Public vs. Private Offers in the Market for Lemons*

Johannes Hörner and Nicolas Vieille

December 2006

Abstract

We analyze a bargaining model in which a sequence of buyers make offers to a long-lived seller endowed with a single unit for sale. Valuation and cost are correlated. We consider both the case in which previous offers are observable and the case in which they are not. When offers are observable, trade may only occur in the first period, so that bargaining typically ends up in an impasse. In the unobservable case, agreement is always eventually reached, although it does involve delay even when frictions disappear.

1 Introduction

While the sensitivity of the predictions of noncooperative game theory to the details of the game form is often perceived as a practical impediment, it has also deepened our understanding of the role and value of various concepts, such as commitment, information and observability. The purpose of this paper is to provide some insights into the role of observability. More precisely, this paper examines the role of observability in the context of bargaining with correlated values.

We investigate the impact of observability on the equilibrium outcome, most notably the probability of reaching agreement. The underlying structure is as in Akerlof's market for 'lemons'. An

*We thank Marco Ottaviani for useful comments, as well as audiences at Columbia University, UQÁM, Johns Hopkins University, London Business School, the Society of Economic Dynamics, and the 2006 Summer Meeting of the European Econometric Society, Vienna.

impatient seller is better informed than the potential buyers about the value of the transaction, but it is nevertheless common knowledge that a mutually beneficial trade exists. All potential buyers share the same valuation for the unit for sale, which is strictly larger than the seller's cost, or reservation utility.

As in the search literature, 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. The search is without recall: a seller turns to another buyer for a new offer as soon as he rejects a previous offer.

We consider both the case in which buyers observe the offers that have been rejected, so that offers are observable, or public, and the case in which buyers do not, but infer from calendar time how many have been submitted so far (the case of hidden, or private offers). Most of our results consider the case of small frictions, i.e. low but positive discounting.

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 allows them to estimate the number of previous offers that must have been declined. However, past offers 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 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, 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 will be submitted. 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 offer that is sufficiently low that the seller would gain by accepting when the quality is low enough? As we show, the problem is one of commitment from the seller's point of view. Since offers are observable, the seller would take advantage of this offer to obtain a better offer from the next potential buyer. In turn, this renders making such an offer unattractive.[1]

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 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 potential 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 submits an offer that is both higher than the seller's lowest possible reservation value 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 will accept the offer if his reservation value is low enough. This, in turn, means that the offer is a profitable deviation for the buyer.

Nevertheless, in any equilibrium with private offers, agreement is not reached immediately, and delay persists even when frictions vanish. The equilibrium is generally not unique. Although we do not characterize all possible equilibria, we show that they all involve mixed strategies and share a common structure.

This dichotomy between public and private offers may be surprising in light of the linkage principle. Indeed, the auctioneer may be better off with private offers. However, it is important to distinguish between how much information is allowed to be revealed by the information structure

[1] Academic departments are well aware of this problem when considering making senior offers. As clearly this example cannot fail any of the necessary rationality criteria, the prevalence of such offers implies that one of our assumptions must fail in this environment.

and the information that is actually revealed in equilibrium. While more 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 in the case of private offers.

Our paper endogenizes the offer distribution that is typically taken as given in the search literature. 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.[2] 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.

Our contribution is further related to three strands of thought in the literature. First, several authors have already considered dynamic versions of Akerlof's model. Second, several papers in the bargaining literature consider interdependent values. Finally, a pair of papers have investigated the difference between public and private offers in the framework of Spence's educational 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. While there are several inessential differences between their model and ours, 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).[3] 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 immediately entails that the price exceeds the valuation to the lowest quality

[2] The offer distribution also depends on the options available to the searcher, such as recall vs. no recall, as we discuss in an example.

[3] More precisely, equality obtains whenever there is a positive measure of goods' qualities traded, because there is a continuum of sellers in their model.

seller, so that trade must occur eventually, if not in a given period. Also related are Taylor (1999), Hendel and Lizzeri (1999), Blouin (2003) 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. In this environment, he shows that, with correlated values, 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. While there are gains of trade for both values in his case, it is not known whether his example admits other equilibria. 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. They show that, when the static incentive constraints preclude first-best efficiency, the limiting bargaining outcome involves agreement but delay, and fails to be second-best efficient. As their paper is the closest to ours, we shall mention and further discuss it on several occasions below. Other related contributions include Riley and Zeckhauser (1983), Cramton (1984), Gul and Sonnenschein (1988), and Vincent (1990).

