one-to-one into the set of possible beliefs about player 1 after equilibrium histories $H_{t'}$. Therefore, measurability with respect to beliefs imposes no additional restriction, for these equilibrium strategies. Observe that, for any $\varepsilon > 0$, we can choose δ sufficiently close to 1, and T and $\{b_t\}$ sufficiently large, so that the payoff of player 1 -and the payoff of player 2's low type- are below ε.

In this example, beliefs are used as a device to mark histories, and, in our opinion, are not directly "payoff-relevant": information is really never revealed, as all types of all players are willing to follow any sequence of actions on the equilibrium path. Ruling out multiple losing bids is not sufficient to eliminate the problem, as the next example shows.

Example 2: Player $i = 1, 2$ has two equiprobable signals: $v_H > v_L \geq 0$ and it is common knowledge that the object is worth at least 1. In periods 0, 3,..., $3t$ for $t \in \mathbb{N}_0$, player 1 wins the object, while in other periods player 2 wins the object. If a player submits an out-of-equilibrium bid, his high type is assigned probability one and the game proceeds in all periods as under one-sided incomplete information. If no out-of-equilibrium is ever submitted, bidding in periods 0, 3, 6... is independent from the bidding in the other periods. Specifically, bidding in periods 0, 3, 6... is recursively defined as follows. Player 2 bids 0 in all these periods. In period 0, player 1 continuously randomizes over the interval $[3/8, 5/8]$ centered around $c_0 = 1/2$, of length $l_0 = \delta/4$. While the support of the bid distribution is independent of his type, the specific distribution function is type-dependent. In period 3, player 1 continuously randomizes over the interval of length $l_1 = \delta/16$ centered around c_1, where c_1 solves $b_0 + \delta c_1 = c_0 + \delta c_0$, and $b_0 \in [3/8, 5/8]$ is the equilibrium bid observed in period 0. More generally, in period $3(t + 1)$, player 1 continuously randomizes over the interval of length $l_t = \delta/4^{t+2}$ centered around c_{t+1}, where c_{t+1} solves $b_t + \delta c_{t+1} = c_t + \delta c_t$ and b_t is the equilibrium bid observed in period $3t$. (See Figure 4)

It is easy to see that player 1 equilibriums bids in all such periods belongs to the interval $(0, 1)$, as the length of the union over all bid supports is bounded above by $\sum_{t=0}^{\infty} 2^t/4^{t+1} = 1/2$. In addition, the construction guarantees that player 1 is indifferent over all sequences of bids.

Finally, as the space of nondegenerate beliefs with two types, $(0, 1)$, and the space of sequences of equilibrium bids, included in $[0, 1]^{\mathbb{N}_0}$, are equinumerous, it follows that one can pick distribution functions such that the posterior probability about player 1's high type uniquely marks any equilibrium history.

A similar construction is performed in the other periods, with player 2 winning. As a result, player 2 wins 2/3 of the time, while player 1 wins only 1/3 of the time. Of course, this division was quite arbitrary: in particularly, for any $\varepsilon > 0$, we can pick $\delta > 1$ and construct an equilibrium such that player 1's payoff is less than or equal to ε.

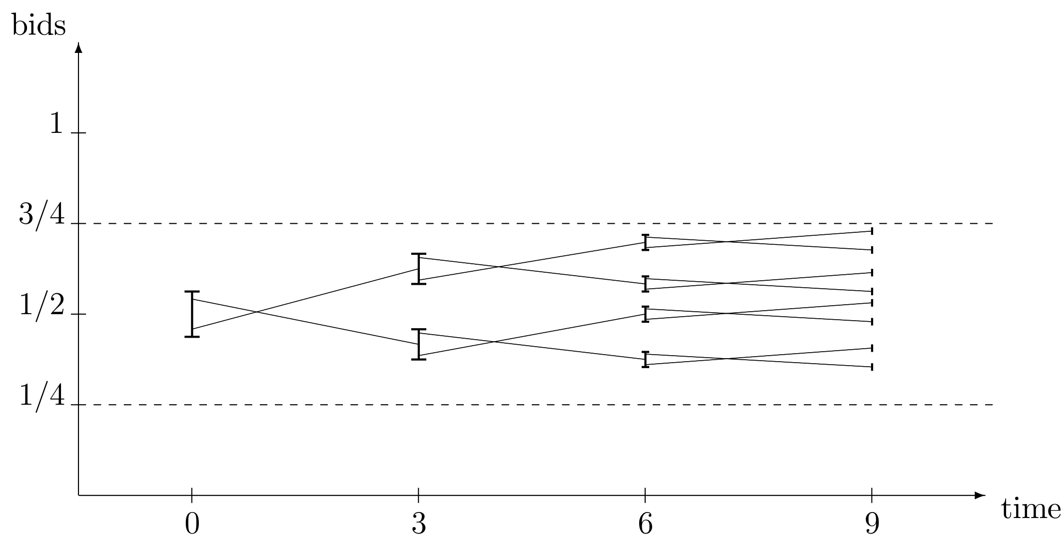

Figure 4: Example 2

In view of example 2, one could still hope for a weak characterization of Markov stationary equilibrium payoffs, by focusing on (a lower bound on) *the sum* of the players' payoffs, rather on each player's payoff. Yet it is possible to build on example 2 to construct examples in which both players' payoffs are arbitrarily low, by using the strategies described in the example as the continuation strategies in a subgame, inducing one of the players' high type to reveal his information with very high probability in the initial period.

Observe also that example 2 does not rely on the fact that the support of beliefs *never* changes. With three types for each player, for instance, some information could be eventually revealed, if, at time $T > 0$ say, player i's intermediate and high type submit one bid, while player i's low type submits another (and continuation strategies are pooling in either case). By choosing T large enough, the same features obtain.[23] To eliminate such equilibria, we therefore need to strengthen Markov stationarity beyond actions being measurable with respect to beliefs only. The second part of this Markov assumption strengthens the original myopia idea to be sure that strategies cannot be arbitrarily chosen across [indifferent] types in such a way as merely to mark time.

We view the Convexity restriction as potentially more serious, although it arises only in subgames in which one of the player's beliefs is degenerate on some type which is not the highest possible. It is illustrated in the following example.

