Public vs. Private Offers in the Market for Lemons

Johannes Hörner and Nicolas Vieille

July 28, 2007

Abstract

We study the role of observability in bargaining with correlated values. Short-run buyers sequentially submit offers to one seller. When previous offers are observable, bargaining is likely to end up in an impasse. In contrast, when offers are hidden, agreement is always reached, although with delay.

1 Introduction

We study the role of observability in bargaining with correlated values. More precisely, we study how the information available to potential buyers affects the probability of reaching an agreement. Our main result is that, if discounting is low and the static incentive constraints preclude first-best efficiency, agreement is always reached when previous offers are kept hidden, while agreement is unlikely to be reached when they are made public.

The structure is as in Akerlof's market for 'lemons'. One seller is better informed than the potential buyers about the value of the single unit for sale. It is common knowledge that a mutually beneficial trade exists. All potential buyers share the same valuation for the unit, which is strictly larger than the seller's cost. The seller bargains sequentially with potential buyers until agreement is reached, if ever, and delay is costly. Each buyer makes a take-it-or-leave-it offer to the seller. That is, the setting is formally a search model: the seller may rationally choose to reject available offers in return for the opportunity to wait for better prospective offers in the future. This search process is without recall.

To take a specific example, consider the sale of residential property. In most countries, houses are sold through bilateral bargaining. Potential buyers come and go, engaging in private negotiations with the seller, until either an agreement with one of them is reached or the house is withdrawn from the market. Typically, potential buyers know the time on market of the house on sale, which provides a rough estimate on the number of rejected offers. However, past offer prices remain hidden, and "only a bad agent would reveal them," in the words of one broker. Similarly, in most labor markets, employers do not observe the actual offers that the applicant may have previously rejected, but they can infer how long he has been unemployed from the applicant's vita. In contrast, in other bargaining settings, such as corporate acquisition via tender offer, previous offers are commonly observed.

Remarkably, our analysis supports the broker's point of view. With public offers, the equilibrium outcome is unique. Bargaining typically ends up in an impasse: only the first buyer submits an offer that has any chance of being accepted. If this offer is rejected, no further serious offer is submitted on the equilibrium path. This is rather surprising, since it is common knowledge that, no matter how low the quality of the unit may be, it is still worth more to the potential buyers than to the seller. Why can a buyer not break the deadlock by making an out-of-equilibrium offer above the seller's lowest cost but below the buyer's lowest possible valuation? The problem is that the relevant benchmark for the seller is no longer his cost, but the offer he can expect in the following period upon rejecting an offer. Because offers are observable, rejecting an out-of-equilibrium offer that is accepted with some probability leads to a favorable updating by the next buyer, and therefore to a higher offer in the next period. As we show, gains from trade between the current buyer and the seller no longer exist once we account for this seller's outside option.[1] The problem is therefore one of commitment by the seller, who would gain by signing a pledge to accept any sufficiently high price.

This result provides an explanation for impasses in bargaining. While standard bargaining models are often able to explain delay, agreement is always reached eventually. Exceptions either rely on behavioral biases (see Babcock and Loewenstein, 1997) or Pareto-inefficient commitments (see Crawford, 1982). Here, it is precisely the inability of the seller not to solicit another offer

[1] Academic departments are well aware of this problem when considering making senior offers. As clearly this example cannot fail any of the underlying rationality assumptions, the prevalence of such offers raises an interesting puzzle.

that discourages potential buyers from submitting serious offers.

By contrast, agreement is *always* reached when offers are private. Because the seller cannot use his rejection of an unusually high offer as a signal to elicit an even higher offer by the following buyer, buyers are not deterred from submitting serious offers. To put it differently, the unique equilibrium outcome with public offers can no longer be an equilibrium outcome here. Suppose, *per impossibile*, that such an equilibrium were to exist. Then consider a deviation in which a potential buyer submitted an offer that is both higher than the seller's lowest possible cost yet lower than the buyer's lowest possible valuation. Future potential buyers would be unaware of the specific value of this out-of-equilibrium offer. Hence, turning it down would not change their beliefs about the unit's value. Thus, given that the seller expects to receive losing offers thereafter, he should accept the offer if his cost is low enough. This, in turn, means that the offer is a profitable deviation for the buyer.

Our main result may appear surprising in light of one of the 'linkage' principles in auction theory, stating that disclosure of additional information increases the seller's expected revenue. In our dynamic set-up, it is important to distinguish between how much information can possibly be revealed given the information structure and the information that is actually revealed in equilibrium. While finer information could be transmitted with public offers than with private offers, this is not what happens in equilibrium: because all offers but the first one are losing offers, no further information about the seller is ever revealed, so that, somewhat paradoxically, more information is communicated with private offers.

As already mentioned, this is a search model. However, unlike in most of the search literature, the distribution of offers is not fixed, but endogenously derived. The analysis shows that random offers can, indeed, be part of the equilibrium strategies. In addition, it shows that the offer distribution depends on the information available to the offerers. Therefore, it also suggests that it is not always innocuous to treat the offer distribution as fixed while considering variants of the standard search models.

The general set-up is described in Section 2. Section 3 solves the case of public offers, and addresses issues of robustness. Private offers are considered in Section 4. Related literature, results and extensions are discussed in Section 5. Proofs are in the appendix.

2 The Model

We consider a dynamic game between a single seller, with one unit for sale, and a countable infinity of potential buyers, or buyers for short. Time is discrete, and indexed by $n = 1, \ldots, \infty$. At each time or period n, one buyer makes an offer for the unit. Each buyer makes an offer only at one time, and we refer to buyer n as the buyer who makes an offer in period n, provided the seller has accepted no previous offer. After observing the offer, the buyer either accepts or rejects the offer. If the offer is accepted, the game ends. If the offer is rejected, a period elapses and it is another buyer's turn to submit an offer.

The quality q of the good is determined by Nature, and is uniformly distributed over the interval $\left[\underline{q}, 1\right]$, for some $\underline{q} < 1$. The value of q is the seller's private information, but its distribution is common knowledge. We refer to q as the seller's type. Given q, the seller's cost of providing the unit is $c(q)$. The valuation of the unit to buyers is common to all of them, and is denoted by $v\left(q\right)$.

We assume that c is (strictly) positive, (strictly) increasing and twice differentiable, with bounded derivatives. We also assume that v is positive, increasing and continuously differentiable. Moreover, we assume that v' is positive. We set $M_{c'} = \sup |c'|$, $M_{c''} = \sup |c''|$, $M_{v'} = \sup |v'|$, $M = \max \left\{M_{c'}, M_{c''}, M_{v'}\right\}$, and $m_{v'} = \min |v'| > 0$.

Observe that the assumption that q is uniformly distributed is made with little loss of generality, given that few restrictions are imposed on the functions v and c.[2]

We assume that gains from trade are always positive: $v\left(q\right) - c\left(q\right) > 0$ for all $q \in [\underline{q}, 1]$. The seller is impatient, with discount factor $\delta < 1$. We are particularly interested in the case in which δ is sufficiently large. To be specific, we set

$$\bar{\delta} := 1 - m_{v'}/3M,$$

and assume throughout that $\delta > \bar{\delta}$.

Buyer n submits an offer p_n that can take any real value. An outcome of the game is a triple (q, n, p_n), with the interpretation that the realized type is q, and that the seller accepts buyer n's offer of p_n (which implies that he rejected all previous offers). The case $n = \infty$ corresponds

[2]In particular, our results are still valid if the distribution of q has a bounded density, bounded away from zero.

to the outcome in which the seller rejects all offers (set p_∞ equal to zero). The seller's von Neumann-Morgenstern utility function over outcomes is his net surplus $\delta^{n-1}(p_n - c(q))$ when $n < \infty$, and zero otherwise. An alternative formulation that is equivalent to the one above is that the seller incurs no production cost but derives a per-period gross surplus, or reservation value, of $c(q)$ from owning the unit. It is immediate to verify that this interpretation yields the same utility function.

Buyer n's utility is $v(q) - p_n$ if the outcome is (q, n, p_n), and zero otherwise (discounting is irrelevant since buyers make only one offer). We define the players' expected utility over lotteries of outcomes, or *payoff* for short, in the standard fashion.

We consider both the case in which offers are public, and the case in which previous offers are private. It is worth pointing out that the results for the case in which offers are public also hold for any information structure (about previous offers) in which each buyer $n > 1$ observes the offer of buyer $n - 1$.

A history (of offers) $h^{n-1} \in H^{n-1}$ in case no agreement has been reached at time n is a sequence $(p_1, \ldots, p_{n-1})$ of offers that were submitted by the buyers and rejected by the seller (we set H_0 equal to $\varnothing$). A behavior strategy for the seller is a sequence $\{\sigma_S^n\}$, where σ_S^n is a probability transition from $[\underline{q}, 1] \times H^{n-1} \times \mathbb{R}$ into $\{0, 1\}$, mapping the realized type q, the history h^{n-1}, and buyer n's price p_n into a probability of acceptance. In the public case, a strategy for buyer n is a probability transition σ_B^n from H^{n-1} to $\mathbb{R}$.[3] In the private case, a strategy for buyer n is a probability distribution σ_B^n over $\mathbb{R}$.

Observe that, whether offers are public or private, the seller's optimal strategy must be of the cut-off type. That is, if $\sigma_S^n(q, h^{n-1}, p_n)$ assigns a positive probability to accepting for some q, then $\sigma_S^n(q', h^{n-1}, p_n)$ assigns probability 1 to it, for all $q' < q$. The proof of this skimming property is standard and can be found in, for example, Fudenberg and Tirole (Chapter 10, Lemma 10.1). The infimum over types q accepting a given offer is called the *indifferent type* (at history (h^{n-1}, p_n) given the strategy profile). Since the specification of the action of the seller's indifferent type does not affect payoffs, we also identify equilibria which only differ in this regard. For definiteness, in all formal statements, we shall follow the convention that a seller's type that is indifferent accepts the offer. For conciseness, we shall usually omit to specify that

[3]That is, for each $h^{n-1} \in H^{n-1}$, $\sigma_B^n(h^{n-1})$ is a probability distribution over $\mathbb{R}$, and the probability $\sigma_B^n(\cdot)[A]$ assigned to any Borel set $A \subset \mathbb{R}$ is a measurable function of h^{n-1}, and similarly for σ_S^n.

some statements only hold 'with probability 1'. For instance, we shall say that the seller accepts the offer when he does so with probability 1.

We use the perfect Bayesian equilibrium concept as defined in Fudenberg and Tirole (Definition 8.2).[4] In both the public and the private case, this implies that upon receiving an out-of-equilibrium offer, the continuation strategy of the seller is optimal.

In the public case, this also implies that, after any history on or off the equilibrium path along which all offers submitted by buyers were smaller than $c(1)$, the belief (over seller's types) of the remaining buyers is common to all of them and computed on the assumption that the seller's reasons for rejecting previous offers were rational. Thus, in the public case, after any such history, the belief of buyer n over q is the uniform distribution over some interval $[\underline{q}_n, 1]$, where $\underline{q}_n$ may depend on the sequence of earlier offers. In the private case, the only non-trivial information sets that are reached with probability 0 occur in periods such that, along the equilibrium path, the probability is 1 that the seller accepts some earlier offer. The specification of beliefs after such information sets turns out to be irrelevant.

Given some (perfect Bayesian) equilibrium, we follow standard terminology in calling a buyer's offer *serious* if it is accepted by the seller with positive probability. An offer is *losing* if it is not serious. Clearly, the specification of losing offers in an equilibrium is, to a large extent, arbitrary. Therefore, statements about uniqueness are understood to be made up to the specification of the losing offers. Finally, an offer is a *winning* offer if it is accepted with probability 1.

We briefly sketch here the static version with one buyer, similar to Wilson (1980). The unique buyer submits a take-it-or-leave-it offer. The game then ends whether the offer is accepted or rejected, with payoffs specified as before (with $n = 1$). Clearly, the seller accepts any offer p provided $p \geq c(q)$. Therefore, the buyer offers $c(q^*)$, where $q^* \in [\underline{q}, 1]$ is the maximizer of

$$\int_{\underline{q}}^{q} (v(t) - c(q))\, dt$$

with respect to q. Observe that the corresponding payoff of the buyer must be positive, because the buyer can always submit an offer in the interval $\left(c\left(\underline{q}\right), v\left(\underline{q}\right)\right)$.

[4] Formally speaking, Fudenberg and Tirole define perfect Bayesian equilibria for finite games of incomplete information only. The suitable generalization of their definition to infinite games is straightforward and omitted.

3 Observable Offers

Throughout this section, we maintain the assumption that offers are public. We prove that the market breaks down: only limited trade takes place in the first stage, and no trade ever takes place beyond that stage.

3.1 Example and Intuition

To shed some intuition on the absence of trade, we first provide a sketch of the argument for a simple parametric example. We let the cost and the valuation functions be given by $c(q) = q$ and $v(q) = (1+\alpha)q$, where $\alpha \in (0,1)$.[5] To ensure the existence of gains from trade, assume here that $\underline{q} > 0$.

Fix some (perfect Bayesian) equilibrium. Given the optimal acceptance rule of the seller, the belief of the current buyer on the seller's type after any history is a uniform distribution over an interval of the form $[q, 1]$. To simplify the exposition, we restrict ourselves here to stationary (or symmetric) pure-strategy equilibria in Markov strategies. (Theorem 1 is proved in appendix without any such restriction.) To be specific, we assume for now that, after any history, the current buyer submits a single offer with probability one that only depends on (what he thinks is) the lowest remaining type q. Such equilibria are described by a function $x : [\underline{q}, 1] \to [\underline{q}, 1]$, where $x(q) \geq q$ stands for the highest type that accepts the current buyer's offer, as a function of the current lowest type q. Accordingly, along the equilibrium path, the lowest remaining type in stage n is given recursively by $\underline{q}_n = x(\underline{q}_{n-1})$.[6]

We start with two simple observations. First, note that no equilibrium offer exceeds $c(1) = 1$. Indeed, following any history, any offer above $v(1) = 1 + \alpha$ is strictly dominated, hence no such offer is ever made. Therefore, following any history, any offer above $1+\delta\alpha$ is accepted by all types

[5] If $\alpha \geq 1$, following the lines of the analysis below, the first buyer offers 1 in any equilibrium, and the seller accepts, irrespective of his type.

[6] Beware that the current lowest type q is not a state variable in the usual sense, since it depends on past offers, and on expectations relative to all future offers as well. As a result, there is some circularity involved here, since the offer of buyer n is taken to depend on the lowest type $\underline{q}_n$, which is computed on the basis of buyer n's strategy, among other things.

since the discounted benefit from turning it down is at most $\delta\alpha$. This implies in turn that any offer which exceeds $1+\delta^2\alpha$ is accepted by all types, etc. Since a winning offer must be at least $c(1)$, it follows that $c(1)$ is the unique winning offer that may be specified in any equilibrium.

Second, a seller with type x accepts an offer p if and only if the benefit $p-x$ derived from accepting the current offer is at least equal to the discounted benefit derived from accepting some later offer. We denote by $p(x)$ the value of p at which type x is indifferent between accepting and declining. This price only depends on x, and not on the currently lowest type, q. Indeed, the offer $p(x)$ will be accepted by all types up to x, so that x will become the lowest remaining type at the next stage. Under the assumption that the equilibrium profile is symmetric and Markov, the highest benefit derived from future offers only depends on x. Plainly, $p(x) \geq c(x) = x$ and $p(\cdot)$ is non-decreasing. The offer $p(x)$ is accepted with probability $\dfrac{x-q}{1-q}$, and the current buyer's payoff, conditional on this offer being accepted, is $(1+\alpha)\dfrac{x+q}{2} - p(x)$. Thus, $x(q)$ solves the maximization problem

$$\max_x \; (x-q)\left((1+\alpha)\,\frac{x+q}{2} - p(x)\right). \tag{1}$$

We shall solve for the game backwards, starting with high values of q, and moving then to lower values of q. In doing go, we shall see that every 'subgame' defined by such a q admits a unique equilibrium (for almost all values of q).