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 the case of 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 shows that, in contrast to the public case, it involves pooling in the limit. While the set-up is rather different, the logic driving these results is similar to ours, at least for public offers. Indeed, in both papers, when offers are observable, firms (buyers) are deterred from submitting mutually beneficial offers because rejecting such an offer sends a strong signal to future firms (buyers) that is so attractive that only very low types would prefer to accept the offer immediately.

The general set-up is described in Section 2. Section 3 provides a simple two-buyer example that illustrates the main insights behind the results. Section 4 characterizes the equilibrium in the case of public offers. Private offers are considered in Section 5. Results and extensions are discussed in Section 6. In particular, we briefly discuss without proofs the results in the case in which buyers do not know calendar time; in which multiple buyers submit offers in every period; in which the seller may engage in cheap talk; and in which the seller, instead of the buyer, is the proposer. Most 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 reservation value of the unit is the seller's private information. This reservation value is denoted by $c(x)$, where the random variable x is determined by nature and uniformly distributed over the interval $[\underline{x}, 1]$, $\underline{x} \in [0, 1)$. We interpret x as an index, such as the quality of the good, and refer to it as the seller's *type*. The valuation (or value) of the unit to buyers is common to all of them, and is denoted by $v(x)$. Buyers do not observe the realization of x, but its distribution is common knowledge.

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

Observe that the assumption that x is uniformly distributed is made with little loss of generality, given that few restrictions are imposed on the functions v and c.[4]

We assume that gains from trade are always positive with $\nu = \inf_x \{v(x) - c(x)\} > 0$. In examples and extensions, we shall often restrict attention to the standard case in which $v(x) = x$ and $c(x) = \alpha x$, with $\underline{x} > 0$, i.e. the reservation value to the seller is a fraction $\alpha \in (0,1)$ of the valuation x to the buyers. 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 $\bar{\delta} = 1 - m/3M$, and will always assume $\delta > \bar{\delta}$. In each period in which the seller owns the unit, he derives a per-period gross surplus of $(1-\delta)\, c(x)$. Therefore, the seller can always guarantee a gross surplus of $c(x)$ by never selling the unit.

Buyer n submits an offer, or *price*, p_n that can take any real value. An outcome of the game is a triple (x, n, p_n), with the interpretation that the realized index is x, 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 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

$$\sum_{i=1}^{n-1} (1-\delta)\, \delta^{i-1} c(x) + \delta^{n-1} p_n - c(x) = \delta^{n-1} (p_n - c(x)), \tag{1}$$

when $n < \infty$, and zero otherwise. An alternative formulation that is equivalent to the one above is that the seller derives no per-period gross surplus from owning the unit, but incurs a production cost of $c(x)$ at the time he accepts the buyer's offer. It is immediate that this interpretation yields the same utility function. Accordingly, we shall often refer to the reservation value $c(x)$ as the seller's *cost*.

Buyer n's utility is $v(x) - p_n$ if the outcome is (x, n, p_n), and zero otherwise (discounting is irrelevant since buyers make only one offer). We define the players' expected utility over

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

lotteries of outcomes, or *payoff* for short, in the standard fashion. Payoff and *profit* will be used interchangeably. We allow for mixed strategies on the part of all players.

We consider both the case in which offers are public, or observable, and the case in which previous offers are not observable, also referred to as *private*, or *hidden*. It is worth pointing out that the results for the case in which offers are public would also hold for any information structure (about previous offers) in which each buyer $n > 1$ observes the price offered by 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 $[\underline{x}, 1] \times H^{n-1} \times \mathbb{R}$ in to $\{\mathtt{Accept}, \mathtt{Decline}\}$, mapping the realized type x, 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 $\mathbb{R}$.[5] 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 $\sigma_S^n(x, h^{n-1}, p_n)$ assigns a positive probability to `Accept` for some x, then $\sigma_S^n(x', h^{n-1}, p_n)$ assigns probability 1 to `Accept`, for all $x' > x$. 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 x accepting a given offer is called the *marginal type* (at history (h^{n-1}, p_n) given the strategy profile). Since the specification of the action of the seller's marginal 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 with marginal type accepts the offer. For conciseness, we shall omit to specify that 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. It follows from iterated deletion of strictly dominated strategies that buyers never submit any offer that is strictly larger than $c(1) = \bar{c}$, the highest possible reservation value

[5] 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}.

to the seller, and that any such offer would be accepted by the seller with probability 1. This argument is also standard and therefore omitted.