> **Example 3**: Player $i = 1, 2$ has two possible signals: $v_H > v_L \geq 0$, which are equally likely for player 1, but player 2 is correctly believed to have signal v_L with probability one. In every period $t \geq 0$, player 1's type v_L bids 0, while player 1's type v_H either bids 0 with positive probability, or continuously randomizes over some interval (α_t, β_t), where $\beta_t > \alpha_t > 0$. Player 2 bids either 0 with positive probability, or α_t with positive probability, or continuously randomizes over the interval (α_t, β_t). If any other bid is submitted by player $i = 1, 2$, his opponent assigns henceforth probability 1 player i being of type v_H, and the equilibrium with one-sided incomplete information follows. The distributions and the extremities of the support are determined by various indifference conditions (including intertemporal ones). This equilibrium resembles the equilibrium of the game with one-sided incomplete information, but player 2, who is "uninformed", cannot bid ε rather than 0 to break the tie in his favor, because player 1 would believe him to be a high type. Therefore, player 1's high type is willing to bid 0 forever, as the probability of winning with such a bid is always positive, and there is no final time after which uncertainty resolves. If player 2 submits the bid 0 with very low probability however, player 1's high type may be exactly indifferent between "separating", by bidding $b \in (\alpha_t, \beta_t)$, and pooling with a bid of 0. Observe in particular that bidding $\alpha_t - \varepsilon$ rather than $\alpha_t + \varepsilon$ is not attractive, as player 2 bids α_t with discrete probability. As for player 2, bidding α_t is preferable to bidding $\alpha_t - \varepsilon$ because the posterior belief associated with the latter bid is bad for him. Although uncertainty may last forever, all information is revealed with probability one, and, for some (equilibrium) choices of $\{\alpha_t, \beta_t\}$, it is revealed rapidly relative to δ, so that the payoffs of patient players are arbitrarily low and the auctioneer's revenue arbitrarily close to its first-best maximum. Indeed, as for the

[23]This version of the example illustrates why it is necessary to use limits rather than exact equality in the Markov assumption.

other assumptions, bidders' payoffs can easily be made as small as 0 or as large as one-half their expected value (i.e. as if pooling at a bid of 0).

This equilibrium violates the intuition that a player's bidding support should be a convex set. A gap is generated in player 2's support by the use of threatening beliefs. To rule out such equilibria, we do not need to *impose* convex supports, but simply rule out a specific kind of non-monotonic beliefs associated with out-of-equilibrium bids; this is B2 above.

6 Concluding Comments

How could an auctioneer reduce or eliminate the tacit collusion exhibited by the equilibrium under two-sided incomplete information? A reserve price is an instrument that, used wisely, would allow the auctioneer to fare better. If he can commit to a reserve price policy, then as the discount factor tends to one, his optimal expected revenue tends toward his revenue from setting an optimal fixed reserve price. Players whose signal is sufficiently high pool at a level slightly above the reserve price, while players with lower signals remain idle. Therefore, the auctioneer's expected revenue is still lower than in the static auction. Finding the optimal sequence of reserve prices in a repeated auction is a formidable task; first steps have been taken by McAfee and Vincent (1997) and by Caillaud and Mezzetti (2003).

Another commonly used procedure in auctioneering is giving the winner the option to purchase future units at the current price (see Cassady 1967). It is clear that such a procedure eliminates the pooling equilibrium, as at least one player would have an incentive to bid a penny more and exercise his option. However, this procedure is not perfect for the seller either, as a player with a low signal may exercise his option and win all units at a price below their real value (assuming that bids are observed before the option decision is made). An analysis of the buyer's option in the context of repeated auctions, as well as risk aversion and the relationship of each to the declining price anomaly, may be found in Black and de Meza (1992).[24]

The auctioneer may also choose to vary the number of items in every lot put up for sale. Ashenfelter (1989) suggests that, in wine auctions, smaller lots are often offered before larger lots, in an attempt to disguise the price decline. Using data from Christie's, Ginsburgh and van Ours (2003) point to the fact that a sequence of lots each of which contains the same number of items seems to generate more revenue than lots with varying number of items.

The auctioneer could also decide, every time that bids are tied, to give the unit to the same (predetermined) bidder; this would obviously destroy the specific pooling equilibrium described in this paper. Another conceivable approach is for the seller to choose not to reveal bids at all. Many other tacks are possible, but, as mentioned, a careful analysis of their effect on equilibrium

[24] Ashenfelter's (1989) survey also discusses both reserve prices and the buyer's option.

strategies and seller's revenue is beyond the scope of this paper (for repeated auctions design, see Abdulkadiroğlu and Chung (2003)).

As in Kyle (1985) and the literature on insider trading, we have restricted attention to the case in which the values of the units are perfectly correlated. This seems to be the most challenging case for information revelation. Indeed, suppose that the value of each unit is an independent draw from some (possibly time-dependent) distribution, for which the static first-price auction admits a unique equilibrium. Then the only equilibrium that is stationary in the repeated game specifies that the static auction be played in each period.

We have restricted attention to the two player case. The pooling equilibrium remains an equilibrium when there are more than two players, but we have not proved uniqueness. As soon as a single player has revealed his information, all other players' private information must be eventually revealed, as such an uninformed bidder cannot be disciplined into pooling (given stationarity) and thus, informed bidders must eventually act. Such information revelation occurs "quickly" relative to the discount factor, which in turn allows players to enforce tacit collusion when none of them has revealed any information. The intuition for uniqueness appears robust: if none of the opponents reveals his private information, then pooling is certainly optimal with low discounting, while if some of them do reveal theirs, it is then still better to pool and become the informed player, for uninformed bidders have a zero payoff if there is more than one of them.

We have verified that the approach outlined in this paper also works within the framework of private values (rather than common) and of second-price auctions (rather than first-price). However, it is clear that considerable work remains to be done in order to provide an integrated theory of repeated auctions.

References

[1] Abdulkadiroğlu, A. and K.-S. Chung (2003), "Auction Design with Tacit Collusion", manuscript, Northwestern University.

[2] Ashenfelter, O. (1989), "How auctions work for wine and art", *Journal of Economic Perspectives*, 3(3), 23-36.

[3] Aumann, R.J. and M. Maschler, with the collaboration of R.E. Stearns (1966-68), Reports of the U.S. Arms Control and Disarmament Agency ST-80, ST-116, ST-143, Washington, D.C., reprinted in *Repeated Games of Incomplete Information*, M.I.T. Press, 1995, Cambridge, MA.

[4] Ausubel, L. and R. Deneckere (1989), "Reputation in Durable-Goods Monopoly", *Econometrica*, 57(3), 511-531.

[5] Baldwin, L., R. Marshall, and J.-F. Richard (1997), "Bidder Collusion at Forest Service Timber Sales", *The Journal of Political Economy*, 105(4), 657-699.

[6] Bikhchandani, S. (1988), "Reputation in Second-Price Auctions", *Journal of Economic Theory*, 46, 97-119.

[7] Black, J.M. and D. de Meza (1992), "Price Differences between Successive Auctions Are No Anomaly", *Journal of Economics and Management Strategy*, 1(4), 607-628.

[8] Caillaud, B. and C. Mezzetti (2003), "Equilibrium Reserve Prices in Sequential Ascending Auctions", manuscript, CERAS-ENPC Paris.

[9] Cassady, R. Jr. (1967), *Auctions and Auctioneering*, University of California Press, Los Angeles, CA.

[10] De Meyer, B. and H.M. Saley (2002), "On the strategic origin of Brownian motion in finance", *International Journal of Game Theory*, 31(2), 285-319.