We first claim that the winning offer is the unique optimal offer provided that the range of remaining types is small enough. Since $p(x) \geq c(x) = x$, the maximand in (1) does not exceed $(x-q)\left((1+\alpha)\,\dfrac{x+q}{2} - x\right)$, which is the payoff that a monopsonistic buyer would obtain by offering x in the static model. Note that the latter payoff is strictly increasing over $x \in [q,1]$ whenever $q \geq 1-\alpha$. In addition, since $p(1) = c(1) = 1$, the two payoffs coincide for $x = 1$. Therefore, the maximization problem in (1) admits as unique solution $x = 1$ for $q \geq 1-\alpha$. Thus, $x(q) = 1$ whenever $q \geq 1-\alpha$. Set $q^* := \sup\{q : x(q) < 1\} \leq 1-\alpha$. Plainly, a winning offer yields a nonnegative payoff whenever it is optimal. Hence, whenever $q > q^*$, the winning offer 1 does not exceed the average valuation of the remaining types, $(1+\alpha)\,\dfrac{1+q}{2}$. Equivalently, $q^* \geq \beta$, where $\beta := \dfrac{1-\alpha}{1+\alpha} > 0$.

Our impasse result rests on three main insights. The first is that a gradual erosion of uncertainty ending with full trade cannot be an equilibrium outcome. Specifically, we claim that there is no value of q for which $x(q) \in (q^*, 1)$. Accordingly, all types in $(q^*, 1]$ behave in the same way on the equilibrium path. To see this, fix $q \leq q^*$. If the current buyer submits an offer $p(x)$ with $x > q^*$, then a winning offer will immediately follow, so that $p(x)$ solves $p(x) - x = \delta(1 - x)$, or

$$p(x) = \delta + (1 - \delta)x.$$

If δ is close to one, this offer $p(x)$ is hardly sensitive to x, and, more importantly, less sensitive to x than the expected valuation, $(1 + \alpha)\frac{x + q}{2}$. Thus, if the offer $p(x)$ yields a nonnegative payoff to the current bidder, any higher offer $p(y)$ $(y > x)$, yields an even higher payoff as both (i) such an offer is accepted with higher probability and (ii) the expected payoff conditional on acceptance increases as well. Formally, the expression $(x - q)\left((1 + \alpha)\frac{x + q}{2} - p(x)\right)$ is strictly convex in $x \in [q^*, 1]$ when $p(x) = \delta + (1 - \delta)x$.

This enables us to determine q^*. By definition, there are values of $q < q^*$ arbitrarily close to q^* for which $x(q) < 1$, and so for which $x(q) \leq q^*$, given that $x(q) \notin (q^*, 1)$. The corresponding optimal offer must yield a payoff arbitrarily close to zero, since it is accepted with an arbitrary small probability. Hence, it cannot be the case that a winning bid yields a positive payoff when $q = q^*$: that is, $1 \geq (1 + \alpha)\frac{1 + q^*}{2}$ or equivalently, $q^* \leq \beta$. Because offering 1 is unprofitable for all values of $q < \beta$, it follows that $x(q) \leq q^*$ for all $q < q^*$. Remember also that $q^* \geq \beta$, so that $q^* = \beta$.

To summarize our findings so far, one has $x(q) = 1$ if $q > \beta$, $x(q) \leq \beta$ if $q < \beta$, and $x(\beta) \in \{\beta, 1\}$. In particular, if $\underline{q} > \beta$, the first buyer submits a winning offer of 1, which is accepted for sure by the seller. Instead, we assume from now on that $\underline{q} < \beta$.[7]

If $x(\beta) = \beta$, our findings imply that $\underline{q}_n \leq \beta$ for all n whenever $\underline{q} < \beta$: no quality above β is ever traded, and not all gains from trade can be realized. It turns out, and this is the second main insight of the proof, that the same impasse arises if $x(\beta) = 1$. To see this, we will show

[7] In the knife-edge case $\underline{q} = \beta$, the equilibrium value of $x(q) \in \{\beta, 1\}$ is indeterminate.

that $p(\cdot)$ is discontinuous at β if $x(\beta) = 1$:

$$\lim_{x \nearrow \beta} p(x) < p(\beta). \tag{2}$$

Intuitively, the cost of making an offer accepted by all types up to β is bounded above the cost of making an offer accepted by all types up to $\beta - \varepsilon$, for all $\varepsilon > 0$, because a winning offer immediately follows in the first case, while the next offer is accepted by all types up to (at most) β in the second case.

Before proving (2), observe that it implies that the offer $p(\beta)$ is worse than the offer $p(\beta - \varepsilon)$, for $\varepsilon > 0$ small enough, since the probability of trade and the expected valuation conditional on trade are arbitrarily close for the two offers, yet the offer prices are not. Therefore, $x(\beta) = 1$ would imply that $x(q) < \beta$ whenever $q < \beta$, so that $\underline{q}_n < \beta$ for every n on the equilibrium path.

Why does (2) hold? On the one hand, under the assumption that $x(\beta) = 1$, an offer of $p(\beta)$ would immediately be followed by a winning offer 1, hence

$$p(\beta) - \beta = \delta(1 - \beta). \tag{3}$$

On the other hand, any offer $p(x)$ with $x < \beta$ would be followed by an offer that does not exceed $p(\beta)$ which, if rejected, would itself be followed by offers that do not exceed 1. Thus, the discounted payoff to type x when declining $p(x)$ is at most $\max\left\{\delta\left(p(\beta) - x\right), \delta^2\left(1 - x\right)\right\}$, hence

$$p(x) - x \leq \max\left\{\delta\left(p(\beta) - x\right), \delta^2\left(1 - x\right)\right\}.$$

Equation (3) implies that the maximum is achieved by the first term, so that, letting x tend to β, $\lim_{x \nearrow \beta} p(x) - \beta \leq \delta\left(p(\beta) - \beta\right)$. Using (3) once more implies (2).

This already implies our result on the impasse, since whether $x(\beta) = \beta$ or 1, $\underline{q}_n \leq \beta$ holds for all n in either case. To complete the equilibrium description, we now argue that the existence of an equilibrium is actually inconsistent with the specification $x(\beta) = 1$. This claim follows by adapting one of our earlier arguments. For $q < \beta$, and since $p(x) \geq x$, the payoff

$$\gamma(x) := \frac{x - q}{1 - q}\left((1 + \alpha)\frac{x + q}{2} - p(x)\right)$$

derived from an offer of $p(x)$ does not exceed the payoff $\frac{x - q}{1 - q}\left((1 + \alpha)\frac{x + q}{2} - x\right)$ that a monopsonistic buyer would obtain when offering x in the static case. In addition, and since $x(q) < \beta$

for all $q < \beta$, one can check that $\lim_{x \nearrow \beta} p(x) = c(\beta) = \beta$, so that the difference between the two expressions vanishes as $x \nearrow \beta$. For $q \geq \beta(1-\alpha)$, the former payoff is increasing over the interval $[q, \beta]$. Hence, assuming for now that a perfect Bayesian equilibrium exists, the function $\gamma(\cdot)$ has a unique maximum for $x = \beta$. This implies that $x(\beta)$ must be equal to β, for otherwise $\gamma(\cdot)$ would be discontinuous at β, with $\lim_{x \nearrow \beta} \gamma(\cdot) > \gamma(\beta)$.

The above equilibrium picture is only partial, as it is yet unclear whether types below β are traded at all. The last insight in the proof is that this picture repeats itself on the interval $[\beta^2, \beta]$, and then inductively on each interval $[\beta^{n+1}, \beta^n]$. Thus, assuming $\underline{q} \in (\beta^{n+1}, \beta^n)$ for some n, there is a unique candidate for a pure, stationary equilibrium in Markov strategies. It is characterized by the following feature: the current buyer offers β^k whenever the lowest remaining type q belongs to the interval $(\beta^{k+1}, \beta^k]$, which the seller accepts if and only if his type does not exceed β^k. It is immediate to verify that this specification constitutes an equilibrium.

As a result, only the first buyer submits a serious offer on the equilibrium path. In terms of comparative statics, observe that a decrease in the severity of the information asymmetry -that is, a slight increase in $\underline{q} < \beta^n$- actually leads to a decrease in the probability of trade.

3.2 General Statement and Related Literature

We build upon the intuition provided above, and state our main result in the public case. We start by defining a finite sequence (q_k), which generalizes the sequence (β^k) of the previous section.

Given q, and since v is increasing, there exists at most one value of $x \in \left[\underline{q}, q\right)$ solving

$$c(q) = \frac{1}{q-x} \int_x^q v(t)\, dt.$$

That is, for this value of x, the expected valuation over types in the interval $[x, q]$ is equal to the highest cost in this interval, $c(q)$. We then denote by $f(q) := x \in \left[\underline{q}, q\right)$ this value, as a function of q, if it exists. The value $f(q)$ is well-defined whenever asymmetric information is severe enough that the buyer's expected value does not exceed the highest seller's cost. This is the case if either the range of existing qualities is large, or valuations do not exceed costs by much.

The mapping f plays a key role in the analysis of the static model of adverse selection between one seller and two or more buyers. Indeed, consider the 'static' game between $n \geq 2$ buyers simultaneously announcing prices and the seller who then either keeps the unit or sells it at the highest price p. Obviously, the seller only sells the unit if $p \geq c(q)$, or $q \leq c^{-1}(p)$. Therefore, it must be that $p' \geq p$ implies that $p' \geq \mathbb{E}\left[v(q)\,|q \leq c^{-1}(p')\right]$ (where $c^{-1}(p) = 1$ for $p \geq c(1)$) with equality for $p' = p$, so that the winner barely breaks even and no higher price yields a positive payoff. That is, $\underline{q} = f\left(c^{-1}(p)\right)$ (unless a winning offer yields a positive payoff).[8] Note that any such offer p strictly exceeds the optimal offer that a lone buyer would submit.

Define the strictly decreasing sequence q_k as follows: $q_0 = 1$, $q_{k+1} = f(q_k)$, as long as $f(q_k)$ is defined (that is, as long as $f(q_k) \geq \underline{q}$). Because $\min_q \{v(q) - c(q)\} > 0$, it must be that, for all q, $q - f(q) > \kappa$ for some $\kappa > 0$. Hence this sequence must be finite, and we denote the last and smallest element of this sequence by $q_K \geq \underline{q}$. Note that this sequence is always well-defined, as $q_0 = 1 > \underline{q}$.

While it is possible that $q_K = \underline{q}$, this proposition is stated here for simplicity for the generic case in which $q_K > \underline{q}$.

Theorem 1 *Assume that $q_K > \underline{q}$, and $\delta > \bar{\delta}$. There is a unique equilibrium outcome, which is independent of δ. On the equilibrium path, the first buyer submits the offer $c(q_K)$, which the seller accepts if and only if $q \leq q_K$. If this offer is rejected, all buyers $n > 1$ submit a losing offer.*

Although the equilibrium outcome is unique, the equilibrium is not. However, there is a unique stationary equilibrium in Markov strategies achieving this outcome. Given any history leading to beliefs that are (necessarily) uniform over some interval $[q, 1]$, the buyer submits the offer $c(q_k)$, where $q_k = \min_l \{q_l : q_l \geq q\}$, and the seller's behavior is given by the obvious best-reply. More generally, all equilibria share the following property: whenever the currently lowest type q is such that $q \in (q_{k+1}, q_k)$ for some $k \leq K$, the current buyer offers $c(q_k)$, which the seller accepts if and only his type does not exceed q_k, and all subsequent offers are losing ones.

[8] More generally, in all equilibria, there are at least two buyers whose offers all have the property that higher offers cannot yield a positive payoff. Because f^{-1} need not be uniquely defined, the equilibrium need not be unique, pure or symmetric.

As mentioned, other equilibria achieving the same outcome exist, but they are rather trivial modifications of this stationary equilibrium.[9]

The proof is in appendix. It is a detailed elaboration of the sketch given in Section 3.1.

It is now possible to draw a comparison between the dynamic version with public offers and the static version with one buyer. Observe that, depending on the exact value of $\underline{q}$, q^* could be anywhere in the interval $\left(\underline{q}, q_{K-1}\right)$, so that both $q_K > q^*$ and $q_K \leq q^*$ may occur, where $c(q^*)$ is the optimal offer in the static version.[10] This means that, from the seller's point of view, the comparison between the dynamic version and the static version with a unique buyer is ambiguous. The probability of sale and the expected revenue could be larger in either format depending on $\underline{q}$. However, it makes more sense to compare this equilibrium outcome with the equilibria in the static game with multiple buyers. The comparison is then immediate, as the offer in the static case must be at least as large as $c\left(q_K\right)$. Thus, the seller is better off in the static version, having the different buyers compete simultaneously for the unit, rather than one at a time. The probability of trade is higher in the static version. Only the first buyer is better off in the dynamic version, while all other buyers are indifferent. As an immediate consequence, the bargaining outcome generically fails to be *ex ante* efficient, that is there exists an incentive-compatible and individually rational mechanism that yields higher expected gains from trade. Indeed, with a single buyer, consider the mechanism in which the seller must accept or reject the fixed price $c\left(q\right)$, where q is the largest root of $f\left(q\right) = \underline{q}$.

The result that trade does not necessarily occur is surprising. In a sense, there is infinite delay as soon as the lowest type reaches q_k. The most recent contribution to this literature, Deneckere and Liang (2006), considers the case of a single long-run buyer, with the same discount factor as the seller, rather than a sequence of short-run buyers. They define a finite sequence (q_k) similar

[9]For completeness, consider here the knife-edge case in which $\underline{q} = q_K$. If the lowest type is q_K, any randomization over the offers $\{q_K, q_{K-1}\}$ is optimal, the payoff of either offer being zero. Because $\underline{q} = q_K$, equilibrium considerations do not uniquely 'pin down' the mixture, as in the case $\underline{q} < q_K$. Indeed, the only reason why the equilibrium (as opposed to the equilibrium outcome) in the case $\underline{q} < q_K$ is not unique is that the offer when the indifferent type is q_k, $k \leq K$, is indeterminate, *following an out-of-equilibrium offer*. The case $\underline{q} = q_K$ is otherwise identical to the case $\underline{q} < q_K$. Along the equilibrium path, the seller rejects all offers provided $q \geq q_{K-1}$.

[10]In the linear case, $\underline{q} = (1-\alpha)q^*$ while the only constraint on $\underline{q}$ is $\dfrac{1-\alpha}{1+\alpha}q_K < \underline{q} \leq q_K$. For low (resp. high) values of $\underline{q}$ in this interval, the dynamic (resp. static) version yields a higher payoff to the seller.

(but not identical) to ours, and find that, in the unique stationary equilibrium, as the lowest type approaches q_k from below, the (single) buyer repeatedly submits offers accepted with small probability by the seller, so that delay ensues. As the time period between successive offers vanishes, this delay remains finite and bounded away from zero. In the limit, bursts of trade alternate with long periods of delay. This implies that the equilibrium acceptance function of the seller becomes increasingly similar to the equilibrium acceptance of the seller in Theorem 1. Note however that, because delay remains finite in their model, the price accepted by type q_k exceeds $c\left(q_k\right)$, so that the actual values of the two sequences differ. Furthermore, this limit comparison is only illustrative, since the conclusion of Theorem 1 is valid for every $\delta > \bar{\delta}$.