We use the perfect Bayesian equilibrium concept as defined in Fudenberg and Tirole (Definition 8.2).[6] 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 $\bar{c}$, 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 x is the uniform distribution over some interval $[\underline{x}_n, 1]$, where $\underline{x}_n$ may depend on the sequence of earlier offers. We do not impose any (other) restriction on beliefs.

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, a buyer's offer is *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 a 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.

Given some equilibrium, there is a one-to-one mapping between the price that is submitted and the marginal type that, by definition, is indifferent between accepting and rejecting it. As it will often be convenient to think of buyers choosing a particular marginal type, rather than a particular price, we usually use the word 'offer' to refer to the marginal type corresponding to a given price submitted by the buyer (as opposed to this price itself).

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

We briefly sketch here the static version with one buyer of the dynamic game described above. 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 $n = 1$). The model considered by Akerlof (1970) is not the static version of this game, because the market mechanism adopted there is Walrasian. Much closer is the second variant analyzed by Wilson (1980), although he considers a continuum of buyers. Clearly, the seller accepts any offer p provided $p \geq c(x)$. Therefore, the buyer offers $c(x^*)$, where x^* maximizes

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

over $x \in [\underline{x}, 1]$. More generally, given $t \in [\underline{x}, 1)$, let $x^*(t)$, or x^* when no confusion should arise, denote the marginal type given the optimal offer when the distribution is uniform over $[t, 1]$. Observe that $x^*(t) > t$ for all $t \in [\underline{x}, 1)$, and that the corresponding profit of the buyer must be positive (because the buyer can always submit a price in $(c(\underline{x}), v(\underline{x}))$. The static case with multiple bidders is discussed below.

3 A two-period Example

In this section, we provide an analysis of a simple case in which there are only two periods, and thus two buyers, in order to provide some intuition about the effect on the equilibrium outcome of the type of information available to the buyers. While some features of the equilibrium are due to the restriction to two periods, this analysis already reveals some and suggests other properties present in the general model. Given that the purpose of the example lies in its simplicity and tractability, we restrict attention to the standard case: the payoff is given by (1), with the valuation $v(x) = x$ and the cost $c(x) = \alpha x$, $\alpha \in (0, 1)$. Further, the seller's discount factor is at least $\delta > \alpha/\beta$, where $\beta := 2\alpha - 1$. Moreover, we assume that the buyer's expected value does not exceed the highest seller's cost, that is, $\underline{x} < \beta$. This means that adverse selection is severe enough to prevent first-best efficient trade (see Lemma 1 of Liang and Deneckere).

3.1 Observable Offers

Assume first that the second buyer, buyer 2, observes the offer submitted by the first buyer, buyer 1. That is, buyer 2's belief is uniformly distributed over $[\underline{x}_2, 1]$, where $\underline{x}_2$ is the marginal type corresponding the observed offer submitted by buyer 1.[7]

Since buyer 2 is the last buyer, his optimization problem is the familiar static one. The seller will accept buyer 2's price offer p if and only if his type x is such that $p \geq \alpha x$. Therefore, buyer 2's optimal 'type' offer (in terms of marginal type x) must maximize

$$\int_{\underline{x}_2}^{x} (t - \alpha x)\, dt,$$

over x in $[\underline{x}_2, 1]$. Indeed, this integral term is his payoff (up to a constant factor $(1 - \underline{x}_2)^{-1}$ ignored throughout), if all seller's types up to x accept. It is easy to see that the optimal type offer equals $\min\{1, \alpha x/\beta\}$. More precisely, if $\underline{x}_2 < \beta/\alpha$, the optimal type offer is $\alpha \underline{x}_2/\beta$, and the corresponding price is $\alpha^2 \underline{x}_2/\beta$; if instead $\underline{x}_2 \geq \beta/\alpha$, the optimal type offer is 1, and the price is α.