[11] Devaney, R. (1989), *An Introduction to Chaotic Dynamical Systems*, 2nd Ed., Addison-Wesley, Reading, MA.

[12] Engelbrecht-Wiggans R. and R.J. Weber (1983), "A sequential auction involving asymmetrically-informed bidders", *International Journal of Game Theory*, 12(2), 123-127.

[13] Engelbrecht-Wiggans R., P. Milgrom, and R.J. Weber (1983), "Competitive Bidding and Proprietary Information", *Journal of Mathematical Economics*, 11, 161-169.

[14] Freixas, X., R. Guesnerie, and J. Tirole (1985), "Planning under incomplete Information and the Ratchet Effect", *Review of Economic Studies*, 52(2), 173-191.

[15] Fudenberg, D. and J. Tirole (1991), *Game Theory*, M.I.T. Press, Cambridge, MA.

[16] Ginsburgh, V. (1998), "Absentee Bidders and the Declining Price Anomaly in Wine Auctions", *The Journal of Political Economy*, 106(6), 1302-1319.

[17] Ginsburgh, V. and J. van Ours (2003), "How to Organize Sequential Auctions: Results of a Natural Experiment by Christie's", working paper, Tilburg University.

[18] Hart, O. and J. Tirole (1988), "Contract Renegotiation and Coasian Dynamics", *The Review of Economic Studies*, 55(4), 509-540.

[19] Hausch, D. (1986), "Multi-Object Auctions: Sequential vs. Simultaneous Sales", *Management Science*, 32(12), 1599-1610.

[20] Hon-Snir, S., D. Monderer, and A. Sela (1998), "A learning approach to auctions", *Journal of Economic Theory*, 82(1), 65-89.

[21] Kwiek, M. (2002), "Reputation and Collusion in the Repeated Second-price Auctions", Northwestern University.

[22] Kyle, A. (1985), "Continuous Auctions and Insider Trading", *Econometrica*, 53(6), 1315-1336.

[23] Maskin, E. and J. Riley (2000), "Equilibrium in Sealed High-bid Auctions", *The Review of Economic Studies*, 67(30), 455-83.

[24] Maskin, E. and J. Riley (2003), "Uniqueness of Equilibrium in Sealed High-Bid Auctions", Princeton.

[25] McAfee, R.P. and D. Vincent (1997), "Sequentially Optimal Auctions", *Games and Economic Behavior*, 18, 246-76.

[26] Milgrom, P. and R.J. Weber (1982), "A Theory of Auctions and Competitive Bidding", *Econometrica*, 50(6), 1089-1122.

[27] Oren, S. and M. Rothkopf (1975), "Optimal Bidding in Sequential Auctions", *Operations Research*, 23(6), 1080-90.

[28] Ortega-Reichert, A. (1968), "Models for Competitive Bidding under Uncertainty", Department of Operations Research Technical Report No. 8, Stanford University.

[29] Pesendorfer, M. (2000), "A Study of Collusion in First-Price Auctions", *The Review of Economic Studies*, 67, 381-411.

[30] Pezanis-Christou, P. (2000), "Sequential Descending-Price Auctions with Asymmetric Buyers: Evidence from a Fish Market", Pompeu Fabra.

[31] Raviv, Y. (2003), "New Evidence on Price Anomalies in Sequential Auctions", manuscript, Princeton University.

[32] Van den Berg, G., J. van Ours, and M. Pradhan (2001), "The Declining Price Anomaly in Dutch Dutch Rose Auctions", *American Economic Review*, 91(4), 1055-62.

[33] Virág, G. (2003), "Repeated Common Value Auctions", manuscript, Princeton University.

7 Appendix

7.1 Static case

Proof of Theorem 1: Because player 1's type $m(\kappa)$ is indifferent between bidding $\alpha(\kappa)$ and $\alpha(\kappa+1)$,

$$\begin{aligned}&\left(\sum_{j=0}^{n(\kappa)-1} q(j) + q(n(\kappa)) G_{n(\kappa)}(\alpha(\kappa))\right)(\alpha(\kappa+1)-\alpha(\kappa)) \\ = \;& q(n(\kappa))\left(G_{n(\kappa)}(\alpha(\kappa+1)) - G_{n(\kappa)}(\alpha(\kappa))\right) v(m(\kappa), n(\kappa)),\end{aligned}$$

and, similarly, because player 2's type $n(\kappa)$ is indifferent between bidding $\alpha(\kappa)$ and $\alpha(\kappa+1)$,

$$\begin{aligned}&\left(\sum_{i=0}^{m(\kappa)-1} p(i) + p(m(\kappa)) F_{m(\kappa)}(\alpha(\kappa))\right)(\alpha(\kappa+1)-\alpha(\kappa)) \\ = \;& p(m(\kappa))\left(F_{m(\kappa)}(\alpha(\kappa+1)) - F_{m(\kappa)}(\alpha(\kappa))\right) v(m(\kappa), n(\kappa)).\end{aligned}$$

Therefore, upon dividing,

$$\frac{q(n(\kappa))\left(G_{n(\kappa)}(\alpha(\kappa+1)) - G_{n(\kappa)}(\alpha(\kappa))\right)}{\sum_{j=0}^{n(\kappa)-1} q(j) + q(n(\kappa)) G_{n(\kappa)}(\alpha(\kappa))} = \frac{p(m(\kappa))\left(F_{m(\kappa)}(\alpha(\kappa+1)) - F_{m(\kappa)}(\alpha(\kappa))\right)}{\sum_{i=0}^{m(\kappa)-1} p(i) + p(m(\kappa)) F_{m(\kappa)}(\alpha(\kappa))}.$$

For $\kappa \leq m+n+1$, define $x(\kappa) = \sum_{i=0}^{m(\kappa)-1} p(i) + p(m(\kappa)) F_{m(\kappa)}(\alpha(\kappa))$ and $y(\kappa) = \sum_{i=0}^{n(\kappa)-1} q(j) + q(n(\kappa)) G_{n(\kappa)}(\alpha(\kappa))$. It follows that:

$$\frac{x(\kappa+1)-x(\kappa)}{x(\kappa)} = \frac{y(\kappa+1)-y(\kappa)}{y(\kappa)},$$

and thus, because $x(m+n+1) = y(m+n+1) = 1$, $x(\kappa) = y(\kappa)$, for all $\kappa \leq m+n+1$. Because for all $\alpha(\kappa+1)$, either $G_{n(\kappa)}(\alpha(\kappa+1))$ or $F_{m(\kappa)}(\alpha(\kappa+1))$ equals 1, this establishes that $s(\kappa) = x(\kappa)$ can be defined recursively as: $s(0) = 0$, and, for $P(i) = \sum_{l=0}^{i} p(l)$, $Q(j) = \sum_{l=0}^{j} q(l)$, recall that $s(\cdot)$ is recursively defined by $s(0) = 0$ and, for all $1 \leq \kappa \leq m+n+1$, for $\kappa - 1$ the largest integer for which $s(\kappa-1)$ has been defined so far, if $\min\{x \leq 1 \mid x = P(i) > s(\kappa-1) \text{ for some } i\} \neq \min\{x = Q(j) > s(\kappa-1) \text{ for some } j\}$, then $s(\kappa)$ equals the lowest of these minima. Otherwise, $s(\kappa) = s(\kappa+1)$ equals their common value. Obviously, $s(m+n+1) = 1$.