Theorem 1 remains valid in the case of a single long-run buyer, provided the buyer is much less patient than the seller. Hence, when combined, these two results point out that the possibility of trade depends on the relative patience of the buyer relative to the seller, an insight already hinted at by Evans (1989) in the case of binary values.

Just as Theorem 1 remains valid with a single buyer who is sufficiently more impatient than the seller, it is also valid if the number of buyers is finite, as long as the probability that each of them is selected to make the offer in each of countably many periods is sufficiently small.

The results of Vincent (1989) and Deneckere and Liang (2006) rely on the screening of types that bargaining over time affords. Because delay is costly for the seller, buyers become more optimistic over time, so that uncertainty is progressively eroded. Our no-trade result points to another familiar force in dynamic games; namely, the absence of commitment. Indeed, if the horizon were finite, the last buyer would necessarily submit a serious offer. However, since Coase's (1972) original insight, the inability to commit has always been associated with an *increase* in the probability of trade. To quote Deneckere and Liang (p. 1313), the "absence of commitment power implies that bargaining agreement will eventually be reached". This is because the traditional point of view emphasized the inability of the buyer to commit to not making another offer. Instead, the driving force here is the inability of the seller not to solicit another offer. This leads to a fall in the probability of trade, and an increase in the inefficiency.

3.3 Robustness[11]

As described, the (essentially) unique equilibrium of the game with observable offers calls for all buyers but the first one to submit losing offers, while they would have 'nothing to lose' by submitting instead offers in the interval $(c(q_K), v(q_K))$. This raises then the issue of robustness of Theorem 1. As these buyers are deterred from submitting moderate offers above the seller's lowest possible cost because they correctly anticipate such offers to be turned down, it is legitimate to ask what happens when buyers expect the seller to 'tremble'. That is, we assume now that, in every period, and independently across periods, with some small probability $\varepsilon > 0$, the seller accepts any (positive) offer, independently of his type.[12] With complementary probability, the seller behaves strategically, as before. One may interpret the probability ε as the probability of a liquidity shock, for instance, and we correspondingly refer to the seller as a liquidity seller – as opposed to a strategic seller – in this event.[13] In a given period, the events unfold as follows. First, the buyer makes his offer. Second, a liquidity shock occurs or not. If it does occur, the seller then accepts the offer. If it does not, the seller then accepts or rejects, depending on his type and his expected continuation payoff. Although the strategic seller cannot anticipate when a liquidity shock will occur in the future, he takes this possibility into account while evaluating the payoff from turning down a given offer. Observe that the standard model studied in the previous sections corresponds to the special case $\varepsilon = 0$. Since we are motivated by the question of robustness, we are especially interested in the limit $\varepsilon \to 0$. For expositional convenience, we restrict attention to the linear example studied in Section 3.1.

Obviously, the payoff of every buyer is now positive, as even an offer of 0 is accepted with probability $\varepsilon > 0$ by the seller. Also, the unit will be eventually sold, as the seller is bound to become a liquidity seller sooner or later. The interesting question, then, is whether an im-

[11] We thank two referees for raising the question addressed in this subsection.

[12] We could alternatively consider the case in which, with some small probability, the seller accepts an offer if and only if it exceeds the seller's cost. In this case, posterior beliefs are no longer uniform (unless the exogenous shock is itself observable).

[13] Note that the offer might fall short of the cost, but this is not inconsistent with the interpretation of a liquidity shock, if this cost is either non-monetary (for instance, an effort cost), or if the asset is not as liquid as the offer. Also, the lower bound on offers has been set to zero for notational convenience, but the result remains valid for any other lower bound not exceeding $\underline{q}$.

passe persists as long as the seller is strategic. The following proposition, proved in appendix, establishes that this is indeed the case.

Proposition 1 *For ε close to zero, and for δ close to one, there exists two decreasing sequences $(\beta_k)_k$, $(\gamma_k)_k$ in $[\underline{q}, 1]$, with $\gamma_0 = \beta_0 = 1$, and $\gamma_k \in \left(\beta_{k+1}, \beta_k\right)$ (whenever defined), such that the following holds.*

Denote by β_K and γ_L the last elements in the two sequences, so that $K - 1 \leq L \leq K$, and assume that $K \geq 1$.[14]

If $\underline{q} \neq \beta_K$ and $\underline{q} \neq \gamma_L$, there exists a unique equilibrium outcome, which is as follows:

- *For $\underline{q} \in (\beta_K, \gamma_{K-1})$, the first buyer offers $\eta_\varepsilon q'$, for some $q' \in [\gamma_K, \beta_K]$ and $\eta_\varepsilon < 1$, which the strategic seller accepts if and only if $q \leq q'$. If this offer is rejected, all buyers $n > 1$ submit the offer 0.*
- *For $\underline{q} \in (\gamma_K, \beta_K)$, all buyers submit the offer 0.*

Further, $\lim_{\varepsilon \to 0} \gamma_k = \lim_{\varepsilon \to 0} \beta_k = \beta^k$ and $\lim_{\varepsilon \to 0} \eta_\varepsilon = 1$, hence the equilibrium outcome converges to the equilibrium outcome of the standard model as $\varepsilon \to 0$.

The definitions of the sequences γ_k, β_k, are provided in the proof. These sequences are obtained independently of $\underline{q}$, hence the case $\gamma_K \neq \underline{q}$ is generic.

It is necessary to introduce two further sequences in order to describe the buyers' strategies somewhat more precisely. There exists two sequences λ_k, μ_k, with $\beta_{k+1} < \lambda_k < \mu_k < \gamma_k$, and $\lim_{\varepsilon \to 0} \lambda_k = \lim_{\varepsilon \to 0} \mu_k = \beta^k(1 - \alpha)$, such that, as a function of the lowest remaining type q, the current buyer offers (i) $\eta_\varepsilon \gamma_k$ if $q \in \left(\beta_{k+1}, \lambda_k\right)$; (ii) $\eta_\varepsilon q'$, where $q' \in (\gamma_k, \beta_k)$ is increasing in q, if $q \in (\lambda_k, \mu_k)$; (iii) $\eta_\varepsilon \beta_k$ if $q \in (\mu_k, \gamma_k)$; (iv) 0 if $q \in (\gamma_k, \beta_k)$. The case $k = 0$ is quite special: if q is in $(\beta_1, 1)$, the buyer makes the winning offer 1. Continuity as ε tends to 0 holds here as well. The behavior of the buyer is summarized in Figure 1.

[14]The description of the equilibrium for $K = 0$ is omitted here, but it can be immediately deduced from the proof of Proposition 1.

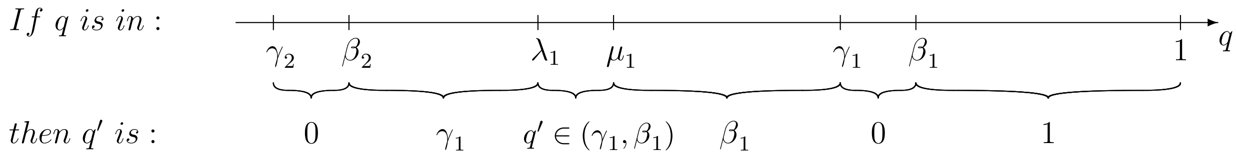

Figure 1 : Indifferent type q' of buyer $n+1$ given indifferent type q of buyer n.

The intuition for this result is as in the standard model. When a buyer is sufficiently optimistic, it is optimal for him to buy with probability one, as in the monopoly case. When the seller is patient enough, competition then kicks in as we consider lower values of $\underline{q}_n$, and the average quality for which it is optimal to buy for sure unravels, up to the point at which a buyer is indifferent between a winning offer and a 'losing' offer that is only accepted in case of a liquidity shock. For slightly lower average qualities, such a losing offer remains optimal: a serious offer that would be likely to be accepted by the strategic type would be prohibitively costly, since if such an offer were rejected, the average quality would sufficiently improve to trigger a winning offer in the following period; and a serious offer that has low probability of being accepted by the strategic seller is still bounded away above the losing offer. For a range of still lower average qualities, it is then optimal to submit an offer that drives up the average quality to a level that triggers losing offers, which ensures that such an offer is not too expensive.

4 Private Offers

4.1 General Properties

Throughout this section, we maintain the assumption that offers are private. As we are unable to construct equilibria in general, we first argue that an equilibrium exists. We will apply a fixed-point argument on the space of buyers' strategies. Given a profile of buyers' strategies σ_B, the buyers' payoffs are computed using the optimal response of the seller to the profile σ_B.

If no later buyer sets a price exceeding $c(1)$, it is suboptimal for a given buyer to set such a price. Hence, for the purpose of equilibrium existence, we can limit the set of buyers' mixed (or behavior) strategies to the set $\mathcal{M}([0, c(1)])$ of probability distributions over the interval $[0, c(1)]$, endowed with the weak-$\star$ topology. The set of strategy profiles for buyers is thus the countable product $\mathcal{M}([0, c(1)])^{\mathbb{N}}$. It is compact and metric when endowed with the product topology. Since the random outcome of buyer n's choice is not known to the seller unless he has rejected the first $n-1$ offers, buyer n's payoff function is not the usual multi-linear extension of the payoff induced by pure profiles. Given a period n, we denote by $q_n(p, \sigma_B)$ the indifferent type when the offer p is submitted in stage n, given the strategy profile σ_B. It is jointly continuous in p and σ_B. As a result, the belief of any given buyer (viewed as a probability distribution over $[c(\underline{q}), c(1)]$) is continuous with respect to σ_B, in the weak-$\star$ topology. Hence, the set $B_n(\sigma_B) \subset \mathcal{M}([0, c(1)])$ of best replies of buyer n to σ_B is both convex-valued and upper hemi-continuous with respect to σ_B. The existence of a (Nash) equilibrium follows from Glicksberg's fixed-point theorem. To any such equilibrium, there corresponds a perfect Bayesian equilibrium, because all buyers' histories are on the equilibrium path until trade occurs, if ever.

While an equilibrium exists, it need not be unique. As an example, for $c(q) = q$, $v(q) = (1+\alpha)q$, $\alpha = 1/3$, $\delta = 3/4$ and $\underline{q} = .4249$, it can be shown that the following two (and probably more) equilibria exist.

- The first buyer offers $p_1 \simeq .61$ and attracts all types up to $\beta = 1/2$ for sure. The second buyer makes a losing offer. Buyer $n \geq 3$ mix between a winning and a losing offer, offering the winning price 1 with probability 3/20.

- The first buyer randomizes between a losing offer and the offer $p_1 \simeq .63$ with indifferent type $q_1 \simeq .51$, assigning a probability $\mu_1 \simeq .40$ to the losing offer. The second buyer mixes between a losing offer and the offer $p_2 \simeq .72$ with indifferent type $q_2 \simeq .62$, assigning a probability $\mu_2 \simeq .76$ to the losing offer. Buyer $n \geq 3$ mix between a winning and a losing offer, offering the losing price with probability $\mu \simeq .94$.

We summarize the discussion so far in the following lemma.

Lemma 1 *An equilibrium exists. For some parameters, the equilibrium is not unique.*

More precisely, the equilibrium is unique if and only if $\underline{q} > q_1$, as defined in Section 3. The two equilibria described above differ quantitatively in terms of delay, revenue, and payoffs. Nevertheless, there are some qualitative similarities. Both equilibria involve mixed strategies. Also, trade occurs with probability 1. This is no coincidence.

Proposition 2 *In all equilibria, trade occurs with probability 1.*

Proof: Fix some equilibrium. Given $q \in \left[\underline{q}, 1\right]$, let $F_n(q)$ denote the unconditional probability that the seller is of type $t \leq q$ and has rejected all offers submitted by buyers $i = 1, \ldots, n-1$, and denote by F the pointwise limit of the sequence (F_n). Suppose that $F(q) \neq 0$ for some $q < 1$. In particular, the probability that the seller accepts buyer n's offer, conditional on having rejected the previous ones, converges to zero as n increases. Hence, the successive buyers' payoffs also converge to zero. Choose q such that $F(q) > 0$ and

$$\int_{\underline{q}}^{q} \left(v(t) - c(q) - \frac{\nu}{2}\right) dF(t) > 0, \tag{4}$$

where $\nu := \min_{x \in [\underline{q},1]} \{v(x) - c(x)\}$. Note that $\dfrac{F(q) - F_n(q)}{F_n(q)}$ is the probability that type q accepts an offer from a buyer after n, conditional on having rejected all previous offers. Since $F(q) > 0$, this probability converges to zero and the offer $p_n(t)$ with indifferent type t tends to no more than $c(q)$. Thus, $p_n(t) \leq c(q) + \frac{\nu}{2}$ for n large and, using (4), buyer n's payoff is bounded away from zero, a contradiction. ■

Observe that Proposition 2 holds independently of δ and establishes that offers arbitrarily close to 1 are eventually submitted.

According to Proposition 2, agreement is always reached in finite time. This raises the question of delay. That is, let $\tau(1)$ denote the random period in which a winning offer is first submitted. The next proposition places bounds on $\mathbb{E}\left[\delta^{\tau(1)}\right]$, the expected delay until agreement. In particular, it implies that $\tau(1)$ is finite a.s.: a winning offer is submitted in finite time, with probability 1.

Proposition 3 *Assume that $\delta > \bar{\delta}$. There exists constants $0 < c_1 < c_2 < 1$ such that, in all equilibria,*

$$c_1 \leq \mathbb{E}\left[\delta^{\tau(1)}\right] \leq c_2.$$

The proof of this and all remaining results can be found in appendix. Delay ($c_2 < 1$) should not come as a surprise. Since the seller can wait until the first winning offer is submitted, and serious offers until then must yield a nonnegative payoff to the buyers submitting them, delay must make waiting for the winning offer a costly alternative to the seller's lower types. Slightly less obvious is the second conclusion; namely, that delay does not dissipate all gains from trade ($c_1 > 0$).

In the first example of this section, the first buyer enjoys a positive payoff, but all other buyers get a zero payoff. More complicated examples of equilibria can be constructed in which more than one buyer gets a positive payoff. However, in all equilibria, all buyers' payoffs are small.

Proposition 4 *There exists a constant $M_1 > 0$ such that, for every $\delta > \bar{\delta}$ and every equilibrium, the payoff of any buyer n is at most*

$$(1-\delta)^2 M_1.$$

According to the next proposition, buyers with a positive equilibrium payoff are infrequent, in the sense that two such buyers must be sufficiently far apart in the sequence of buyers.

Proposition 5 *There exists a constant $M_2 > 0$ such that, for every $\delta \geq \bar{\delta}$ and every equilibrium, the following holds. If buyer n_1 and buyer n_2 both get a positive payoff, then*

$$|n_2 - n_1| \geq \frac{M_2}{1-\delta}.$$

4.2 Equilibrium Strategies

As mentioned, we do not provide an explicit characterization of an equilibrium for general parameters. Nevertheless, all equilibria share common features. Given some equilibrium, let $F_n(q)$ denote the (unconditional) probability that the seller's type t is less than or equal to q and that all offers submitted by buyers $i = 1, \ldots, n-1$ are rejected. Set $\underline{q}_n := \inf\{q : F_n(q) > 0\}$. Buyer n's strategy is a probability distribution over offers in $[c(\underline{q}), c(1)]$. We denote by P_n its support and by T_n the corresponding (closed) set of indifferent types. That is, if buyer n's strategy has finite support, $q \in T_n$ if it is an equilibrium action for buyer n to submit some p_n that seller's type t accepts if and only if $t \leq q$.