Buyer 1 faces a more interesting problem. If he submits a type offer $x < \beta/\alpha$, the seller's marginal type x can expect to get a price $\alpha^2 x/\beta$ in the second period, so that the price $p(x)$ buyer 1 must offer satisfies

$$p(x) - \alpha x = \delta \left(\frac{\alpha^2 x}{\beta} - \alpha x \right), \text{ or } p(x) = \left(\delta \frac{\alpha^2}{\beta} + (1 - \delta)\, \alpha \right) x,\ x < \beta/\alpha.$$

Observe that $\alpha^2 - \beta = (1 - \alpha)^2 > 0$. Therefore, if δ is large enough, the coefficient of x in $p(x)$ strictly exceeds 1. In that case, $p(x) > x$ and such offers are strictly unprofitable if $x > \underline{x}$.[8]

On the other hand, if buyer 1 submits a type offer $x \geq \beta/\alpha$, the seller's marginal type x can expect a price of α by rejecting, so that the price $p(x)$ that buyer 1 must offer satisfies

$$p(x) - \alpha x = \delta (\alpha - \alpha x), \text{ or } p(x) = \delta \alpha + (1 - \delta)\, \alpha x,\ x \geq \beta/\alpha,$$

an affine function in x. In that case, buyer 1's payoff, as a function of his type offer, is

$$\int_{\underline{x}}^{x} (t - p(x))\, dt.$$

[7] Since $\underline{x} < \beta$, it follows that $\underline{x}_2 < 1$.

[8] More precisely, $\delta\alpha^2/\beta + (1 - \delta)\,\alpha > 1$ if and only if $\delta > \alpha/\beta$, which has been assumed.

Hence, his marginal payoff is

$$x - p(x) - p'(x)(x - \underline{x}) = (1 - 2(1-\delta)\alpha)x - \alpha\delta + (1-\delta)\alpha\underline{x},$$

an affine function in x that is strictly increasing in x, since $1 - 2(1-\delta)\alpha > 0$ is precisely equivalent to $\delta > \alpha/\beta$. That is, buyer 1's payoff is a strictly *convex* function of his offer. The maximizer over the interval $[\beta/\alpha, 1]$ is therefore either β/α or 1. The former yields a strictly negative payoff, because $p(x)$ is continuous at $\beta/\alpha > \underline{x}$, yet we have seen that all serious offers below β/α are strictly unprofitable. The latter is also strictly unprofitable, since by assumption the expected value falls short of the seller's higher cost.[9] It follows that buyer 1 necessarily submits a losing offer.

We summarize this discussion in the following Proposition.

Proposition 3.1 *Assume $v(x) = x$, $c(x) = \alpha x$, $\delta > \alpha/\beta$ and $\underline{x} < \beta$. The unique equilibrium outcome in the two-period game with observable offers is such that:*

(i) the first buyer submits a losing offer with probability 1;

(ii) the second buyer submits the offer that is optimal in the static game; i.e. he offers the price αx^*, which is accepted by all types up to $x^* = \alpha\underline{x}/\beta$.

According to Proposition 3.1, the first buyer makes no profit and no sale. This may be somewhat surprising, since there are known gains from trade between the buyer and the seller. After all, why doesn't the first buyer 'preempt' the second buyer, by offering immediately what the second buyer offers in the second period? The seller would be able to collect her surplus without delay, and the first buyer would make a positive profit. The problem, of course, is that the second buyer only offers αx^* *if* the first buyer submits a losing offer. If the first buyer submitted a higher offer instead, buyer 2 would 'up the ante' by making a larger offer. Anticipating this, the seller would only accept an offer from the first buyer that is nearly as large as the second offer. Indeed, if the second buyer offers a price of α, the discount that the

[9] It follows from this analysis that, under the alternative assumption $\underline{x} > \beta$, the first buyer would make a winning offer with probability 1.

seller's type x is willing to accept in the first period is given by $\alpha - p(x) = \alpha(1-\delta)(1-x)$, which vanishes as δ tends to 1. This explains why the first buyer's payoff is strictly convex in the relevant interval. If some serious offer were profitable in that interval of types, then larger (type) offers would be even more so, because the corresponding price increase would be insignificant compared to the impact on the average quality. Yet the only offer that the second buyer would not top is a winning offer, which is necessarily unprofitable.