Consider again the indifference of player 2's type $n(\kappa)$ between bids $\alpha(\kappa)$ and $\alpha(\kappa+1)$:

$$\begin{aligned}&\sum_{i=0}^{m(\kappa)} p(i)(v(i, n(\kappa)) - \alpha(\kappa)) + p(m(\kappa)) F_{m(\kappa)}(\alpha(\kappa))(v(m(\kappa), n(\kappa)) - \alpha(\kappa)) \\ = \;& \sum_{i=0}^{m(\kappa)} p(i)(v(i, n(\kappa)) - \alpha(\kappa+1)) + p(m(\kappa)) F_{m(\kappa)}(\alpha(\kappa+1))(v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)).\end{aligned}$$

It follows that, for all $0 \leq \kappa \leq m+n$:

$$\frac{s(\kappa)}{s(\kappa+1)} = \frac{v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)}{v(m(\kappa), n(\kappa)) - \alpha(\kappa)},$$

and therefore, $\alpha(0) = 0$, and, for all $1 \leq \kappa \leq m+n+1$:

$$\alpha(\kappa) = \sum_{l=0}^{\kappa-1} (s(l+1) - s(l)) v(m(l), n(l)).$$

In the same way, we determine the distribution functions. If $b \in [\alpha(\kappa), \alpha(\kappa+1)]$,

$$\begin{aligned} p(m(\kappa)) F_{m(\kappa)}(b) &= s(\kappa+1) \frac{v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)}{v(m(\kappa), n(\kappa)) - b} - \sum_{j=0}^{m(\kappa)} p(i), \\ q(n(\kappa)) G_{n(\kappa)}(b) &= s(\kappa+1) \frac{v(m(\kappa), n(\kappa)) - \alpha(\kappa+1)}{v(m(\kappa), n(\kappa)) - b} - \sum_{j=0}^{n(\kappa)} q(j). \end{aligned}$$

We can now compute the expected revenue R. For $0 \leq \kappa \leq m+n$, let $R(\alpha(\kappa), \alpha(\kappa+1))$ be the expected revenue from bids $b \in [\alpha(\kappa), \alpha(\kappa+1)]$. Since $\sum_{j=0}^{n(\kappa)} q(j) + q(n(\kappa)) G_{n(\kappa)}(b) = \sum_{j=0}^{m(\kappa)} p(i) + p(m(\kappa)) F_{m(\kappa)}(b)$, it follows that

$$\begin{aligned} R(\alpha(\kappa), \alpha(\kappa+1)) &= s(\kappa+1)^2 \left(\frac{v(\kappa) - \alpha(\kappa+1)}{v(\kappa) - b} \right) (2b - v(\kappa)) \Big|_{\alpha(\kappa)}^{\alpha(\kappa+1)} \\ &= s(\kappa+1)^2 (2\alpha(\kappa+1) - v(\kappa)) - s(\kappa)^2 (2\alpha(\kappa) - v(\kappa)), \end{aligned}$$

where for simplicity, $v(\kappa) = v(m(\kappa), n(\kappa))$. It follows that

$$\begin{aligned} R &= \sum_{\kappa=0}^{m+n} R(\alpha(\kappa), \alpha(\kappa+1)) \\ &= 2\alpha(m+n+1) - \sum_{l=0}^{m+n} \left(s(l+1)^2 - s(l)^2 \right) v(l) \\ &= \sum_{l=0}^{m+n} 2(s(l+1) - s(l)) v(l) - \sum_{l=0}^{m+n} \left(s(l+1)^2 - s(l)^2 \right) v(l) \\ &= v(m+n+1) - \sum_{l=0}^{m+n} s(l)(2 - s(l))(v(l+1) - v(l)) \text{ (summation by parts)} \\ &= \sum_{l=0}^{m+n} (1 - s(l))^2 (v(l+1) - v(l)). \quad \blacksquare \end{aligned}$$

7.2 Binary one-sided case (calculations)

Additional details for the 2×1 case: The equilibrium bid distributions along histories $h_t \in H_t'$ are given by:

$$U_t(b) = \delta^{T-t}\frac{1-p_T}{1-b} \text{ and } p_t H_t(b) = \delta^{T-t}\frac{1-p_T}{1-b} - (1-p_t).$$

Define $\pi^I(p)$ and $\pi^U(p)$ to be the payoff of the informed and uninformed bidder, respectively, given belief p. Obviously, $\pi^I(p) = (1-\delta)\delta^T(1-p_T) = (1-\delta)\delta^{T/2}(1-p)^{\frac{1}{T+1}}$, where we define $T = \min\left\{t \in N; 1-p \geq \delta^{(T+1)(T+2)/2}\right\}$. It is simple to verify that $\pi^I(p)$ is decreasing in p, $\lim_{p\to 0}\pi^I(p) = 1-\delta$, $\lim_{p\to 1}\pi^I(p) = 0$, and $\pi^I(p)$ is decreasing in δ. Observe that $1-p_t = \left(\delta^{-t/2}(1-p)^{\frac{1}{T+1}}\right)^{T+1-t}$, and, denoting the odds ratio $p_t/(1-p_t)$ by l_t, we have $(p_t-\beta_t)/p_t = l_{t+1}/l_t$. By bidding 0 repeatedly, the uninformed player gets:

$$\begin{aligned}
\pi^U(p) &= (1-\delta)p\left[\frac{p_0-\beta_0}{p_0}\left(1+\delta\frac{p_1-\beta_1}{p_1}\left(1+\delta\frac{p_2-\beta_2}{p_2}\left(\cdots+\delta\frac{p_{T-1}-\beta_{T-1}}{p_{T-1}}\right)\right)\right)\right] \\
&= (1-\delta)\delta^T\left[\sum_{t=1}^{T}\delta^{\frac{(t-2)(T+1-t)}{2}}(1-p)^{\frac{t}{T+1}}\right] - (1-p)\left(1-\delta^T\right).
\end{aligned}$$

Observe that, because $T = \min\left\{t \in N; 1-p \geq \delta^{(T+1)(T+2)/2}\right\}$, $\delta^T \to 1$ as $\delta \to 1$. Notice also that

$$\pi^U(p) = (1-\delta)p\left[\frac{l_1}{l_0} + \delta\frac{l_2}{l_0} + \cdots + \delta^{T-1}\frac{l_T}{l_0}\right] \leq p\left(1-\delta^T\right),$$

so that $\pi^U(p) \to 0$ as $\delta \to 1$.