The following proposition complements Proposition 5, as together they imply that the number of buyers with a positive payoff is bounded above, uniformly in the discount factor and the equilibrium.

Proposition 6 *Assume that $\delta > \bar{\delta}$. Given some equilibrium, let $N_0 := \inf\{n \in \mathbb{N} \cup \{\infty\} : 1 \in T_n\}$. There exists a constant $M_3 > 0$ such that, in all equilibria, $N_0 \leq M_3/(1-\delta)$. Further, given some equilibrium, $T_n \subset \{\underline{q}_{N_0}, 1\}$ for all $n \geq N_0$. For all $n > N_0$, buyer n's equilibrium payoff is zero.*

Observe that, by Propositions 4, 5 and 6, the (undiscounted) sum of *all* buyers' payoffs is at most $(1-\delta)^2 M_3 M_1/M_2$, and therefore vanishes as the seller becomes more patient.

Thus, from period N_0 on, buyers only make winning or losing offers, and all but the first of these have a payoff of zero. In fact, it follows readily from the proof that the equilibrium payoff of buyer N_0 is zero as well, as long as $\underline{q} \leq q_1$.

If $\underline{q} > q_1$, it follows from Proposition 6 that in the unique equilibrium outcome, the first buyer offers $c(1)$, which the seller accepts. Indeed, provided that he is called upon to submit an offer, any buyer is guaranteed a positive payoff, because he can always offer $c(1)$.

In all other cases, there exist multiple equilibria. In particular, there always exists an equilibrium in which agreement is reached in bounded time, as well as an equilibrium in which agreement is reached in unbounded, yet finite, time. Indeed, given any equilibrium profile σ, the probabilities assigned by buyers $n > N_0$ to the winning offer can be modified so that the modified profile is still an equilibrium, with the desired property. While the equilibria obtained in this way are payoff-equivalent, the first example in this section shows that this is not true across all equilibria.

For $\underline{q} < q_1$, the next proposition formalizes the idea that all equilibria are in mixed strategies.

Proposition 7 *Assume that $\delta > \bar{\delta}$ and $\underline{q} < q_1$. No buyer $n \leq N_0$ uses a pure strategy, except possibly buyer 1. All buyers $n \leq N_0$ submit a serious offer with positive probability.*

Indeed, buyer 1 need not use a mixed strategy, as the first example given in this section illustrates. Without further assumptions, it is difficult to establish additional structural properties on equilibrium strategies. However, it can be shown that, if v is concave and c is convex over

$(\underline{q}, 1)$, with either v or c being strictly so, then each buyer's strategy is a distribution with *finite* support, so that each buyer randomizes over finitely many offers only.[15]

Propositions 6 and 7 allow us to circumscribe the equilibrium strategies as follows. During a first phase of the game (until period $N_0 - 1$), buyers' strategies assign positive probability to more than one offer (with the possible exception of the first buyer's strategy); in particular, they all assign positive probability to serious, but not winning, offers. Some of these buyers may enjoy a small positive payoff, while all others have zero payoff; in fact, the number of those not submitting a losing offer with positive probability is finite as well. In a second phase (from period N_0 on), all buyers' payoffs are zero, and they randomize between the winning offer and a losing offer, with relative probabilities that are to a large extent free variables. Thus, as long as offers are rejected, the unit's expected value increases until N_0 and is constant thereafter.

It is tempting to investigate the existence of equilibria in which all but the first buyers' strategies assign positive probability to two offers, as in the second example given at the beginning of this section. Such equilibria need not exist, suggesting that either some buyers' strategies assign positive probability to more than two offers, or that the lower offer in the support of some buyers' strategies is serious as well. It is possible to construct equilibria of the second kind, but doing so generally is intractable, even numerically.

Other comparisons between the two scenarios are less clear-cut. Since it is possible that $q^* = q_K$, Samuelson's (1984) Proposition 1 implies that the outcome with public offers is the preferred one among the outcomes of all bargaining procedures from the first buyer's viewpoint. In particular, since eventual agreement in the hidden case implies that serious (but not winning) offers involve prices higher than the cost of the corresponding indifferent type, the first buyer prefers the outcome of the game with public offers to the outcome in the game with private outcome whenever q^* is sufficiently close to q_K. The same argument applies to the *aggregate* buyers' payoff. Buyers $n \geq 2$ prefer the outcome with hidden offers, although any difference disappears as the discount factor tends to one (see Proposition 4).

From the seller's point of view, the first example of an equilibrium in Section 4 is preferred to the outcome with public offers by all types of the seller, so that this equilibrium outcome is *ex ante* more efficient than the unique outcome with public offers. We have not found any example

[15]See Hörner and Vieille (2007).

in which this conclusion is reversed. However, it is straightforward to show that no equilibrium is second-best efficient. In terms of interim efficiency, the comparison is ambiguous. Considering the second example in Section 4, it is easy to check that very low types prefer the outcome under observable offers, while very high types prefer the outcome under hidden offers.

5 Related Literature and Concluding Comments

Our contribution is related to three literatures. First, several authors have already considered dynamic versions of Akerlof's model. Second, several papers in the bargaining literature consider interdependent values. Third, several papers in the literature on learning address the conditions under which learning occurs. In particular, two papers have investigated the difference between public and private offers in the framework of Spence's signaling model.

Janssen and Roy (2002) consider a dynamic, competitive durable good setting, with a fixed set of sellers. They prove that trade for *all* qualities of the good occurs in finite time. The critical difference lies in the market mechanism. In their model, the price in every period must clear the market. That is, by definition, the market price must be at least as large as the good's expected value to the buyer conditional on trade, with equality if trade occurs with positive probability (this is condition (ii) of their equilibrium definition). This expected value is derived from the equilibrium strategies when such trade occurs with positive probability, and it is assumed to be at least as large as the lowest unsold value even when no trade occurs in a given period (this is condition (iv) of their definition). This implies that the price exceeds the valuation to the lowest quality seller, so that trade must occur eventually. Also related are Taylor (1999), Hendel and Lizzeri (1999), House and Leahy (2004) and Hendel, Lizzeri and Siniscalchi (2005).

In the bargaining literature, Evans (1989), Vincent (1989) and Deneckere and Liang (2006) consider bargaining with interdependent values. Evans (1989) considers a model in which the seller's unit can have one of two values, and assumes that there is no gain from trade if the value is low. He shows that the bargaining may result in an impasse when the buyer is too impatient relative to the seller. In his appendix, Vincent (1989) provides another example of equilibrium in which bargaining breaks down. As in Evans, the unit can have one of two values. It follows from Deneckere and Liang (2006) that his example is generically unique. Deneckere and Liang (2006) generalize these findings by considering an environment in which the unit's quality takes values in an interval. They characterize the (stationary) equilibrium of the game between a buyer and a seller with equal discount factors, in which, as in ours, the uninformed buyer makes all the offers. When the static incentive constraints preclude first-best efficiency, the limiting bargaining outcome involves agreement but delay, and fails to be second-best efficient. Other related contributions include Riley and Zeckhauser (1983), Cramton (1984) and Gul and Sonnenschein (1988).

There is a large literature on learning. Some papers have examined the conditions under which full learning occurs as time passes under rather general conditions (See Aghion, Bolton, Harris and Jullien, 1991). These models, however, are typically cast as decision problems in which Nature's response is exogenously specified. Here instead, the information that is being revealed is a function of the seller's best-reply, which is endogenous. Nöldeke and van Damme (1990) and Swinkels (1999) develop an analogous distinction in Spence's signalling model. Both consider a discrete-time version of the model, in which education is acquired continuously and a sequence of short-run firms submit offers that the worker can either accept or reject. Nöldeke and van Damme consider the case of public offers, while Swinkels focuses mainly on private offers. Nöldeke and van Damme show that there is a unique equilibrium outcome that satisfies the never-a-weak-best-response requirement and that the equilibrium outcome converges to the Riley outcome as the time interval between consecutive periods shrinks. For private offers, Swinkels proves that the sequential equilibrium outcome is unique and that it involves pooling in the limit. The logic driving these results is similar to ours, at least for public offers. In both papers, firms (buyers) are deterred from submitting mutually beneficial offers because rejecting such an offer sends a strong signal to future firms and elicits wage offers so attractive that only very low types would accept the current offer.

There are many conceivable variations to the simple set-up considered here and the reader may wonder how far the results extend to more complex environments. For instance, one may wish to model buyers as being unaware of the number of prior offers that have been submitted. One may want to allow for multiple buyers in each period. One could also allow the seller to choose how much information to disclose; to make offers, or to send messages that are more or less committing, such as list prices in the case of real estate. These extensions are further discussed in the working paper (See Hörner and Vieille, 2007). It might also have occurred to the reader that the impasse result for the public case cannot possibly hold if the horizon is finite, as the situation of the last buyer reduces then to the static case, for which the unique optimal offer is serious. However, it is not hard to see that, for a fixed horizon T, there exists a discount factor $\bar{\delta}_T < 1$ such that, if $\delta > \bar{\delta}_T$, all offers but the last one are necessarily losing (provided $\underline{q} < q_1$): because the last buyer's offer is serious, the objective function of the penultimate buyer inherits the convexity property stressed in Section 3.1. that compels him to submit a losing offer, and the result follows by backward induction.[16]

A crucial assumption throughout the analysis has been that buyers do not receive private signals. Allowing for private information is both economically relevant and likely to drastically affect the results in the public case. In this case, later buyers do not only learn about the seller's type through the seller's earlier actions, but also through the offers that were made by previous buyers, as in the literature on cascades (See Bikhchandani, Hirshleifer and Welch, 1991). A unit might remain on the market for a longer time period either because the seller is particularly reluctant to give up a good that he knows to be of high quality, or because all previous buyers' signals have been unfavorable. The possibility of the latter event is likely to depress later offers, but we conjecture that this might paradoxically help trade by prompting the seller into accepting earlier offers.

References

[1] Aghion, P., Bolton, P., Harris, C. and B. Jullien, 1991. Optimal Learning by Experimentation, *Review of Economic Studies*, **58**, 621-654.

[2] Akerlof, G., 1970. The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism, *Quarterly Journal of Economics*, **84**, 488-500.

[3] Babcock, L. and G. Loewenstein, 1997. Explaining Bargaining Impasse: The Role of Self-Serving Biases, *Journal of Economic Perspectives*, **11**, 109-126.

[16] One could also fix the discount factor $\delta > \bar{\delta}$ and let T tend to infinity. The complete analysis of this case is quite intricate and we have not been able to satisfactorily resolve it.

[4] Bikhchandani, S., Hirshleifer, D. and I. Welch, A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades, *Journal of Political Economy*, **100**, 992-1026.

[5] Coase, R. H., 1972, Durability and Monopoly, *Journal of Law and Economics*, **15**, 143-149.

[6] Cramton, P., 1984. Bargaining with Incomplete Information: An Infinite-Horizon Model with Continuous Uncertainty, *Review of Economic Studies*, **51**, 573-591.

[7] Crawford, V. P., 1982. A Theory of Disagreement in Bargaining, *Econometrica*, **3**, 607-638.

[8] Deneckere, R. and M.-Y. Liang, 2006. Bargaining with Interdependent Values, *Econometrica*, **74**, 1309-1364.

[9] Evans, R., 1989. Sequential Bargaining with Correlated Values, *Review of Economic Studies*, **56**, 499-510.

[10] Fudenberg, D. and J. Tirole, 1991. *Game Theory*. Cambridge, MA: MIT Press.

[11] Gul, F. and H. Sonnenschein, 1991. On Delay in Bargaining with One-Sided Uncertainty, *Econometrica*, **56**, 601-611.

[12] Hendel, I. and A. Lizzeri, 1999. Adverse Selection in Durable Goods Markets, *American Economic Review*, **89**, 1097-1115.

[13] Hendel, I., Lizzeri, A. and M. Siniscalchi, 2005. Efficient Sorting in a Dynamic Adverse-Selection Model, *Review of Economic Studies*, **72**, 467-497.

[14] Hörner, J. and N. Vieille, 2007. Public vs. Private Offers in the Market for Lemons, *Center for Mathematical Studies DP*, Northwestern University.

[15] House, C.L. and J.V. Leahy, 2004. An *sS* Model with Adverse Selection, *Journal of Political Economy*, **112**, 581-614.

[16] Janssen, M. C. W., and S. Roy, 2002. Dynamic Trading in a Durable Good Market with Asymmetric Information, *International Economic Review*, **43**, 1, 257-282.

[17] Nöldeke, G. and E. van Damme, 1990. Signalling in a Dynamic Labour Market, *Review of Economic Studies*, **57**, 1-23.

[18] Riley, J. G. and R. Zeckhauser, 1983. Optimal Selling Strategies: When to Haggle, and When to Hold Firm, *Quarterly Journal of Economics*, **76**, 267-287.

[19] Samuelson, W., 1984. Bargaining under Asymmetric Information, *Econometrica*, **52**, 995-1005.

[20] Swinkels, J., 1999. Educational Signalling with Preemptive Offers, *Review of Economic Studies*, **66**, 949-970.

[21] Taylor, C. 1999. Time-on-the-Market as a Signal for Quality, *Review of Economic Studies*, **66**, 555-578.

[22] Vincent, D., 1989. Bargaining with Common Values, *Journal of Economic Theory*, **48**, 47-62.

[23] Wilson, C., 1980. The Nature of Markets with Adverse Selection, *Bell Journal of Economics*, **11**, 108-130.

Appendix

A Proofs for public offers

A.1 Proof of Theorem 1

The proof is by induction over K. The proof for $K = 0$ is, in most respects, identical to the proof of the induction step, and we therefore provide only the latter. We let a (perfect, Bayesian) equilibrium be given. Recall that $\underline{q}_n$ stands for the lowest type which rejected all offers up to stage n.

Assume that for some $k < K$, the following holds: for every period $n \geq 1$ and after any history h^{n-1}, buyer n offers $c(q_l)$ whenever $\underline{q}_n \in (q_{l+1}, q_l)$ for some $l < k$, and all subsequent offers are losing ones. We will prove that the same conclusion holds for $k + 1$. The proof is broken into the following four steps:

- whenever $\underline{q}_n \in (q_{k+1}, q_k)$, no equilibrium offer of buyer n is accepted by some type $s > q_k$;

- whenever $\underline{q}_n \in (q_{k+1}, q_k)$, if an equilibrium offer of buyer n is accepted by $s = q_k$, then all subsequent offers are losing ones;

- if $\underline{q}_n \in (q_{k+1}, q_k)$ is close enough to q_k, the unique equilibrium offer of buyer n is $c(q_k)$, which the seller accepts if and only his type does not exceed q_k;

- more generally, if $\underline{q}_n = q \in (q_{k+1}, q_k)$, the unique equilibrium offer of buyer n is $c(q_k)$, which the seller accepts if and only his type does not exceed q_k.

Step 1: We let a stage n and an history h^{n-1} be given, such that $q := \underline{q}_n \in (q_{k+1}, q_k)$.

If buyer n submits an offer $p(s)$ with indifferent type $s \in (q_{l+1}, q_l]$ for some $l < k$, the following offer is $c(q_l)$ by the induction hypothesis. Hence, $p(s)$ must solve

$$p(s) - c(s) = \delta(c(q_l) - c(s)),$$

so that buyer n's payoff is

$$\frac{1}{1-q}\int_q^s \{v(t) - \delta c(q_l) - (1-\delta)c(s)\}\, dt.$$

As a function of s, the integral is twice differentiable over the interval $(q_{l+1}, q_l]$, with first and second derivatives given by (up to the constant $1-q$)

$$v(s) - \delta c(q_l) - (1-\delta)c(s) - (1-\delta)c'(s)(s-q),$$

and

$$v'(s) - 2(1-\delta)c''(s)(s-q).$$

Since $(2M_{c'} + M_{c''})(1-\delta) < m$, buyer n's payoff is strictly convex over $(q_{l+1}, q_l]$. Since buyer n's payoff is negative for $s = q_l$, the claim follows.