While the second buyer is able to take advantage of his monopsonistic situation and collect a positive payoff, we shall see in the general model that the more typical situation for a buyer is that of the first buyer. Before doing so, we shall consider the case of private, or hidden offers with two buyers.

3.2 Hidden Offers

The analysis of the case of private offers is more difficult. We shall proceed by a series of claims.

The second buyer's profit is positive. To see this, let $\underline{x}_2$ be the lowest type in the support of buyer 2's beliefs. That is, the seller accepts any equilibrium price offered by the first buyer if and only if his type is below $\underline{x}_2$. Note that $\underline{x}_2 < 1$ (as the expected value of the unit falls short of the seller's higher cost), and thus the first offer is rejected with positive probability. By offering a price in $(\alpha\underline{x}_2, \underline{x}_2)$, buyer 2 can guarantee himself a positive payoff. For future reference, observe that buyer 2's lowest equilibrium offer, denoted x_2', must strictly exceed $\underline{x}_2$.

The first buyer's profit is zero, and he submits a losing offer with positive probability. Consider the first buyer's payoff as a function of an offer $x \in (\underline{x}, x_2')$. By definition, the second buyer's offer necessarily exceeds x_2', so that the price that the first buyer should offer must solve

$$p(x) - \alpha x = \delta\left(\mathbb{E}[p_2] - \alpha x\right), \text{ or } p(x) = \delta\mathbb{E}[p_2] + (1-\delta)\alpha x,\ x \in (\underline{x}, x_2'),$$

where $\mathbb{E}[p_2]$ is the expected equilibrium price offered by the second buyer. By the same argument as above, the corresponding payoff is strictly convex in x, so that its maximum over this interval is either attained at $\underline{x}$ or at x_2'. However, we have seen that $x_2' > \underline{x}_2$, so that $\underline{x}_2 = \underline{x}$, establishing both assertions.

Both buyers' strategies are (totally) mixed. Suppose that buyer 2 follows a pure strategy, making the offer x_2' (at a price $\alpha x_2'$) with probability 1. Observe that this price would also be accepted by types $x \leq x_2'$ if it were offered by the first buyer, so that the first buyer could make a positive profit, which is a contradiction. Suppose now that buyer 1 follows a pure strategy and therefore makes the losing offer $\underline{x}$ with probability 1. Then the second buyer's optimal offer would be unique (and equal to $\alpha \underline{x}/\beta$, as in the public case), so that his strategy would be pure.

Buyer 2's strategy assigns positive probability to exactly two offers. First, if buyer 2 assigns positive probability to two distinct offers x_2' and x_2'', buyer 1 must submit some offer in (x_2', x_2''). If not, then buyer 2's payoff is a strictly concave function in x over (x_2', x_2'') (as the density of the type distribution is uniform over this interval), which is continuous over $[x_2', x_2'']$. Therefore, it cannot achieve its maximum at both endpoints simultaneously. Second, observe that buyer 1's payoff from offer x can be written as

$$(x - \underline{x}) \left(\frac{x + \underline{x}}{2} - p(x) \right), \tag{2}$$

where $p(x)$ solves $p(x) - \alpha x = \delta \left(\mathbb{E}\left[(p_2 - \alpha x) 1_{p_2 \geq \alpha x} \right] \right) = \delta \sup_{y \in \mathbf{R}} \mathbb{E}\left[(p_2 - \alpha x) 1_{p_2 \geq y} \right]$. As a supremum of affine functions, p is convex over any interval $(x_1', x_1'') \subset [\underline{x}, 1]$, and is affine over this interval if and only if buyer 2 does not submit any offer in this interval with positive probability. If buyer 1 submits both offers x_1' and x_1'' with positive probability, it must be that $p(x) \geq (x + \underline{x})/2$ with equality for both $x = x_1', x_1''$. Thus, this is only possible if buyer 2 does not submit any offer in this interval with positive probability. We now combine these two observations. Assume that buyer 2 submits three offers x_2', x_2'' and x_2''' with positive probability, with $x_2' < x_2'' < x_2'''$. Hence, buyer 1's strategy must assign positive probability to offers both in $[x_2' + \varepsilon, x_2'' - \varepsilon]$, and in $[x_2'' + \varepsilon, x_2''' - \varepsilon]$, for some $\varepsilon > 0$. In turn, this implies that buyer 2 cannot submit any offer in $(x_2'' - \varepsilon, x_2'' + \varepsilon)$, which is a contradiction.