Finally, we study the variations of the expected bids. Let h_t, u_t be the densities correspondingly respectively to the distributions H_t and U_t. The density of the maximum bid, m_t, is given by

$$m_t(b) = u(t)(1-p_t+p_t H_t(b)) + p_t h_t(b) U_t(b) = \frac{2\delta^{2(T-t)}(1-p_T)^2}{(1-b)^3},$$

and its expectation (conditional, as usual, on the informed bidder having bid 0 up to $t-1$) is

$$\begin{aligned}
E_t &= 2\delta^{2(T-t)}(1-p_T)^2\int_0^{1-\delta^{T-t}(1-p_T)}\frac{tdt}{(1-t)^3} \\
&= 2\delta^{2(T-t)}(1-p_T)^2\left(\frac{1}{2}+\frac{1-2\delta^{T-t}(1-p_T)}{2\left(\delta^{T-t}(1-p_T)\right)^2}\right) \\
&= \left(1-\delta^{T-t}(1-p_T)\right)^2,
\end{aligned}$$

which is decreasing in t. Next, observe that the *unconditional* expectation of the winning bid in period $t \geq 1$, $t \leq T$, F_t, satisfies

$$\begin{aligned} F_t &= 1 - \prod_{i=0}^{t-1} (1 - \beta_i)(1 - E_t) \\ &= 1 - (1 - p_T)(1 - E_t) \sum_{i=0}^{t-1} \delta^{T-i} \\ &= 1 - \delta^T \left(1 - p^T\right) \frac{\delta^{-t} - 1}{\delta^{-1} - 1} \left(1 - \left(1 - \delta^{T-t}(1 - p_T)\right)^2\right), \end{aligned}$$

which is decreasing in t as well. Of course, for $t > T$, it is equal to the prior, p, and is larger than the corresponding expectation for all $t \leq T$.

7.3 Two-sided incomplete information

Proof of Theorem 3: We first prove existence by showing that a fully pooling equilibrium outcome (from the first period on) exists at any bid level $b^* \leq \lambda$. To see that this is the case, simply pick some such b^* and assume that any player who bids strictly more than b^* is thought to be the highest type possible of that player (i.e. either m or n respectively), so that bids immediately jump to a strictly larger level and the one-sided separating equilibrium follows.[25] For large enough δ, Theorem 2 implies that his continuation payoffs are (at most) approaching zero[26], whereas all types of both players – except possibly the lowest types – are making strictly positive profits along the equilibrium path, and in particular profits that are bounded away from zero independently of δ. For a type 0 who expects zero profits anyway (which only occurs if $b^* = \lambda$), bidding above b^* leads to a loss in the current period (because his expected value for the good must be exactly λ) and continuation payoffs of at most zero, because all bids from then on will be above his value. Hence there is no incentive for him to deviate either. This particular updating is not unreasonable (the highest type has the strongest incentive to try to win at any stage), but note that the equilibrium outcome survives with much weaker protocols: any revised distribution that first-order stochastically dominates the prior (i.e. such that high bids are "good news" about one's opponent's type) will lead to larger expected values and thus typically a higher pooling bid in the continuation. For high enough δ, this outweighs any possible one-shot gains. Finally, note that all of these equilibria trivially satisfy our assumptions.

We prove the second half of the theorem by induction on the total number $l = m' + n'$ of types who are actually present, where m' (resp. n') is the number of types in the support of

[25] Of course, if the deviant player is not actually of the highest type, then he's not willing to make the high bids required in the continuation game, but that simply means that he will make sure he loses in every period, and thereby earn nothing. His opponent, however, will continue to believe that he is the highest type.

[26] We choose the continuation equilibrium which satisfies $\mathbf{R}$ and hence for which this is true.

player 2's (resp. 1's) beliefs about player 1 (resp. 2). In fact, we will need to prove a slightly stronger statement than is given in the theorem: we also claim that in any equilibrium the total discounted probability of winning is arbitrarily close to 1/2 for each player (so in fact it really does look very much like pooling). We proceed in several steps.

1. The basis of the induction is $l = 2$, in which case the condition in the statement of the theorem implies that we must have $p_1(j) = p_2(k) = 1$, with $j < m$ and $k < n$. Since belief supports are singletons here, beliefs can never change in equilibrium and thus by Markov stationarity any equilibrium must involve the same strategies in every period. Neither player is willing to lose forever unless his opponent is bidding at least $\lambda = v(j, k)$, and neither player is willing to bid λ if his opponent is bidding less than λ (since bidding λ yields zero profit but by deviating and winning at a lower price he could get a positive profit).[27] Hence they must be making the same bid $b^* \leq \lambda$, and by the above argument we know that all such pooling outcomes are indeed supported by equilibria. Therefore in this case we must actually have full pooling immediately, and ε is irrelevant since the bound on payoffs is exact.

2. We now assume $l > 2$ and that the inductive hypothesis holds. If in equilibrium the support of beliefs never changes, then either we have full pooling from period 0 onward, or one of the players has a degenerate belief. The reasoning is as follows: if this condition holds, all types of each player are using the same strategies in every period, so by taking ϵ arbitrarily small in the Markov assumption, we see that all types must in fact be using those same strategies with exactly the same probabilities, meaning that beliefs themselves (not just the supports thereof) never change. Hence, by the first part of Markov, the equilibrium itself must be stationary. No type of either player is willing to lose forever (unless perhaps he is the only possible type on his side, in which case the other player's beliefs are indeed degenerate), so both players must be making only serious bids. If each player is making only a single serious bid, then they are the same and full pooling occurs from period 0 onward (at some bid $b^* \leq \lambda$). But if a player is making more than one serious bid, and is doing the same thing in every period, then single-crossing would imply that not all of his types can be indifferent between all equilibrium bids – unless he has only one type, as claimed.

3. If the support of beliefs never changes, and one of the players (say player 1) has degenerate beliefs, then once again we have full pooling from period 0 on. This must also be a stationary equilibrium, and $l > 2$ implies that there are at least two types of player 1, so he can't be losing with probability one in every period. But now single-crossing applied to the various possible types of player 1 implies that he must be making only a single bid $b^* \leq \lambda$ (where the inequality follows since player 1's lowest type is willing to make the bid b^* and it wins with positive probability). Player 2 behaves myopically within his equilibrium bidding support, so he will never bid below b^* and of course he can't be bidding only above b^*. If he bids both b^* and some $b > b^*$ with positive

[27]Note that they cannot be using nontrivial mixed strategies either, as standard results from the analysis of Bertrand competition with symmetric constant marginal costs tell us that these only exist if demand is nonzero at arbitrarily high prices, and in our context there is a lower bound of 0 on bids. More precisely, you could make a negative bid, but the seller wouldn't pay you if you won.

probability, then Convexity implies that bidding $(b^* + b)/2$ leads to the same beliefs and hence the same continuation value, in which case he would strictly prefer it to bidding b, given player 1's behavior (i.e. pooling at b^*). But this would contradict that b is in his equilibrium bidding support, so in fact he too must bid only b^*, as claimed.