Step 2: We argue by contradiction. We thus assume that, for some stage n and some history h^{n-1} with $\underline{q}_n \in (q_{k+1}, q_k)$, there is a positive probability that an equilibrium offer p_n by buyer n with indifferent type q_k is eventually followed by a serious offer. This implies $p_n > c(q_k)$. Let $\bar{p} > c(q_k)$ be the supremum of all such offers (with indifferent type q_k), where the supremum is taken over all such stages n and histories h^{n-1}.

Given a stage n and an history h^{n-1}, note that the price $p_n(q)$ with indifferent type $q \leq q_k$ does not exceed $\bar{p}$. Indeed, denote by p^* the supremum of all such prices, and denote by $\tau(q)$ the first buyer submitting an offer $\tilde{p}_{\tau(q)}$ that type q accepts, so that

$$p_n(q) - c(q) = \mathbf{E}\left[\delta^{\tau(q)-n}(\tilde{p}_{\tau(q)} - c(q))|\tau(q) > n\right] \leq \delta\left(\max\{\bar{p}, p^*\} - c(q)\right),$$

hence

$$p^* - c(q) \leq \delta\left(\max\{\bar{p}, p^*\} - c(q)\right),$$

and therefore $p^* \leq \bar{p}$.

Consider a buyer and a history (still denoted by n and h^{n-1}) who submits an equilibrium offer $p_n > (1-\delta)c(q_k) + \delta\bar{p}$ with indifferent type q_k. If instead, buyer n deviates to a serious offer $p(s)$ with indifferent type $s < q_k$, then $p(s)$ does not exceed

$$p(s) \leq (1-\delta)c(s) + \delta\bar{p} \leq (1-\delta)c(q_k) + \delta\bar{p}.$$

By choosing s close enough to q_k, buyer n's payoff, $\frac{1}{1-q}\int_q^s \{v(t) - p(s)\}\,dt$, is thus higher than his equilibrium payoff, $\frac{1}{1-q}\int_q^{q_k} \{v(t) - p_n\}\,dt$, which is a contradiction.

Step 3: Let n and $h^{n-1} \in H^{n-1}$ be given, with $\underline{q}_n = q < q_k$. Consider a potential offer $p(s)$, with indifferent type s. Obviously, $p(s) \geq c(s)$. Observe also that by Step 2, $p(s) - c(s)$ converges to zero as s increases to q_k. Hence, buyer n's payoff, $\frac{1}{1-q}\int_q^s \{v(t) - p(s)\}\,dt$, is at most $\frac{1}{1-q}\int_q^s \{v(t) - c(s)\}\,dt$, and the difference between the two integrals converges to zero, as s increases to q_k. The latter integral, as a function of s, is differentiable, with derivative $v(s) - c(s) - c'(s)(s-q)$, which is positive whenever $s - q < \frac{m}{M_{c'}}$. Thus, for q close enough to q_k, the upper bound, $\frac{1}{1-q}\int_q^s \{v(t) - c(s)\}\,dt$, is increasing over $[q, q_k)$. Hence, for such q, and since an equilibrium is assumed to exist, buyer n's equilibrium offer is $c(q_k)$, which the seller accepts if and only if his type does not exceed q_k.

Step 4: Again, we argue by contradiction. We assume that, for some n and h^{n-1} with $\underline{q}_n > q_{k+1}$, buyer n's strategy assigns a positive probability to serious offers with indifferent type below q_k. Among all such n and h^{n-1}, let $\tilde{q} \in (q_{k+1}, q_k]$ be the supremum of $\underline{q}_n$.

Consider now any n and h^{n-1} with $q := \underline{q}_n \leq \tilde{q}$. By definition of $\tilde{q}$, any offer $p(s)$ with indifferent type $s > \tilde{q}$ is followed by an offer $c(q_k)$ from the next buyer, so that $p(s)$ must satisfy

$$p(s) - c(s) = \delta(c(q_k) - c(s)),$$

and buyer n's payoff writes

$$\frac{1}{1-q}\int_{q}^{s}\left\{v(t)-\delta c(q_k)-(1-\delta)c(s)\right\}dt.$$

As in Step 1, the integral is a strictly convex function of s. Therefore, the indifferent type of any equilibrium offer is either equal to q_k, or lies in the interval $[q,\tilde{q}]$. In the former case, buyer n's offer is $c(q_k)$, and his payoff is positive since $\tilde{q}>q_{k+1}$. In the latter case, buyer n's payoff is at most $(\tilde{q}-q)\left(v(\tilde{q})-c(q)\right)$, which is arbitrarily close to zero, provided q is close enough to $\tilde{q}$. As a consequence, for $q<\tilde{q}$ close to $\tilde{q}$, the unique equilibrium offer of buyer n is $c(q_k)$, with indifferent type q_k. This contradicts the definition of $\tilde{q}$.

To conclude the proof, observe that the following strategy profile, defined recursively, is indeed an equilibrium. To each history h^{n-1}, we associate a unique $\underline{q}_n$ as follows. Let $\underline{q}_0=\underline{q}$ and given $\underline{q}_{n-1}$ and p_{n-1}, define $\underline{q}_n$ as follows. If there exists $q_k=\min\left\{q_{k'}\geq\underline{q}_{n-1}:c\left(q_{k'}\right)\geq p_n\right\}$, let $\underline{q}_n=\max\left\{q,\underline{q}_{n-1}\right\}$, where q solves $p_{n-1}-c(q)=\delta(c(q_k)-c(q))$. If such q_k does not exist, let $\underline{q}_n=1$. Observe that $\underline{q}_n$ only depends on h^{n-1}. Given h^{n-1}, we specify buyer n's belief as being the uniform distribution over $\left[\underline{q}_n,1\right]$ (possibly degenerate on 1). Given h^{n-1}, the buyer's strategy $\sigma_B^n\left(h^{n-1}\right)$ assigns probability 1 to the offer $c(q_k)$, where k is uniquely defined by $\underline{q}_n\in(q_{k+1},q_k]$. Given $\left(h^{n-1},p_n\right)$, the seller's strategy $\sigma_S^n\left(q,h^{n-1},p_n\right)$ assigns probability 1 to accepting if and only if $q\leq\underline{q}_{n+1}$. ■

A.2 Proof of Proposition 1

The proof is in most respects similar to the proof of Theorem 1. We therefore refer the reader to the previous section for a number of technical details.

Observe that a strategic seller of type q rejects any offer less than $(1-\delta)\,q$, since he can always reject the current offer and accept the following (nonnegative) offer. Thus, an offer of 0 is only accepted by the liquidity seller. We denote the corresponding payoff by

$$\Pi^0\left(q\right):=\varepsilon\left(1+\alpha\right)\left(1+q\right)/2,$$

where q is the lowest remaining type.

A difficulty that arises here is that a strategic seller of type q might conceivably accept an offer below q. Indeed, as he anticipates that, in case of a future liquidity shock, he would accept any offer, he might be willing in some circumstances to accept an offer immediately that is below his cost. (In fact, this will happen in equilibrium.) We wish to show first that this cannot happen after histories leading to beliefs that are uniform over $[q,1]$, for q large enough. To see this, observe first that, for q large enough (and ε small enough), $(1+\alpha)(1+q)/2-1 > \Pi^0(q)$. That is, by offering 1, the buyer can secure strictly more than what he gets from any offer that would only be accepted by the liquidity seller. (As in the standard model, no buyer ever offers more than 1, and such an offer is accepted with probability one by the seller.) This implies that, in equilibrium, after such a history, the current buyer (and every subsequent buyer) only submits offers that are accepted with positive probability by the strategic seller. Hence, from this period on, all equilibrium offers must exceed $(1-\delta)q$. Therefore, by rejecting an offer, the seller can expect an offer of at least $(1-\delta)q$ in the following period, a strategy that secures a continuation payoff of at least $\left(1-\delta^2\right)q$. In turn, this means that offers that are accepted with positive probability by the strategic sellers (after such histories) must be at least equal to $\left(1-\delta^2\right)q$, and more generally, by an immediate induction argument, at least equal to $(1-\delta^m)q$ for all m. That is, after any history leading to beliefs that are uniform over $[q,1]$, for q large enough, all offers must be at least equal to q, which was to be shown.

This implies that, after such a history, the indifferent type q' corresponding to any equilibrium offer p is such that $p \geq q'$. Therefore, the payoff of such an offer is (up to the constant $1-q$)

$$(1-\varepsilon)\int_q^{q'}(t-p)\,dt+\varepsilon\int_q^1(t-p)\,dt \leq (1-\varepsilon)\int_q^{q'}\left(t-q'\right)dt+\varepsilon\int_q^1\left(t-q'\right)dt,$$

with equality for $q'=1$. For ε small enough, the right-hand side is strictly increasing in q' for q close enough to 1. Therefore, there exists $\tilde{q}<1$ such that, for any history leading to beliefs that are uniform over $[q,1]$, $q \geq \tilde{q}$, the buyer's unique equilibrium offer is 1. Let β_1 denote the infimum over such $\tilde{q}<1$.

We next claim that when the lowest type is β_1, the current buyer is indifferent between an offer of 0 and a winning offer, so that β_1 is the (unique) solution to $(1+\alpha)(1+q)/2-1=\Pi^0(q)$.

Indeed, observe first that, as in the standard model, the payoff from submitting an offer with indifferent type y in $(\beta_1,1)$ is convex in y. Therefore, no such offer is ever submitted.

Observe next that, whenever the lowest type q exceeds β_1, the unique optimal offer is the winning one. Hence, when $q=\beta_1$, the winning offer yields at least as much as an offer of 0. On the other hand, let $q<\beta_1$ be close to β_1, and such that the buyer submits with positive probability an offer which indifferent type at most β_1.[17] If this offer is accepted with positive probability by the strategic seller, it must exceed $(1-\delta)\,q$, yet the probability that it is accepted by the strategic seller is arbitrarily small. Therefore, it would be strictly worse than offering 0, which is only accepted by the liquidity seller. Thus, for such q's, offering 0 is (weakly) preferred to offering 1. By continuity, this is also true when $q = \beta_1$.

We now point out a feature of the equilibrium that is absent from the standard model. Recall that if the lowest type q is below β_1, a null offer is preferred to a winning offer. Besides, if q is close enough to β_1, a null offer is preferred to any offer with indifferent type below β_1, while offers with indifferent type in $(\beta_1, 1)$ are ruled out by convexity. It follows that there exists an interval $(\tilde{q}, \beta_1)$, such that, whenever the lowest type q belongs to $(\tilde{q}, \beta_1)$, the buyer submits an offer of 0 (which the seller accepts if and only if he is a liquidity seller). Let γ_1 denote the infimum of all such $\tilde{q}$.

This allows us to compute the offer with indifferent type $q \in (\gamma_1, \beta_1)$. Indeed, all following offers are 0. Hence a seller will type q is ready to accept the offer $(1-\delta)\left(q + \delta(1-\varepsilon)\,q + \delta^2(1-\varepsilon)^2\,q + \cdots\right)$, that is, to $\eta_\varepsilon q$, where $\eta_\varepsilon := (1-\delta)/(1-(1-\varepsilon)\,\delta)$. Observe that $\eta_\varepsilon \to 1$ as $\varepsilon \to 0$.

It is useful to introduce the function

$$\Pi(x;q) := (1-\varepsilon)\,\frac{x-q}{1-q}\left((1+\alpha)\,\frac{x+q}{2} - \eta_\varepsilon q\right) + \varepsilon\left((1+\alpha)\,\frac{1+q}{2} - \eta_\varepsilon q\right).$$

It is an analog of the payoff function to a monopsonistic buyer, which was used in the analysis of the standard model. For ε small enough, $\Pi(x;q)$ is strictly concave in x. We then set $\Pi^*(q) := \max_x \Pi(x;q)$. The maximizer is

$$x^*(q) = \frac{\eta_\varepsilon}{2\eta_\varepsilon - (1+\alpha)}\left(q - \frac{\varepsilon}{1-\varepsilon}(1-q)\right),$$

which is strictly increasing in q. Also, for ε small enough, $\Pi^*(q)$ is strictly increasing in q and strictly larger than $\Pi^0(q)$.

Define μ_1 to be the solution of $x^*(\cdot) = \beta_1$. We here claim that (i) γ_1 is the larger of the two roots of the polynomial equation $\Pi(\beta_1;\cdot) = \Pi^0(\cdot)$; (ii) whenever the current lowest type q is such that $q \in (\mu_1, \gamma_1)$, the current buyer submits the offer $\eta_\varepsilon \beta_1$, and all subsequent offers are 0.

Indeed, let the current lowest type q be such that $q < \beta_1$. A winning offer yields a negative payoff, while offers with indifferent type in $(\beta_1, 1)$ are ruled out by convexity. A null offer yields the payoff

[17] To be precise: consider any history leading to a belief that is the uniform distribution over the interval, and such that the current buyer does not submit a winning offer with probability 1.

$\Pi^0(q)$, while an offer which attracts all types up to $x < \beta_1$ yields the payoff $\Pi(x; q)$. Depending on the position of q relative to the largest root of $\Pi(\beta_1; \cdot) = \Pi^0(\cdot)$, either the null offer is optimal or not. This proves (i). If $q \in [\mu_1, \gamma_1)$, the payoff $\Pi(x; q)$ is increasing in $x < \beta_1$. Existence of a maximizer requires then that an offer of $\eta_\varepsilon \beta_1$ is accepted by all types up to β_1, which requires that all future offers be equal to zero. This proves (ii).

We continue to solve backwards for the equilibrium. Define λ_1 to be the solution of $x^*(\cdot) = \gamma_1$. Provided that δ is close enough to one, and ε close enough to zero, we claim that, whenever the current lowest type q is in the interval (λ_1, μ_1), the unique equilibrium offer of the current buyer is $\eta_\varepsilon x^*(q)$, and that this offer is accepted by all types up to $x^*(q)$.

Indeed, let the current lowest type q be such that $q \in (\lambda_1, \mu_1)$. Any offer with indifferent type in $(\mu_1, \gamma_1) \cup (\beta_1, 1)$ is ruled out by convexity. Since $\mu_1 - \lambda_1$ converges to zero as ε converges to zero, any offer with indifferent type in (λ_1, μ_1) yields a payoff close to zero. On the other hand, $x^*(q) \in (\gamma_1, \beta_1)$, hence the offer $\eta_\varepsilon x^*(q)$ yields the payoff $\Pi(x^*(q); q)$, which is bounded away from zero, and exceeds $\Pi^0(q)$, for small ε.

Finally, define β_2 to be the smaller of the two roots of the polynomial equation $\Pi(\gamma_1; \cdot) = \Pi^0(\cdot)$. We claim that, for δ close to 1 and ε close to 0, and whenever the lowest type q belongs to (β_2, λ_1), the unique equilibrium of the current buyer is $\eta_\varepsilon \gamma_1$, that all subsequent offers are 0, and that this offer is accepted by all types up to γ_1.