Buyer 1's serious offers are concentrated on $[x_2', x_2'']$. We already know that buyer 1 cannot submit offers in $(\underline{x}, x_2')$, so it remains to show that he cannot do so either over $(x_2', 1)$. Over this interval, $p(x) = \alpha x$, so that his payoff written above is a strictly concave quadratic function. If some $x_1' \in (x_2', 1)$ maximized this function, it would thus have to be that both factors of (2) are zero (since both the payoff at x_1' and its derivative at x_1' would have to vanish), which is impossible, since $x_2' > \underline{x}$.

The following proposition summarizes these findings and completes the characterization of the equilibria. While the proposition describes the main necessary conditions only, the necessary and sufficient conditions can be found in the proposition's proof.

Proposition 3.2 *Assume $v(x) = x$, $c(x) = \alpha x$, $\delta > \alpha/\beta$ and $\underline{x} < \beta$. There is a continuum of equilibria. In all equilibria:*

(i) the second buyer's strategy assigns positive probability to a losing offer and some serious offers in $[x_2', x_2'']$, where $x_2' = \alpha x_2'' = x^*$;

(ii) the second buyer's strategy assigns positive probability to two serious offers, x_2' and x_2'' (with probability $1 - \beta/(2\alpha\delta)$ and $\beta/(2\alpha\delta)$ respectively).

Proof. It remains to be shown that the two offers of buyer 2, x_2' and x_2'', are indeed given by (i) $x_2' = \alpha x^* = \alpha \underline{x}/\beta$ and (ii) $x_2'' = x^* = \underline{x}/\beta$, (iii) that the probability of the latter offer is $\beta/(2\alpha\delta)$, and (iv) that there is a continuum of equilibria.

(i) For any offer x_2 in $[\underline{x}, x_2']$, the distribution over seller's type, conditional on the offer being accepted, is uniform over $[\underline{x}, x_2]$. Since x^* is the optimal offer against a uniform distribution over $[\underline{x}, 1]$, buyer 2 would thus strictly prefer offering x^* to x_2' if $x_2' > x^*$. On the other hand, suppose that $x_2' < x^*$. Since buyer 1 submits no serious offer less than x_2', the density of buyer 2's belief is constant (say, equal to μ) over this interval. Note that, if buyer 1's belief were uniform over $[\underline{x}, \underline{x} + \mu^{-1}]$, he would optimally submit offer $\tilde{x} = \min\{\underline{x} + \mu^{-1}, x^*\} > x_2'$. Observe that, given his actual belief, the average quality conditional on winning with an offer $\tilde{x}$ is larger than with the uniform belief, and that the probability of the seller's accepting such an offer is also higher than with the uniform belief. Therefore, the offer $\tilde{x}$ is strictly more profitable with buyer 2's actual belief than with the belief that is uniformly distributed over $[\underline{x}, \underline{x} + \mu^{-1}]$, while the offer x_2' is equally profitable under either belief. Since, with the uniform distribution, the offer $\tilde{x}$ is strictly more profitable than the offer x_2', the same must be true with buyer 2's actual belief, which contradicts the optimality of x_2'.

(ii) and (iii) The price $p(x)$ that buyer 1 must offer to $x \in (x_2', x_2'')$ is given by

$$p_1(x) = \alpha x + \delta\lambda_2(\alpha x_2'' - \alpha x), \tag{3}$$

where λ_2 is the probability assigned by buyer 2 to offer x_2''. Since buyer 1 submits offers in $[x_2', x_2'']$, it must be that $p(x) \geq (x + \underline{x})/2$ for all $x \in [x_2', x_2'']$, with equality for some $x \in [x_2', x_2'']$. Since $p(\cdot)$ is affine, the equality $p(x) = (x + \underline{x})/2$ must hold throughout the interval $[x_2', x_2'']$. Identification with (3) yields

$$\alpha(1 - \delta\lambda_2) = \frac{1}{2} \text{ and } \frac{\underline{x}}{2} = \delta\alpha\lambda_2 x_2'',$$

which gives the result.