4. Therefore either we are done, or there is some positive probability that at least one player's belief support changes. Without loss of generality we will take it to be player 2, i.e. there is some chance that player 1 partially separates. Let t^* be the first such period in which this occurs. We will further assume for the moment that no bids by player 1 in period t^* are made by only his highest possible type, m. We will prove soon that this highest type is never willing to separate by himself, so this assumption will in fact be valid in any equilibrium. If player 1 does make a bid $\hat{b}$ that causes player 2's belief support to change, then by this assumption we still have $p_1(m) < 1$, and thus the inductive hypothesis applies, with some corresponding 'effective pooling bid level' b^*. Finally, we note that in this eventuality at least two types of player 1 have non-overlapping strategy supports, so by the Markov refinement those strategies must lead to nontrivially distinct probabilities of winning, and hence must differ very quickly (relative to δ). Thus we may take $\delta^{t^*} \approx 1$ in what follows.

5. But this means that all types of player 1, whether or not they bid $\hat{b}$ in equilibrium (in period t^*), know that they could in fact bid $\hat{b}$ and thereby receive payoffs very close to those corresponding to pooling at b^*: this follows since after bidding $\hat{b}$ the probability of winning for each type of each player must be essentially $1/2$ in the sequel. Now consider some type j of player 1 who does bid $\hat{b}$: for any other bid b in the combined support of player 1's period 0 bidding strategy, either type j also makes that bid in equilibrium, or he does not. In the former case he's obviously indifferent between the two, and in the latter case the bid b also leads to a change in player 2's belief support, so the same reasoning as above applies in reverse; in particular for large δ the 'effective pooling bid level' that corresponds to b must be very close to the actual b^* corresponding to $\hat{b}$. Either way, for large δ, all types of player 1 must be almost indifferent between *all* bids in player 1's bidding support, whether or not they make them in equilibrium. Thus, for any $\varepsilon > 0$, we can choose δ sufficiently close to 1 so that each type of player 1 expects profits within ε of his profits in the b^* pooling equilibrium. It is also clear at this point that the analogous discounted probabilities of winning must be arbitrarily close to $1/2$ for any bid that player 1 makes in equilibrium, since there are in fact multiple types of player 1 in player 2's belief support, and single-crossing implies that they couldn't all be nearly indifferent otherwise.[28]

6. It remains only to show that the highest types do not make fully separating bids, i.e. bids that would completely reveal their type. First, it is clear that if there is any probability that player 1 makes a partially separating bid (i.e. one that causes player 2's belief support to change) in period t^* that leads to beliefs other than degenerate on type m, then the highest type cannot

[28]Note that none of this analysis depended on whether or not player 2 was pooling, or indeed in any way on what player 2 was doing. This is because player 1 can't affect player 2's behavior, and the inductive hypothesis applies with boundedly positive probability (given that $p_2(n) < 1$).

be willing to fully separate. This is because the inductive hypothesis implies that (at least with positive probability, depending on player 2's action) such a bid would be followed by something essentially like pooling, and the type m would make boundedly positive profits in such a situation – since he is in fact the highest type – which for large δ is preferable to the vanishingly small profits that he receives according to Theorem 2 after separating.[29] Similarly, there can be no chance that player 2 makes a bid that leads player 1 to have a strictly smaller belief support (other than degenerate on type n), as otherwise player 1 type m would strictly prefer to pool in period 0 and hope for that bid (which again by induction would lead to something like pooling forever) rather than separate and get very little.

7. Therefore, either player 1 type m does not separate in period t^* (and we're done), or all equilibrium bids by each player are made by all types of that player, except possibly for some additional bids made only by type m (and/or type n respectively). This must go on in every period, since if there were ever a chance of receiving an essentially pooling outcome in the future, both highest types would prefer to wait for the possibility of that. Now since $p_1(m), p_2(n) < 1$, the probability that no highest type ever separates is bounded away from zero, and both players know this (whatever their own type). Conditional on that event, at least one of the players must have a discounted expected probability of winning that is at least $1/2$.[30] But then, for large enough δ, the highest type of that player would prefer to facilitate that outcome, which gives him strictly positive surplus since all types of that same player are willing to make those same bids (and there are in fact some lower types), than to make a revealing bid and hence receive an arbitrarily small payoff. Since our supposition (toward a contradiction) is that player 1 type m is willing to separate in period t^*, it must be that player 2 type n does not make fully separating bids in equilibrium. In particular, player 1's belief support never changes, and thus player 1 type m must actually be making some separating bids in every period, or the Markov refinement would apply and we would be done (as in steps **2-3**).

8. Since all types $j < m$ of player 1 are willing to make all of player 1's equilibrium bids (other than type m's separating bids) in every period, we can once again apply single-crossing, and in this case we deduce that they must all be pooling on a single bid α_t in each period; call this strategy $\hat{s}$. Furthermore, we can apply the same reasoning to player 2, since all of his types are making all bids in equilibrium: single-crossing implies that in each period they must all be making a single bid β_t. Since player 1 type m is willing to separate in any given period, but he is also willing to play $\hat{s}$ (which is played by lower types as well), it must be that $\hat{s}$ yields a total

[29] Technically, Theorem 2 shows that for any fixed prior beliefs, as δ goes to 1, bidders' profits go to 0. However, it can be shown that this limit is in fact uniform in the priors; this is clearly seen in the specific calculations for the binary case, as detailed in the appendix. Also, note that technically Theorem 2 does not apply if player 2 happens to fully separate in the same period, but in that case type m gets a payoff of exactly 0, which is indeed bad.

[30] It does not follow that this player, conditioning on his opponent never revealing as a highest type, sees a discounted probability $p \geqq 1/2$; but it does follow that in that case he sees a p that is boundedly greater than 0, which is all that we require.

discounted probability of winning that is positive but very near 0. Hence β_t must be greater than α_t in most periods, but not all periods, and in particular there must be arbitrarily late periods t in which $\beta_t \leq \alpha_t$ but $\beta_{t+1} > \alpha_{t+1}$. If player 1 type m is indifferent between separating in period t, versus pooling in period t (and receiving a boundedly positive instantaneous payoff) and then separating in period $t+1$, it must be that $\beta_{t+1} - \beta_t$ (for such periods) is uniformly bounded below (away from 0).[31] This contradicts the Continuity assumption, since beliefs are obviously converging over time.