Indeed, let the current lowest type q be such that $q \in (\beta_2, \lambda_1)$, and assume first that q is close to λ_1. Consider any offer of the current buyer, and denote by q' the indifferent type of the strategic seller. Values $q' \in (\mu_1, \gamma_1) \cup (\beta_1, 1)$ are ruled out at equilibrium, by convexity. Values of $q' \in (q, \lambda_1]$ correspond to offers that are accepted with arbitrarily low probability. Yet, the corresponding offer is bounded away from zero, hence it is dominated by an offer of zero. We next compare values of $q' \in [\lambda_1, \mu_1]$ to values of $q' \in (\gamma_1, \beta_1)$. On the one hand, the probabilities that such offers are accepted differ by an amount which is of the order of $\gamma_1 - \mu_1$, and hence is bounded away from zero. On the other hand, any offer with $q' \in [\lambda_1, \mu_1]$, is followed by an offer which is at least $\eta_\varepsilon \gamma_1$, while offers with $q' \in (\gamma_1, \beta_1)$ are at most $\eta_\varepsilon \beta_1$, hence the difference between two such offers converges to zero as δ converges to 1, and ε converges to 0. Finally, and since $q \in (\beta_2, \lambda_1)$, the buyer's payoff $\Pi(q'; q)$ is decreasing in $q' \in (\gamma_1, \beta_1)$, and its supremum exceeds $\Pi^0(q)$. Therefore, assuming an equilibrium exists, the supremum of $\Pi(q'; q)$

is achieved at $q' = \gamma_1$.

B Proofs for private offers

B.1 Preliminaries

The remainder of the appendix is organized as follows. As a preliminary, we set up some additional notation, and state a few important facts. We then prove Propositions 3 through 7, though in a different order. We start with Propositions 4 and 5. We then need to prove Proposition 6 – with the exception of the upper bound on N_0. Indeed, it is a logical preliminary to Proposition 7, which we prove next, and its proof is instrumental in the proof of Proposition 3.

A strategy of buyer n is a probability distribution σ_B^n over offers. We denote $\tilde{p}_n$ the random offer by buyer n. Any profile σ_B of such distributions induces a probability distribution over sequences of offers, which we denote $\mathbf{P}_{\sigma_B}$. Expectation with respect to $\mathbf{P}_{\sigma_B}$ is denoted by $\mathbf{E}_{\sigma_B}$.

If a seller with type q declines the first offer, and plans to accept an offer at a (random) time $\tau > 1$, his expected payoff is $\mathbf{E}_{\sigma_B}\left[\delta^{\tau-1}(\tilde{p}_\tau - c(q))\right]$. His optimal continuation payoff is thus $\sup_{\tau>1}\mathbf{E}_{\sigma_B}\left[\delta^{\tau-1}(\tilde{p}_\tau - c(q))\right]$, where the supremum is taken over all stopping times $\tau > 1$, and the offer $p_1(q)$ for which type q is indifferent between accepting and declining, is given by

$$p_1(q) - c(q) = \sup_{\tau>1}\mathbf{E}_{\sigma_B}\left[\delta^{\tau}\left(\tilde{p}_\tau - c(q)\right)\right].$$

For concreteness, we assume that a seller accepts an offer whenever indifferent. Therefore, a seller with type q accepts the offer from buyer $\tau(q) := \inf\{n \geq 1 : \tilde{p}_n \geq p_n(q)\}$. Similarly, the offer $p_n(q)$ with indifferent type q, is given in stage n by

$$\begin{aligned} p_n(q) - c(q) &= \sup_{\tau>n}\mathbf{E}_{\sigma_B}\left[\delta^{\tau-n}(\tilde{p}_\tau - c(q))\right] \\ &= \mathbf{E}_{\sigma_B}\left[\delta^{\tau_n(q)-n}(\tilde{p}_{\tau_n(q)} - c(q))\right] \end{aligned}$$

where $\tau_n(q) := \inf\{k > n : \tilde{p}_k \geq p_n(q)\}$.

It follows that

$$p_n(q) - c(q) \geq \delta\left(p_{n+1}(q) - c(q)\right), \tag{5}$$

with equality if and only if buyer $n+1$ makes no offer above $p_{n+1}(q)$: competition between successive buyers prevents p_n from being much below p_{n+1}.[18] Using a version of the envelope theorem, the function

[18] On the other hand, it may happen that p_{n+1} is much below p_n, as is e.g. the case if buyer $n+1$ makes high offers with high probability, followed by losing offers.

p_n has a left-derivative everywhere, given by

$$D^- p_n(q) = c'(q) \left(1 - \mathbf{E}_{\sigma_B} \left[\delta^{\tau_n(q)-n}\right]\right). \tag{6}$$

Note that $\mathbf{E}_{\sigma_B} \left[1 - \delta^{\tau_n(q)-n}\right]$ is non-increasing in q, and therefore, p_n is convex if the cost function is convex.

The function p_n may be interpreted as an (inverse) offer function faced by buyer n, and (6) provides a direct link between the slope of this offer function at q, and the discounted time at which a seller with type q expects to receives an acceptable offer – the earlier the discounted time, the lower the slope of p_n.

We now comment on the beliefs of the various buyers. Since offers are private, the belief of buyer n need not be a uniform distribution. Recall that $F_n(q)$ is the (unconditional) probability that the seller is of type $t \leq q$, and rejects offers from buyers 1 through $n-1$. Letting $f_n(q) = \frac{1}{1-\underline{q}} \prod_{k=1}^{n-1} \mathbf{P}_{\sigma_B} (\tilde{p}_k < p_k(q))$ denote the (normalized) probability that a seller with type q rejects the first $n-1$ offers, one has

$$F_n(q) = \int_{\underline{q}}^{q} f_n(t)dt.$$

Observe that f_n is left-continuous and non-decreasing, so that F_n is non-decreasing, convex, and admits a left-derivative $D^- F_n(q) = f_n(q)$. We last introduce $\underline{q}_n = \max\{q \in [\underline{q}, 1] : F_n(q) = 0\}$, the lowest type that rejects the first $n-1$ offers with probability 1.

With these notations at hand, the expected payoff $\pi_n(q)$ of buyer n, when submitting the offer $p_n(q)$, is given by

$$\pi_n(q) := \int_{\underline{q}}^{q} (v(t) - p_n(q)) f_n(t)dt. \tag{7}$$

We denote by $\overline{v}_n(q) = \frac{1}{F_n(q)} \int_{\underline{q}}^{q} v(t) f_n(t)dt$ the average valuation of types below q, as seen by buyer n. Then (7) rewrites

$$\pi_n(q) = F_n(q) \left(\overline{v}_n(q) - p_n(q)\right),$$

which reads as the probability that the n-th offer is accepted, $F_n(q)$, times the conditional payoff, given that trade takes place. The payoff function π_n has a left-derivative everywhere, equal to

$$D^- \pi_n(q) = (v(q) - p_n(q)) f_n(q) - D^- p_n(q) F_n(q). \tag{8}$$

We stress that the quantities introduced so far, p_n, F_n, f_n, $\underline{q}_n$, π_n, $\overline{v}_n$ all depend on the profile σ under consideration, although the notation does not indicate this. Throughout the appendix, we let an equilibrium σ^* be given, and no confusion should arise.

For conciseness, we will refer to the indifferent type associated with an offer p as a *type offer*. For instance, the statement *buyer n submits a type offer q* is logically equivalent to the statement *buyer n submits an offer with indifferent type q*.

Since π_n is continuous, the equilibrium payoff π_n^* of buyer n is equal to $\max_{[\underline{q},1]} \pi_n$, and one has $\pi_n(q) = \pi_n^*$ for every $q \in T_n$, where T_n is the support of the random type offer of buyer n.[19]

Finally, we state a preliminary observation that is used repeatedly below. We consider a buyer, $n+1$, who only submits type offers bounded away from $\underline{q}_{n+1}$ – the lowest remaining type. We prove that the previous buyer then makes no serious type offer below the lowest serious type offer of buyer $n+1$.

Lemma 2 *Assume $\underline{q}_{n+2} > \underline{q}_{n+1}$, for some $n \in \mathbb{N}$. Then buyer n submits no type offer in $(\underline{q}_n, \underline{q}_{n+2})$.*[20] *In particular, $\underline{q}_{n+1} = \underline{q}_n$, buyer n submits a losing offer with positive probability, and $\pi_n^* = 0$.*

The inequality $\underline{q}_{n+2} > \underline{q}_{n+1}$ is satisfied whenever $\pi_{n+1}^* > 0$, since $\pi_{n+1}(\underline{q}_{n+1}) = 0$, and $\pi_{n+1}(q)$ is therefore arbitrarily close to zero in a neighborhood of $\underline{q}_{n+1}$. Lemma 2 thus implies that there are no two consecutive buyers with positive equilibrium payoff.

Proof. Let a type $q \in (\underline{q}_n, \underline{q}_{n+2})$ be given. By assumption, a seller with type q plans to accept buyer $n+1$'s offer with probability one, were he to decline buyer n's offer. Thus, the seller's continuation payoff is $\delta \mathbf{E}_{\sigma_B^*}[\tilde{p}_{n+1} - c(q)]$, and therefore, $p_n(q) = (1-\delta)c(q) + \delta \mathbf{E}_{\sigma_B^*}[\tilde{p}_{n+1}]$. Since $\delta > \overline{\delta}$, this implies that $v(q) - p_n(q)$ is increasing.

Set $z := \inf\{q \in \left[\underline{q}_n, 1\right] : v(q) \geq p_n(q)\}$ (with $\inf \emptyset = \underline{q}_n$). Note that $D^-\pi_n(q) = (v(q) - p_n(q)) f_n(q) - c'(q)F_n(q) < 0$ on $\left(\underline{q}_n, z\right]$. On the other hand, on the interval $\left(z, \underline{q}_{n+2}\right]$, $D^-\pi_n$ is upper semicontinuous since f_n is non-decreasing and left-continuous. We now prove that $D^-\pi_n$ is increasing.

Since f_n is non-decreasing, one has

$$\liminf_{x \nearrow q} \frac{(v(q) - p_n(q)) f_n(q) - (v(x) - p_n(x)) f_n(x)}{q - x} \geq (v'(q) - D^- p_n(q)) f_n(q),$$

[19] That is, T_n is the smallest closed set of type offers that is assigned probability one by σ_B^n.

[20] That is, $\sigma_B^{*,n}$ assigns probability zero to type offers in $(\underline{q}_n, \underline{q}_{n+2})$.

thus,

$$\begin{aligned}\liminf_{x\nearrow q}\frac{D^-\pi_n(q)-D^-\pi_n(x)}{q-x} &\geq (v'(q)-(1-\delta)c_n'(q))f_n(q)-(1-\delta)\left(c''(q)F_n(q)+c'(q)f_n(q)\right)\\ &\geq v'(q)-(1-\delta)\left(2c'''(q)f_n(q)\right)\\ &> 0\end{aligned}$$

where the first inequality holds since $F_n(q) \leq f_n(q)$ and the second one since $\delta \geq \bar{\delta}$.

Since $D^-\pi_n$ is upper semicontinuous, this implies that $D^-\pi_n$ is strictly increasing over $\left(z, \underline{q}_{n+2}\right]$, hence π_n is strictly convex over $\left[z, \underline{q}_{n+2}\right]$.

To summarize, π_n is continuous, decreasing over $\left[\underline{q}_n, z\right]$, and strictly convex over $\left[z, \underline{q}_{n+2}\right]$. Therefore, it has no maximum over $(\underline{q}_n, \underline{q}_{n+2})$. This proves the first claim.

If buyer n does not submit a losing offer with positive probability, then his lowest type offer is at least $\underline{q}_{n+2}$, which implies $\underline{q}_{n+1} \geq \underline{q}_{n+2}$ – a contradiction.

In particular, $\pi_n^* = \pi_n(\underline{q}_n) = 0$. This concludes the proof of the lemma. ■

B.2 Proof of Proposition 4

We here prove that equilibrium payoffs are very small. Proposition 8 below implies Proposition 4.

Proposition 8 *The equilibrium payoff of each buyer n is at most*

$$\pi_n^* \leq \frac{2}{m_{v'}(1-\underline{q})}(1-\delta)^2\left(v\left(\underline{q}_{n+1}\right)-c\left(\underline{q}_{n+1}\right)\right)^2.$$

Proof. Consider a buyer n with positive equilibrium payoff, $\pi_n^* > 0$, so that $\underline{q}_{n+1} > \underline{q}_n$ and

$$\pi_n^* = F_n(\underline{q}_{n+1})\left(\overline{v}_n\left(\underline{q}_{n+1}\right)-p_n\left(\underline{q}_{n+1}\right)\right).$$

We bound below each of the two terms.

Note that $\pi_n^* > 0$ implies $\pi_{n+1}^* = 0$, which implies in turn

$$p_{n+1}(\underline{q}_{n+1}) \geq v(\underline{q}_{n+1}) \tag{9}$$

(for otherwise buyer $n+1$ would get a positive payoff when making a type offer just above $\underline{q}_{n+1}$). Note also that $\overline{v}_n(\underline{q}_{n+1}) < v(\underline{q}_{n+1})$, and that $\pi_n^* > 0$ implies $p_n(\underline{q}_{n+1}) \leq \overline{v}_n(\underline{q}_{n+1})$:

$$p_n(\underline{q}_{n+1}) \leq \overline{v}_n(\underline{q}_{n+1}) < v(\underline{q}_{n+1}). \tag{10}$$

Recall finally (5):

$$p_n(\underline{q}_{n+1}) - c(\underline{q}_{n+1}) \geq \delta \left(p_{n+1}(\underline{q}_{n+1}) - c\left(\underline{q}_{n+1}\right)\right). \tag{11}$$

We rely on (9), (10) and (11) to prove first that the expected payoff conditional on trade, $\overline{v}_n(\underline{q}_{n+1}) - p_n(\underline{q}_{n+1})$, is at most of the order $1-\delta$. By (10), then (9), then (11), one has

$$\overline{v}_n(\underline{q}_{n+1}) < v(\underline{q}_{n+1}) \leq p_{n+1}(\underline{q}_{n+1}) \leq \frac{1}{\delta}\left\{ p_n(\underline{q}_{n+1}) - (1-\delta)c(\underline{q}_{n+1})\right\}$$

hence

$$\overline{v}_n(\underline{q}_{n+1}) - p_n(\underline{q}_{n+1}) \leq \frac{1-\delta}{\delta}\left(v\left(\underline{q}_{n+1}\right) - c\left(\underline{q}_{n+1}\right)\right). \tag{12}$$

Next, we argue that $F_n(\underline{q}_{n+1})$ is at most of the order $1-\delta$. Substituting (9) into (11) yields

$$\delta v(\underline{q}_{n+1}) \leq p_n(\underline{q}_{n+1}) - (1-\delta)c(\underline{q}_{n+1}),$$

which then implies, using the first half of (10),

$$\overline{v}_n(\underline{q}_{n+1}) \geq \overline{q} := \delta v(\underline{q}_{n+1}) + (1-\delta)c(\underline{q}_{n+1}). \tag{13}$$

The intuition now goes as follows. If the probability $F_n(\underline{q}_{n+1})$ is non-negligible, then the computation of $\overline{v}_n(\underline{q}_{n+1})$ must involve a significant fraction of low types, and $\overline{v}_n(\underline{q}_{n+1})$ is therefore bounded away from $v(\underline{q}_{n+1})$, which stands in contradiction with (13). To verify formally this claim, we compute the highest value for $F_n(\underline{q}_{n+1})$ which is consistent with (13), and compute the value Ω of the infinite-dimensional, linear problem $(\mathcal{P})$:

$$\mathcal{P} : \sup \int_{\underline{q}}^{q} f(t)dt,$$

where the supremum is taken over the set $\mathcal{F}$ of non-decreasing, left-continuous functions with values in $\left[0, \frac{1}{1-\underline{q}}\right]$, and such that $\int_{\underline{q}}^{q} v(t)f(t)dt \geq \overline{v}\int_{\underline{q}}^{q} f(t)dt$. The analysis of $(\mathcal{P})$ is standard. When endowed with the Levy distance, $\mathcal{F}$ is compact, and the objective of $(\mathcal{P})$, continuous, hence there is an optimal solution, f^*. Since $v(\cdot) - \overline{v}$ is strictly increasing, the solution f^* must be of the form $f^*(t) = 1_{t>q^*} \times \frac{1}{1-\underline{q}}$ for some q^*. The location of q^* is dictated by the constraint: $\int_{q^*}^{q} v(t)dt = \overline{v}(q-q^*)$. Plugging the inequality $v(t) \geq v(q) - M_{v'}(q-t)$ for all $t \in [q^*, q]$ into the constraint yields

$$q - q^* \leq \frac{2}{m_{v'}}(v(q) - \overline{v}).$$

Therefore,

$$\Omega \leq \frac{2(1-\delta)}{m_{v'}(1-\underline{q})}\left(v(\underline{q}_{n+1}) - c(\underline{q}_{n+1})\right). \tag{14}$$

Collecting (12) and (14) then yields

$$\pi_n^* \leq \frac{2}{m_{v'}(1-\underline{q})}(1-\delta)^2\left(v\left(\underline{q}_{n+1}\right) - c\left(\underline{q}_{n+1}\right)\right)^2,$$

as desired. ■

B.3 Proof of Proposition 5

The intuition for the proof is as follows. A seller with type $\underline{q}_{n_1+1} < \underline{q}_{n_2+1}$ expects to receive an acceptable offer at stage n_2 at the latest. Thus, the difference $n_2 - n_1$ is directly linked to the discounted time at which a seller with type $\underline{q}_{n_1+1}$ expects to trade. In Lemma 3, we first provide a lower bound on $D^- p_n(q)$. We then rely on the relation between $D^- p_n\left(\underline{q}_{n_1+1}\right)$ and the discounted time at which type $\underline{q}_{n_1+1}$ trades.