(iv) Observe that, if buyer 2 randomizes according to Proposition 3.2, buyer 1's strategy is optimal if and only if his offers are either losing or in the interval $[x_2', x_2'']$. Therefore, to characterize all equilibria, we must provide the necessary and sufficient condition guaranteeing that the offers x_2' and x_2'' are optimal for buyer 2.

Let G_1 be the cumulative distribution function of buyer 1's offer, and define F_2 by $F_2(x) = \frac{1}{1-\underline{x}} \int_{\underline{x}}^{x} G_1(t)dt$. $F_2(x)$ it the (unconditional) probability that the seller is of type $t \leq x$, and has rejected buyer 1's offer. Thus, buyer 2's payoff when offering x is equal to

$$\pi_2(x) = \int_{\underline{x}}^{x} (t - \alpha x) dF_2(t).$$

Since $x_2' = x^*$, the offer x_2' is more profitable than lower offers. Thus, buyer 2's strategy is optimal if and only if

$$\pi_2(x_2') = \pi_2(x_2'') \text{ and } \pi_2(x) \leq \pi_2(x_2'), \text{ for all } x \geq x_2'.$$

Integration by parts yields that

$$\pi_2(x) = \pi_2(x_2') + (1-\alpha)xF_2(x) - (x_2' - \alpha x) F_2(x_2') - \int_{x_2'}^{x} F_2(t)dt.$$

The function F_2 is convex, nondecreasing and continuous over $[x_2', 1]$, with $F_2(x_2') \leq (x_2' - \underline{x})/(1-\underline{x})$ and $F_2(x) = F_2(x_2'') + (x - x_2'')/(1-\underline{x})$ for $x \geq x_2''$. Conversely, any such positive function is associated with some strategy of buyer 1. Among these functions, a necessary and sufficient equilibrium condition is thus

$$(1-\alpha)xF_2(x) - (x_2' - \alpha x)\,F_2(x_2') \leq \int_{x_2'}^{x} F_2(t)dt$$

for all $x \geq x_2'$, with equality for $x = x_2''$. Plainly, there is a continuum of such strategies. ■

According to Proposition 3.2, there is a continuum of equilibria. In fact, the support of buyer 1's offers may be finite or not.[10] As the careful reader may have noticed, the multiplicity in this example is closely tied to the linearity that is common to the valuation and the cost function (that is, player 1 would randomize over two offers exactly in general). While this is correct, we shall see that multiple equilibria are no longer 'nongeneric' in the general model with private offers.

Despite this multiplicity, the comparison between both scenarios is now easy to make. Buyer 1 is indifferent between both cases, because his payoff is zero in either one. Buyer 2 prefers public offers, because in both cases he submits the offer x^* at a price αx^* with positive probability, but this offer is more likely to be accepted with public offers. The seller prefers private offers, because buyer 2 offers a price of at least αx^* in the private case, and possibly more. Finally, the social surplus is also larger in the case of private offers, because the probability of trade is larger and trade occurs no later with private offers, relative to public offers.

As we shall see, some but not all of these properties hold true in the general model. With public offers, buyers get deterred from submitting a serious offer, because the seller can expect a much higher offer in the following period by rejecting it. As a consequence, buyers' payoffs are low, and so is the probability of trade. With private offers, buyers' payoffs may not be large, but they typically submit serious offers, which results in a higher probability of trade.