Thus player 1 type m is never willing to separate after all, and the proof is complete. ■

7.4 Symmetric binary types

Proof of Theorem 4: For a separating equilibrium to exist, the only incentive constraint to verify is that high types have no incentive to mimic low types. By deviating, a high type gets a payoff of $(1-\delta)(1-p)/2 + \delta(1-\delta)$. By abiding by the equilibrium strategy, he gets $(1-\delta)(1-p)$. Therefore, it is necessary and sufficient that $2\delta \leq 1-p$. For a pooling equilibrium to exist, the only incentive constraint to verify is that a high type does not want to deviate and bid slightly more. By deviating, he can get up to $1-\delta+\delta\pi^U(p)$, and by abiding by the equilibrium, he gets $\frac{1}{2}$. Therefore, a necessary and sufficient condition for a pooling equilibrium to exist is that $1-\delta+\delta\pi^U(p) \leq \frac{1}{2}$.

Consider now a semi-pooling equilibrium. Let γ be the probability that one's opponent is of the high type and makes a bid larger than p' (a *separating* bid), which is both the lowest bid (the *pooling* bid) and the posterior belief about the opponent's type, conditional on the pooling being observed. Then, for high types to be indifferent between the pooling bid and a slightly higher bid, we need:

$$\begin{aligned} V(p) &= \underbrace{\gamma\delta\pi^I(p') + (1-\gamma)\left((1-\delta)\frac{1}{2} + \delta V(p')\right)}_{\text{payoff from bidding } p'} \\ &= \underbrace{(1-\gamma)\left((1-\delta)+\delta\pi^U(p')\right)}_{\text{payoff from bidding } p'+\varepsilon,\text{ for small } \varepsilon>0}, \end{aligned} \tag{1}$$

where $V(p)$ is the value given belief p, and $V(p')$ is the value given p' (uniqueness will be shown), and Bayes' rule gives that

$$p' = \frac{p-\gamma}{1-\gamma}, \text{ or } 1-\gamma = \frac{1-p}{1-p'}.$$

[31] Indeed, for many parameter values, it won't be possible to make him indifferent at all.

Rearranging gives

$$1-p = \frac{\delta\pi^{I}(p')}{(1-\delta)\frac{1}{2}+\delta\left(\pi^{U}(p')+\pi^{I}(p')-V(p')\right)}\left(1-p'\right), \tag{2}$$

$$V(p) = \frac{\delta\pi^{I}(p')\left(1-\delta+\delta\pi^{U}(p')\right)}{(1-\delta)\frac{1}{2}+\delta\left(\pi^{U}(p')+\pi^{I}(p')-V(p')\right)}.$$

We show that such an equilibrium cannot exist for $\delta<1/2$ and $2\delta\leq 1-p$ (the latter clearly implies the former), nor can it exist for $\delta\geq 1/2$ and $1-\delta+\delta\pi^{U}(p)\leq\frac{1}{2}$ (same remark). Consider the first case. Observe that a (decreasing) sequence $(p,p',p'',\ldots)$ of consecutive semi-pooling equilibria cannot have an accumulation point, for by picking a term arbitrarily far in the sequence, its value $V(p)$ would be arbitrarily close to $\frac{1}{2}$, while the deviation payoff for a high type would be arbitrarily close to $1-\delta+\delta\pi^{U}(p)>\frac{1}{2}$, a contradiction (this argument covers also the case where the alleged accumulation point is 0); therefore, any sequence of consecutive equilibria must have a largest term, after which, by necessity, the equilibrium is separating. Pick the largest such term. Equation (1) implies that

$$\left(p-p'\right)\delta\pi^{I}\left(p'\right)+(1-p)\left(\delta(1-\delta)\left(1-p'\right)\right)=(1-p)\left((1-\delta)\frac{1}{2}+\delta\pi^{U}\left(p'\right)\right).$$

Because $\pi^{I}(p')\leq 1-\delta$ and $\pi^{U}(p')\geq 0$, this requires

$$1-p'(2-p)\geq\frac{1-p}{2\delta}.$$

However, this is impossible if $2\delta<1-p$, and only possible for $2\delta=1-p$ if $p'=0$, that is, if the equilibrium is in fact separating. Consider now the second case. Suppose first that there is a sequence $(p,p',p'',\ldots)$ with infinitely many terms that are semi-pooling. Then we can pick terms such that the payoff must be arbitrarily close to $1/2$ and $1-\delta+\delta\pi^{U}(p)$, implying that the limit of such a sequence cannot be such that $1-\delta+\delta\pi^{U}(p)<\frac{1}{2}$. We can then suppose without loss of generality that, if a semi-pooling exists with $1-\delta+\delta\pi^{U}(p)<\frac{1}{2}$, then (conditional on both players making a pooling bid), players pool (forever) at belief p'. Equation (1) then implies that

$$\left(p-p'\right)\delta\pi^{I}\left(p'\right)+(1-p)\frac{1}{2}=(1-p)\left(1-\delta+\delta\pi^{U}\left(p'\right)\right), \tag{1'}$$

which is impossible because $1-\delta+\delta\pi^{U}(p')<1-\delta+\delta\pi^{U}(p)<\frac{1}{2}$, and $\pi^{I}(p')\geq 0$. In fact, the same argument rules out a semi-pooling equilibrium with p such that $1-\delta+\delta\pi^{U}(p)=\frac{1}{2}$. This establishes the second claim. It may be worthwhile at this point to observe that, if a semi-pooling equilibrium exists for $1-\delta+\delta\pi^{U}(p)>\frac{1}{2}$, then the sequence of continuation plays (conditional on players making the pooling bid) is an infinite sequence consisting exclusively of semi-pooling equilibria. If there is a limit, it must obviously satisfy $1-\delta+\delta\pi^{U}(p)=\frac{1}{2}$, so this limit is unique. As pooling equilibria cannot exist for $1-\delta+\delta\pi^{U}(p)>\frac{1}{2}$, all we have to show is that such a sequence cannot be finite, i.e. such that there exists p, after which (conditional on both players making pooling bids) the equilibrium is pooling (i.e., p' is such that $1-\delta+\delta\pi^U(p') \leq \frac{1}{2}$). This immediately follows from (1').

We are left with establishing existence (and uniqueness) of semi-pooling equilibria for $\delta < 1/2$ and $2\delta > 1-p$, as well as for $\delta > 1/2$ and $1-\delta+\delta\pi^U(p) > \frac{1}{2}$. As shown, if such a semi-pooling equilibrium exists, then it must belong to a sequence of semi-pooling equilibria which is finite in the first case, ending up with a separating equilibrium, and infinite in the second case, converging to the unique p solving $1-\delta+\delta\pi^U(p) = \frac{1}{2}$.