Lemma 3 *For any buyer n, and any serious type offer $q > \underline{q}_n$ in T_n, one has $D^- p_n(q) \geq \frac{m_{v'}}{2}(1-\underline{q})$.*

The proof of Lemma 3 uses the following technical inequality.

Lemma 4 *Let $h : [\underline{q}, 1] \to \mathbb{R}_+$ be non-decreasing. Then, for any $[a, b] \subseteq \left[\underline{q}, 1\right]$, one has*

$$\frac{h(b)}{\int_a^b h(t)dt}\int_a^b v(t)h(t)dt + \frac{m_{v'}}{2}\int_a^b h(t)dt \leq v(b)h(b) \tag{15}$$

(with the convention $\frac{0}{0} = 0$).

Proof. By induction over the cardinality of the range of v, it is easily checked that (15) holds whenever v is a step function. The result then follows using a limit argument. ■

Proof of Lemma 3. Since π_n is maximal at q, one has $D^-\pi_n(q) \leq 0$, that is,

$$(v(q) - p_n(q))f_n(q) - D^- p_n(q)F_n(q) \leq 0. \tag{16}$$

On the other hand, since $q \in T_n$, one has $\pi_n(q) = \pi_n^*$ which, since q is a serious type offer, implies $v(q) \geq \overline{v}_n(q) \geq p_n(q)$. Plugging these inequalities into (16), one obtains $f_n(q)\left(v\left(q\right) - \overline{v}_n\left(q\right)\right) - D^- p_n(q)F_n(q) \leq 0$ or, equivalently,

$$f_n(q)v(q) \leq f_n(q)\frac{\int_{\underline{q}}^q v(t)f_n(t)dt}{\int_{\underline{q}}^q f_n(t)dt} + D^- p_n(q)\int_{\underline{q}}^q f_n(t)dt.$$

The result then follows by applying Lemma 4. ■

Proof of Proposition 5. Recall from (6) that $D^- p_n(q) = c'(q)(1 - \mathbf{E}_{\sigma_B^*}\left[\delta^{\tau_n(q)-n}\right])$, where $\tau_n(q) = \inf\{k > n : \tilde{p}_k \geq p_k(q)\}$. Consider buyer $n = n_1$, and $q = \underline{q}_{n_1+1} > \underline{q}_{n_1}$. Since $q \in T_n$, and by Lemma 3, one has $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau_n(q)-n}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. On the other hand, $\tau_n(q) \leq n_2$, $\mathbf{P}_{\sigma_B^*}$-a.s. This implies $\delta^{n_2-n_1} \leq 1 - \dfrac{m_{v'}}{2M_{c'}}$ and thus,

$$n_2 - n_1 \geq \frac{1}{1-\delta} \times \frac{m_{v'}}{m_{v'} + 2M_{c'}},$$

as desired. ■

For later use, we note that the very same argument, applied to $n = 1$, and to any serious offer $q \in T_1$, ensures that $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau(1)}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. This yields the upper bound in Proposition 3.

B.4 Proof of Proposition 6

We here prove Proposition 9 below. It corresponds to Proposition 6, except for the upper bound on N_0, which is established in the proof of Proposition 3.

Proposition 9 *There is a stage N_0 such that:*

P1 $T_n \subseteq \{\underline{q}_{N_0}, 1\}$*, for all $n \geq N_0$;*

P2 $\max T_n < 1$*, for all $n < N_0$.*

In addition, $\pi_n^* = 0$, for all $n \geq N_0$.

There may be several (consecutive) stages consistent with **P1** and **P2**. Without further notice, we choose N_0 to be the first of those stages.

Proof. Define $N := 1 + \max\{n : \max T_n < 1 \text{ and } F_n(1) \geq \frac{\nu}{M_{c'}}\}$, where $\nu = \min_{[\underline{q},1]} (c(\cdot) - c(\cdot))$. Since $\lim_n F_n(1) = 0$, the stage N is well-defined, and either $F_N(1) < \frac{\nu}{M_{c'}}$, or $\max T_N = 1$.

To start with, assume that the latter holds. We prove that **P1** and **P2** hold with $N_0 = N$.

Since $1 \in T_{N_0}$, one has $\pi_{N_0}(1) = F_{N_0}(1)\,(\overline{v}_{N_0}(1) - c(1)) \geq 0$, and thus $\pi_n(1) \geq 0$, for all $n \geq N_0$. We argue by contradiction, and assume that $T_n \cap (\underline{q}_{N_0}, 1) \neq \emptyset$, for some $n \geq N_0$, and we call n the first such stage. One thus has $\overline{v}_{n+1}(1) > \overline{v}_n(1)$, hence $\pi_{n+1}(1) > 0$ so that $\pi_{n+1}^* > 0$, and thus $\underline{q}_{n+2} > \underline{q}_{n+1}$. This implies that buyer $n+1$ makes a winning offer with probability one. (Otherwise indeed, the equilibrium payoff of buyer $n+2$ would be at least $\pi_{n+2}(1) \geq \pi_{n+1}(1)$, and both buyers $n+1$ and $n+2$ would have a positive payoff – a contradiction.) Put otherwise, $\underline{q}_{n+2} = 1$. By Lemma 2, buyer n makes no serious type offer in $(\underline{q}_n, \underline{q}_{n+2}) = (\underline{q}_{N_0}, 1)$. This is the desired contradiction. Therefore, $T_n \subseteq \{\underline{q}_{N_0}, 1\}$, for all $n \geq N_0$.

Assume now that $F_N(1) < \frac{\nu}{M_{c'}}$, and that $\overline{q} := \max_{n<N} \max T_n < 1$. We prove that no buyer $n \geq N_0$ ever submits a serious type offer in $(\overline{q}, 1)$. Since $\lim_n F_n(1) = 0$, this implies that some buyer $n \geq N$ eventually submits a winning offer with positive probability. Letting N_0 be the first such buyer, one then has $T_n \subseteq \{\underline{q}_{N_0}, 1\}$, using the same proof as above, and the result follows.

We prove our claim inductively. Let $n \geq N$ be given, and assume that none of the buyers $N, \ldots, n-1$ submits a type offer in $(\overline{q}, 1)$. Since this is also true for buyers $k < N$, one has $f_n(t) = \frac{1}{1-\underline{q}}$, for $t \in [\overline{q}, 1]$.

For $q \geq \overline{q}$, define

$$\begin{aligned} \tilde{\pi}_n(q) &= \int_{\underline{q}}^{q} (v(t) - c(q)) f_n(t) dt \\ &= \tilde{\pi}_n(\overline{q}) + \frac{1}{1-\underline{q}} \int_{\overline{q}}^{q} (v(t) - c(q)) dt. \end{aligned}$$

This is the payoff that would accrue to buyer n if he were the last buyer, or alternatively if all buyers following n would only submit losing offers. Thus, $\tilde{\pi}_n(q) \geq \pi_n(q)$, with equality if $q = 1$. The derivative of $\tilde{\pi}_n$ is given by

$$\begin{aligned} \tilde{\pi}'_n(q) &= \frac{1}{1-\underline{q}} (v(q) - c(q)) - c'(q)(F_n(q) - F_n(\overline{q})) \\ &\geq \frac{\nu}{1-\underline{q}} - M_{c'} F_n(1) > 0. \end{aligned}$$

Therefore, $\tilde{\pi}_n$ is increasing over the interval $[\overline{q}, 1]$, so that $\pi_n(1) > \pi_n(q)$ for every $q \in [\overline{q}, 1)$: buyer n makes no offer in $(\overline{q}, 1)$.

It remains to prove that $\pi_n^* = 0$ for all $n \geq N_0$. It suffices to prove that $\pi_{N_0}^* = 0$. Assume to the contrary that $\pi_{N_0}^* > 0$. Then buyer N_0 would make a winning offer with probability one, for otherwise $\pi_{N_0+1}^*$ would also be positive. Therefore, by Lemma 2, buyer $N_0 - 1$ would make no serious type offer in $(\underline{q}_{N_0-1}, \underline{q}_{N_0+1}) = (\underline{q}_{N_0-1}, 1)$, which would stand in contradiction to the definition of N_0. ■

B.5 Proof of Proposition 7

For convenience, we recall the statement of Proposition 7.

Proposition 10 *All buyers $n < N_0$ submit a serious offer with positive probability. No buyer $n < N_0$ uses a pure strategy, with the possible exception of the first buyer.*

Proof. By definition of N_0, buyer $N_0 - 1$ makes a serious, non-winning offer with positive probability.

We start with the first statement. We argue by contradiction, and assume that $T_n = \{\underline{q}_n\}$, for some $n < N_0$. (In particular, $\pi_n^* = 0$.) Let $n_* > n$ be the first buyer following n who submits a serious offer with positive probability. Since $n_* \leq N_0 - 1$, one has $\overline{q}_{n_*} := \max T_{n_*} < 1$.

Because of discounting and using (5) inductively, one has

$$p_n(\overline{q}_{n_*}) - c(\overline{q}_{n_*}) = \delta^{n_* - n}(p_{n_*}(\overline{q}_{n_*}) - c(\overline{q}_{n_*})),$$

hence $p_n(\overline{q}_{n_*}) < p_{n_*}(\overline{q}_{n_*})$.

Since buyers $n \leq k < n_*$ only submit losing offers, the distribution of types faced by buyers n and n_* is the same, and $\overline{v}_n(\overline{q}_{n_*}) = \overline{v}_{n_*}(\overline{q}_{n_*})$. It follows that $\pi_n(\overline{q}_{n_*}) > \pi_{n_*}(\overline{q}_{n_*}) = \pi_{n_*}^* \geq 0$ – a contradiction. This concludes the proof of the first statement.

Consider now an arbitrary buyer n, with $1 < n < N_0$. If buyer n assigns probability one to a specific type offer, it must be to a serious type offer $q_n > \underline{q}_n$, and then $\underline{q}_{n+1} = q_n$. On the other hand, $p_{n-1}(q_n) - c(q_n) = \delta(p_n(q_n) - c(q_n))$, hence $p_{n-1}(q_n) < p_n(q_n)$. By Lemma 2, $\pi_{n-1}^* = 0$ and buyer $n-1$ makes no type offer in $(\underline{q}_{n-1}, q_n)$, hence $\overline{v}_{n-1}(q_n) = \overline{v}_n(q_n)$. As above, this implies $\pi_{n-1}(q_n) > 0$ – a contradiction. ■

B.6 Equilibrium delay

Recall that $\tau(1) = \inf\{n : \tilde{p}_n = c(1)\}$ is the first buyer who submits a winning offer. From the proof of Proposition 5, we know that $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau(1)}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. We here proceed to provide a lower bound for $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau(1)}\right]$, which is independent of δ and of the equilibrium σ^*.

Let N_0 be given by Proposition 9, and define $N_1 := \inf\{n : F_n(1) \leq \frac{\nu}{M_{c'}}\}$. We proceed in three steps:

Step 1 : one has $N_1 \leq \frac{C_1}{1-\delta}$, where C_1 is independent of $\delta \geq \bar{\delta}$ and of the equilibrium σ^*.

Let $\overline{q} := \max_{n<N_1} \max T_n$ be the highest type offer that may be submitted by some buyer $n < N_1$. From the proof of Proposition 9, we know that either $\overline{q} = 1$, in which case $N_0 = N_1$, or that no buyer ever submits a type offer in $(\overline{q}, 1)$.

Step 2 : One has $N_0 - N_1 \leq \frac{C_0}{1-\delta}$, where C_0 is independent of $\delta \geq \bar{\delta}$ and of the equilibrium σ^*;

Step 3 : One has $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau(1)-N_0}\right] \geq C_2$, where $C_2 > 0$ is independent of $\delta \geq \bar{\delta}$, and of the equilibrium σ^*.

By **Steps 1–3**, one thus has $\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau(1)}\right] \geq C_2 \times \left(\delta^{\frac{1}{1-\delta}}\right)^{C_0+C_1}$, and the result follows since $\delta^{\frac{1}{1-\delta}} \geq e^{-\bar{\delta}}$ for every $\delta \geq \bar{\delta}$.

Step 1 and **Step 2** make use of the following technical result, stated without proof, which links the discounted expectation of a random variable (hereafter, r.v.) to its tail distribution.

Lemma 5 *Let τ be a random time with integer values, and such that*

$$\mathbf{E}\left[\delta^{\tau-n} | \tau > n\right] \geq a, \tag{17}$$

for some $a > 0$ and all $n \geq 1$. Then $\mathbf{P}\left(\tau \geq \frac{1}{ab(1-\delta)}\right) \leq b$, for all $b > 0$.

Our computation of C_1 and C_0 involves three parameters. We choose ϕ, ψ s.t. $0 < \phi < \psi < \nu$, and $\eta \in (0, \frac{\psi-\phi}{2v(1)})$. Next, set $K = 1 + \left\lceil \frac{M_{v'}(1-\underline{q}}{2(\nu-\psi)} \right\rceil$, $\varepsilon = \frac{1}{2}\eta^K$, and $a = \frac{\phi\delta}{c(1)-c(\underline{q})}$. We express C_1 and C_0 as functions of these constants. We make no attempt at optimizing the choice of ϕ, ψ and η.

Step 1: Define $C_1 := 1 + \ln \frac{M_{v'}}{\nu} \frac{1}{a} \frac{1}{\varepsilon^2}$. We prove that $N_1 \leq C_1/(1-\delta)$.

For a given stage n, we let $q_n := \max\{q \in [\underline{q}, 1] : \overline{v}_n(q) > c(q) + \phi\}$. Since $\overline{v}_n(1) \leq c(1)$ one has $q_n < 1$. On the other hand, since $\overline{v}_n(\underline{q}_n) = v(\underline{q}_n)$, one also has $q_n > \underline{q}_n$.