Consider first $\delta < 1/2$ and $2\delta > 1-p$. We show that (2) determines uniquely $(V(p), p)$ as a function of $(V(p'), p')$, and that the (first coordinate of the) sets of consecutive solutions of (2), with boundary conditions given by $(p, (1-\delta)(1-p))$ for $p \in (0, 1-2\delta]$ partition $(1-2\delta, 1)$. To see this, suppose that $V(p')$ is decreasing in p' and smaller than $(1-\delta)$. It follows from the first equation of (2) that $(1-p)/(1-p')$ is decreasing in p'. This follows from the fact that

$$\frac{(1-\delta)\frac{1}{2}+\delta\left(\pi^U(p')+\pi^I(p')-V(p')\right)}{\delta\pi^I(p')} = 1+\frac{\pi^U(p')}{\pi^I(p')}+\frac{(1-\delta)\frac{1}{2}-\delta V(p')}{\delta\pi^I(p')}$$

is increasing (π^U and $1/\pi^I$ are increasing and positive, and so is $(1-\delta)\frac{1}{2}-\delta V(p')$, because $\frac{1}{2} > \delta$, and, by hypothesis, $(1-\delta) \geq V(p')$, and $V(p')$ is decreasing. The fact that is $(1-p)/(1-p')$ is decreasing in p' implies in particular that p increases with p', but also, from $(1')$, since

$$V(p) = \left(1-\frac{1-p}{1-p'}\right)\delta\pi^I(p') + \frac{1-p}{1-p'}\left((1-\delta)\frac{1}{2}+\delta V(p')\right),$$

as $\delta\pi^I(p') \leq (1-\delta)\frac{1}{2} \leq (1-\delta)\frac{1}{2}+\delta V(p')$, and both $\delta\pi^I(p')$ and $(1-\delta)\frac{1}{2}+\delta V(p')$ are decreasing in p', that $V(p)$ decreases with p' (as a weighted average of two decreasing functions with increasing weight on the smaller one).Because $V(p)$ is decreasing, it follows in particular that $V(p) \leq 1-\delta$, provided that V is continuous, which follows by induction as well once it is established for the first iteration. It is immediate to verify that for $p' = \varepsilon > 0$ arbitrarily small, and associated $V(p') = (1-\delta)(1-p')$, there exists $(p, V(p))$ solving (2) arbitrarily close to $(1-2\delta, 2\delta(1-\delta))$. Because (the projection on the first coordinate space of) the image by (2) of the interval $(0, 1-2\delta]$ is an interval $(1-2\delta, p^*]$, for some $p^* > 1-2\delta$, and the value V is continuous on that interval, it follows by induction that the intervals of probabilities constructed this way have neither "gaps" nor "overlaps", and that V is continuously decreasing in p. Observe that, as $p' = 1$ is a fixed point of the first equation of (2), the union of these intervals never stretches above one. Conversely, because, for $p' \geq 2\delta(1-\delta)$ (which is certainly a probability that is "reached", since it belongs to the "first" interval),

$$1-p = \frac{\delta\pi^I(p')}{(1-\delta)\frac{1}{2}+\delta(\pi^U(p')+\pi^I(p')-V(p'))}\left(1-p'\right) < \frac{1-p'}{1+(1-\delta)(1-2\delta)},$$

for any $p < 1$ is eventually included in the union of intervals recursively obtained by application of (2). This proves that for every $p > 1-2\delta$, there exists one and only one equilibrium outcome, specifying in particular that players semi-pool as long as they have pooled, until the common belief p is less than $1-2\delta$, at which point separation occurs.

Let us now study the case $\delta \geq 1/2$ and $1-\delta+\delta\pi^U(p) > \frac{1}{2}$. Defining $q = 1/(1-p)$, $w = V(p)/(1-p)$, and $f(q) = 1-\delta+\delta\pi^U(1-1/q)$, we get from (2) the following pair of difference equations

$$\begin{cases} q_{n+1} - q_n = \frac{\left(f(q_n)-\frac{1-\delta}{2}\right)q_n - \delta w_n}{\delta\pi^I(q_n)} \\ w_{n+1} = f(q_n), \end{cases}$$

where calendar time is reversed, that is, q_n corresponds to the posterior belief given semi-pooling and a prior belief q_{n+1}. Observe that the unique critical point of this system is $\bar{x} := (\bar{q}, \bar{w}) = \left((1-\bar{p})^{-1}, w(\bar{p})\right)$, where $\bar{p}$ is the unique root of $1-\delta+\delta\pi^U(p) = \frac{1}{2}$. For later use, observe also that this system is equivalent to the second-order difference equation

$$q_{n+1} - q_n = \frac{(f(q_n) - (1-\delta)/2)\, q_n - \delta f(q_{n-1})}{\delta\pi^I(q_n)},$$

from which it is apparent that $q_n > q_{n-1} > \bar{q}$ implies that also $q_{n+1} > q_n > \bar{q}$ (because in that case $(f(q_n) - (1-\delta)/2)\, q_n - \delta f(q_{n-1}) > (1-\delta)\left(f(q_n) - \frac{1}{2}\right)$). Computing the Jacobian evaluated at this fixed point, we get:

$$\begin{bmatrix} 1 + \frac{f'(\bar{q})\bar{q} + f(\bar{q}) - (1-\delta)/2}{\delta\pi^I(\bar{q})} & -\frac{1}{\pi^I(\bar{q})} \\ f'(\bar{q}) & 0 \end{bmatrix},$$

whose roots are real conjugate, one of which has modulus strictly less than one, the other one has modulus strictly larger than one. Indeed, the discriminant is positive because $\left(1 + \frac{f'(\bar{q})\bar{q} + f(\bar{q}) - (1-\delta)/2}{\delta\pi^I(\bar{q})}\right)^2 > 4f'(\bar{q})/\pi^I(\bar{q})$, as the term which is squared exceeds $\left(1 + f'(\bar{q})/\pi^I(\bar{q})\right)^2$ (using $f(\bar{q}) - (1-\delta)/2 > 0$ and $\bar{q}/\delta > 1$), and the ordering of moduli is easily established using the same bounds. Therefore, $\bar{x}$ is a hyperbolic fixed point, the map defined by the system of difference equations has a saddle point at $\bar{x}$, and so has its inverse map (the eigenvalues of the inverse matrix are the inverses of the eigenvalues). By the stable manifold theorem (see Devaney (1989)), there exists a neighborhood of $\bar{x}$ such that, for each q in this neighborhood, there exists a unique w such that the limit of the system starting from (q, w) is $\bar{x}$. Because $q_n = (1-p_n)^{-1}$ is strictly increasing in p_n, and $w_n = f(q_{n-1})$ is similarly increasing in p_{n-1}, we may therefore conclude that, using standard calendar time, there exists a neighborhood of $\bar{p}$, such that, for each p_n in this neighborhood, there exists a unique p_{n+1} in this neighborhood, such that the sequence (p_k) going consecutively through p_n and p_{n+1} tends to $\bar{p}$. Evidently, $p_{n+1} > \bar{p}$. As p_{-n} is monotonic since q_{-n} is, and $\bar{p}$ is the unique fixed point of the second-order difference equation, it follows that through all points $p \in (\bar{p}, 1)$ such a sequence exists, and uniqueness follows from trivial continuity and monotonicity observations. ■