The proof is organized as follows. In Claim 1 below, we first prove that, conditional on the seller having rejected all previous offers, the probability that the seller's type does not exceed q_n, is bounded away from zero. That is, from buyer n's viewpoint, types below q_n have a significant probability.

<u>Claim 1</u>: One has $F_n(q_n) \geq 2\varepsilon F_n(1)$.

Next, we observe that, since $\overline{v}_n(q_n)$ is bounded away from q_n, the offer $p_n(q_n)$ with independent type q_n is also bounded away from $c(q_n)$. Therefore, it is likely that type q_n receives acceptable offers shortly after stage n (for otherwise, he would accept a price close to $c(q_n)$). In Claim 2, we use this insight to prove that conditional on the seller having rejected all previous offers, it is likely that type q_n accepts an offer within $1/a\varepsilon(1-\delta)$ additional stages.

Claim 2: One has $F_{n+N_{a,\varepsilon}}(q_n) < \varepsilon F_n(1)$, where $N_{a,\varepsilon} := \dfrac{1}{a\varepsilon(1-\delta)}$ is defined as in the proof of Lemma 5.

The assertion of **Step 1** immediately follows from Claims 1 and 2. Indeed, observe that, for a given q, $F_n(1) - F_n(q)$ is the probability that the seller rejects all offers from buyers $1, 2, \ldots, n-1$, and has a type t in $[q, 1]$. This difference is non-increasing in n, hence

$$F_{n+N_{a,\varepsilon}}(1) - F_{n+N_{a,\varepsilon}}(q_n) \leq F_n(1) - F_n(q_n).$$

By Claims 1 and 2, this yields

$$F_{n+N_{a,\varepsilon}}(1) \leq (1-\varepsilon)F_n(1)$$

and thus also, $F_{1+iN_{a,\varepsilon}}(1) \leq (1-\varepsilon)^i$, for all $i \geq 1$.

In particular, $F_n(1) < \frac{\nu}{M_{c'}}$ as soon as $n \geq 1 + C_1/(1-\delta)$, as desired.

Proof of Claim 1. We introduce an auxiliary sequence of types which is defined by $y_0 = \underline{q}$ and

$$y_{j+1} = \max\{q \in [0,1] : \mathbf{E}\left[v(t) | t \in [y_j, q]\right] \geq c(q) + \psi\},$$

until $y_J = 1$. In particular, $\mathbf{E}\left[v(t)|t \in [y_j, y_{j+1}]\right] = c(y_{j+1}) + \psi$ for $j < J-1$. On the other hand, observe that

$$\mathbf{E}\left[v(t)|t \in [y_j, y_{j+1}]\right] = \frac{1}{y_{j+1} - y_j} \int_{y_j}^{y_{j+1}} v(t)dt \geq v(y_{j+1}) - \frac{M_{v'}}{2}(y_{j+1} - y_j).$$

Since $v(y_{j+1}) - c(y_{j+1}) \geq \nu$, this implies that

$$y_{j+1} - y_j \geq \frac{2(\nu - \psi)}{M_{v'}}, \text{ for } j < J-1,$$

hence $J \leq K$.

For a given stage n, let $j_n := \min\{j : F_n(y_j) \geq \eta^{K-j} F_n(1)\}$. We now check that $q_n \geq y_{j_n}$, which yields $F_n(q_n) \geq F_n(y_{j_n}) \geq 2\varepsilon F_n(1)$, as desired.

There is nothing to prove if $j_n = 0$, hence assume $j_n > 0$. By definition of j_n, one has $F_n(y_{j_n-1}) < \eta F_n(y_{j_n})$: conditional on $t \leq y_{j_n}$, it is very likely that buyer n faces a type in $[y_{j_n-1}, y_{j_n}]$. Hence,[21]

$$\left|\mathbf{E}_{F_n}\left[v(t)|t \leq y_{j_n}\right] - \mathbf{E}_{F_n}\left[v(t)|t \in [y_{j_n-1}, y_{j_n}]\right]\right| \leq 2\eta v(1), \tag{18}$$

[21] denoting by $\mathbf{E}_{F_n}$ the expectation under the belief held by buyer n

since $|\mathbf{E}(X) - \mathbf{E}(X1_A)| \leq 2\mathbf{P}(\overline{A}) \sup |X|$ for every bounded r.v. X and every event A. The first expectation in (18) is $\overline{v}_n(y_{j_n})$, while the second one is at least $c(y_{j_n})+\psi$. Therefore, $\overline{v}_n(y_{j_n}) \geq c(y_{j_n})+\phi$ and thus, $q_n \geq y_{j_n}$. ■

Proof of Claim 2. For clarity, we abbreviate q_n to q. Recall that $\tau_n(q) = \inf\{m > n : \tilde{p}_m \geq p_m(q)\}$ denotes the first buyer after n, who submits an offer that is acceptable to type q. For any given stage $m \geq n$, and since $p_m(q) - c(q) = \mathbf{E}_{\sigma_B^*}\left[\delta^{\tau_n(q)-m}(\tilde{p}_{\tau_n(q)} - c(q))|\tau_n(q) > m\right]$, one has

$$p_m(q) - c(q) \leq (c(1) - c(q))\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau_n(q)-m}|\tau_n(q) > m\right]. \tag{19}$$

On the other hand, $p_m(q) = \overline{v}_m(q) \geq \overline{v}_n(q)$ if $\pi_m^* = 0$, and then $p_m(q) - c(q) \geq \phi$, while $p_m(q) - c(q) \geq \delta(p_{m+1}(q) - c(q)) \geq \overline{\delta}\phi$ if $\pi_m^* > 0$, since then, $\pi_{m+1}^* = 0$. Using (19), this implies

$$\mathbf{E}_{\sigma_B^*}\left[\delta^{\tau_n(q)-m}|\tau_n(q) > m\right] \geq \frac{\overline{\delta}\phi}{c(1) - c(q)} \geq a.$$

Apply now Lemma 5 to obtain $\mathbf{P}_{\sigma_B^*}(\tau(q) \geq n + N_{a,\varepsilon}|\tau(q) \geq n) \leq \varepsilon$.

Finally, observe that

$$\begin{aligned}
\mathbf{P}_{\sigma_B^*}(\tau(q) \geq n + N_{a,\varepsilon}) &\geq \mathbf{P}_{\sigma_B^*}(\tau(q) \geq n + N_{a,\varepsilon}, t \in [\underline{q}, q]) \\
&\geq \mathbf{P}_{\sigma_B^*}(\tau(t) \geq n + N_{a,\varepsilon}, t \in [\underline{q}, q]) \\
&= F_{n+N_{a,\varepsilon}}(q),
\end{aligned}$$

whereas

$$\begin{aligned}
\mathbf{P}_{\sigma_B^*}(\tau(q) \geq n) &= \mathbf{P}_{\sigma_B^*}(\tau(q) \geq n, t \in [\underline{q}, 1]) \\
&\leq \mathbf{P}_{\sigma_B^*}(\tau(1) \geq n, t \in [\underline{q}, 1]) \\
&= F_n(1).
\end{aligned}$$

Therefore,

$$\frac{F_{n+N_{a,\varepsilon}}(q)}{F_n(1)} \leq \mathbf{P}_{\sigma_B^*}(\tau(q) \geq n + N_{a,\varepsilon}|\tau(q) \geq n) \leq \varepsilon,$$

as desired. ■

Step 2: Define $C_0 = \left\lceil \frac{2M_{c'}}{(1-\underline{q})m_{v'}} \right\rceil \times \frac{1}{a}\frac{1}{\varepsilon^2}$. We prove that $N_0 - N_1 \leq C_0/(1-\delta)$.

Denote $q_* := \max_{n<N_1} \max T_n$ the highest type offer that may be submitted before stage N_1. In particular, one has $\overline{v}_n(1) \geq \overline{v}_n(q_*) \geq c(q_*)$.

On the other hand, $N_0 = \inf\{n : \overline{v}_n(1) = c(1)\}$. Indeed, $\overline{v}_{N_0}(1) = c(1)$ since $\pi^*_{N_0} = 0$, and $\overline{v}_{N_0}(1) > \overline{v}_{N_0-1}(1)$ since buyer $N_0 - 1$ makes a serious, non-winning offer with positive probability.

Thus, $N_1 - N_0$ is bounded by the time it takes for $\overline{v}_n(1)$ to increase from $c(q_*)$ to $c(1)$. Between stages N_1 and N_0, and using the proof of Proposition 9, no buyer ever submits a serious type offer above q_*. Hence, $\overline{v}_n(1)$ increases steadily with time, at a speed which is related to the probability with which successive buyers do trade. Lemma 6 below provides a precise estimate of this relationship.

Lemma 6 *Let $n < m \leq N_0$ be any two stages, and denote by $\pi_{n,m} := \frac{F_n(1)-F_m(1)}{F_n(1)}$ the probability that the seller accepts an offer from some buyer $n, n+1, \ldots, m-1$, conditional on having declined all previous offers. Then:*

$$\overline{v}_m(1) - \overline{v}_n(1) \geq \frac{m_{v'}(1-\underline{q})}{2} \times (1-q_*)\frac{\pi_{n,m}}{F_m(1)}. \tag{20}$$

The proof of Lemma 6 is tedious and somewhat lengthy. It is postponed to the end of the section. By Lemma 5, one has, as in **Step 1**, $\pi_{n,n+N_{a,\varepsilon}} \geq \varepsilon$. By Lemma 6, one thus has

$$\begin{aligned} \overline{v}_{n+N_{a,\varepsilon}}(1) - \overline{v}_n(1) &\geq \frac{m_{v'}(1-\underline{q})}{2} \times (1-q_*)\varepsilon \\ &\geq \frac{m_{v'}(1-\underline{q})}{2} \frac{c(1)-c(q_*)}{M_{c'}}\varepsilon. \end{aligned}$$

Hence, in any block of $N_{a,\varepsilon}$ consecutive stages $k < N_0$, the average quality increases by at least $\dfrac{m_{v'}(1-\underline{q})\varepsilon}{2M_{c'}}$ times $c(1) - c(q_*)$. In particular, it takes no more than $\left\lceil \dfrac{2M_{c'}}{m_{v'}(1-\underline{q})\varepsilon} \right\rceil$ such blocks to increase from $c(q_*)$ to $c(1)$. The result follows.

Step 3: In the light of the results obtained so far, this last step is straightforward. Observe first that $p_{N_0}(\underline{q}_{N_0}) \geq v(\underline{q}_{N_0}) \geq c(\underline{q}_{N_0}) + \nu$, for otherwise buyer N_0 would get a positive payoff when submitting a type offer slightly above $\underline{q}_{N_0}$.

Since no buyer $n \geq N_0$ ever submits a serious offer below 1, one has

$$p_{N_0}(\underline{q}_{N_0}) - c(\underline{q}_{N_0}) = (c(1) - c(\underline{q}_{N_0}))\mathbf{E}_{\sigma^*_B}\left[\delta^{\tau(1)-N_0}\right],$$

which yields

$$\mathbf{E}_{\sigma^*_B}\left[\delta^{\tau(1)-N_0}\right] \geq \frac{\nu}{c(1) - c(\underline{q}_{N_0})}.$$

This concludes the proof.

Proof of Lemma 6. Fix the distribution f_n of types faced by buyer n, and the value of $\pi_{n,m}$. We minimize $\overline{v}_m(1)$ over all distributions of types that buyer m may possibly be facing. It is convenient to parameterize such distributions by $g(t)$, the probability that a seller with type t would reject all offers from buyers $n, n+1, \ldots, m-1$, so that $f_m(t) = g(t) f_n(t)$.

Hence, $\overline{v}_m(1)$ is minimal when $\int_{\underline{q}}^1 g(t) f_n(t) v(t) dt$ is minimal. The minimum is computed over all non-decreasing functions g, with values in $[0,1]$, and such that

(i) $g(t) = 1$ over $[q_*, 1]$ (since there is no serious offer beyond q_*);

(ii) $\int_{\underline{q}}^1 g(t) f_n(t) dt = (1 - \pi_{n,m}) F_n(1)$.

Since v is increasing, the minimum is obtained when g is constant over the interval $[\underline{q}, q_*]$, that is, $g(t) = \omega$ if $t < q_*$, and $g(t) = 1$ if $t \geq q_*$. The value of ω is deduced from **(ii)**, and is given by $\omega F_n(q_*) = \pi_{n,m} F_n(1)$.

Thus,

$$\overline{v}_m(1) - \overline{v}_n(1) \geq \frac{1}{F_m(1)} \int_{\underline{q}}^1 v(t) g(t) f_n(t) dt - \frac{1}{F_n(1)} \int_{\underline{q}}^1 v(t) f_n(t) dt. \tag{21}$$

The rest of the proof consists in showing that the right-hand side of (21) is at least equal to the right-hand side in (20).

Plugging g into (21), one has

$$\begin{aligned}
\overline{v}_m(1) - \overline{v}_n(1) &\geq \frac{1}{F_m(1)} \left\{ \int_{\underline{q}}^1 v(t) f_n(t) dt - \omega \int_{\underline{q}}^{q_*} v(t) f_n(t) dt \right\} - \overline{v}_n(1) \\
&= \frac{\omega F_n(q_*)}{F_m(1)} \left\{ \frac{1}{F_n(1)} \int_{\underline{q}}^1 v(t) f_n(t) dt - \frac{1}{F_n(q_*)} \int_{\underline{q}}^{q_*} v(t) f_n(t) dt \right\} \\
&= \frac{\pi_{n,m}}{F_m(1) F_n(q_*)} \left\{ F_n(q_*) \int_{\underline{q}}^1 v(t) f_n(t) dt - F_n(1) \int_{\underline{q}}^{q_*} v(t) f_n(t) dt \right\} \\
&= \frac{\pi_{n,m}}{F_m(1) F_n(q_*)} \left\{ F_n(q_*) \int_{q_*}^1 v(t) f_n(t) dt - (1 - q_*) \int_{\underline{q}}^{q_*} v(t) f_n(t) dt \right\}
\end{aligned}$$

where the first equality follows from the identity $\dfrac{a-a'}{b-b'} = \dfrac{b'}{b-b'}\left(\dfrac{a}{b} - \dfrac{a'}{b'}\right)$, the second from the value of ω, and the third from $F_n(1) = F_n(q_*) + (1-q_*)$.

We now use the inequality $v(t) \geq v(q_*) + m_{v'}(t-q_*)$ $(t \in [q_*, 1])$ to bound the first integral, and the inequality $v(t) \leq v(q_*) + m_{v'}(t-q_*)$ $(t \in [\underline{q}, q_*])$ to bound the second one. After simplification, this yields

$$\overline{v}_m(1) - \overline{v}_n(1) \geq \frac{\pi_{n,m}}{F_m(1)F_n(q_*)}(1-q_*)m_{v'}\left\{F_n(q_*)\frac{1-q_*}{2} + \int_{q_*}^{\underline{q}}(q_*-t)f_n(t)dt\right\}. \tag{22}$$

Consider finally the right-hand side of (22). For a given value of $F_n(q_*)$, the integral is minimized when f_n is constant over $[\underline{q}, q_*]$, and equal to $F_n(q_*)/(q_* - \underline{q})$. The integral is then equal to $\frac{1}{2}(q_* - \underline{q})F_n(q_*)$. Substituting into (22), this yields

$$\overline{v}_m(1) - \overline{v}_n(1) \geq \frac{\pi_{n,m}}{F_m(1)}(1-q_*)m_{v'} \times \frac{1-\underline{q}}{2},$$

as desired. ■