While a detailed analysis of the model with recall is beyond the scope of this paper, it is worth mentioning at this point that the assumption of no recall does affect the equilibrium strategies in

[10] Among these equilibria, there is a unique one in which buyer 1 submits exactly one serious offer with positive probability. In this equilibrium, buyer 1 submits offer $x_1 = \underline{x}\left(1+\frac{1-\alpha}{\alpha\lambda_1}\right)$ with probability $\lambda_1 = \left(\sqrt{\alpha\,(4+\alpha)}+\alpha-2\right)/\,(2\alpha)$.

the case of hidden offers. In the simple two-period example developed here, the equilibrium with recall also involves randomization by both bidders. In this case as well, the first bidder's payoff is zero, while the second bidder's payoff is positive. However, the support of the bid distributions is different in the two cases, and so is the expected revenue of the seller. Perhaps surprisingly, in specific examples, the seller is better off without recall, primarily because the second seller is more likely to submit the lower offer with recall than without it. This contrasts with the obvious ranking in the standard models of search in which the distribution of offers is exogenous, and suggests here as well that endogenizing this distribution may overturn familiar conclusions. By contrast, in the case of observable offers, in which the equilibrium involves pure strategies, both in the simple example and in the general case, the results are identical with or without recall.

4 General Case: Observable Offers

We return now to the general model described in Section 2, in which there are countably many periods, so that there is no 'last' buyer as in Section 3. Throughout this section, we maintain the assumption that offers are public.

Given x, observe that there exists at most one value of $y \in [\underline{x}, x]$ solving

$$\int_y^x (v(t) - c(x))\, dt = 0.$$

That is, if it exists, type y is such that the expected value to the buyer over all the types in $[y, x]$ is equal to the highest cost over types in this interval, $c(x)$. See Figure 1(a) for an illustration in the case of $c(x) = \alpha v(x) = x$, $\alpha > 1/2$, in which the inverse f^{-1} of f is well-defined. The value y exists if and only if adverse selection is severe enough that the buyer's expected value does not exceed the highest seller's cost. If it exists, y is unique because v is strictly increasing.

Let $f(x) = y \in [\underline{x}, x]$ denote this value, as a function of x, whenever it exists. Because $\inf_x \{v(x) - c(x)\} > 0$, it must be that $x - y > \kappa$ for some $\kappa > 0$, for all x. Observe that the mapping f plays a key role in the analysis of the static model of adverse selection between one seller and many (more precisely, two or more) buyers. Following logic *à la Bertrand*, any equilibrium outcome of this static version involves two or more buyers making an offer x such that $\underline{x} = f(x)$ at the price $c(x)$ (as long as $f(1)$ is defined; otherwise, they offer $c(1)$). Note that any such offer x strictly exceeds x^*, the optimal offer that a lone buyer would submit.

Define the strictly decreasing sequence x_k as follows: $x_0 = 1$, $x_{k+1} = f(x_k)$, as long as $f(x_k)$ is defined. Since $x - f(x) > \kappa$, this sequence must be finite, and we denote the smallest element of this sequence by x_K.[11] For instance, if $c(x) = \alpha v(x) = \alpha x$, with $\alpha > 1/2$, it is easy to verify that $x_k = \beta^k$, $\beta = 2\alpha - 1$.

The sequence $\{x_k\}$ plays an important role in Proposition 4.1. While it is possible that $x_K = \underline{x}$, this proposition is stated here for simplicity for the generic case in which $x_K > \underline{x}$.

Proposition 4.1 *Assume that $x_K > \underline{x}$, and $\delta > \bar{\delta}$. There is a unique equilibrium outcome, which is independent of δ. On the equilibrium path, the first buyer submits the price $c(x_K)$, which the seller accepts if and only if $x \leq x_K$. If this price is rejected, all buyers $n > 1$ submit a losing offer.*

Before turning to the proof, it is worth offering an intuition for why all buyers but the first necessarily submit losing offers. The logic here is similar to the one that which we applied to the first buyer in the two-period example. Consider the problem faced by buyer $n > 1$. While both this buyer and the seller know that there are gains from trade, they are unable to reach an agreement. The reason is that, while buyer $n + 1$ submits a losing offer on the equilibrium path, he only does so as long as buyer n does as well. If buyer n were to submit a serious offer $x \in (x_k, x_{k-1})$, for some $k \leq K$, buyer $n + 1$'s equilibrium action would actually be to offer the price $c(x_{k-1})$, accepted by the seller if and only if his type is below x_{k-1}. Later buyers would then all submit losing offers. That is, if buyer n were to submit such a serious offer x, play would proceed as if x were the initial value $\underline{x}$.

Assuming that buyers $n + 1, \ldots$ follow such strategies, why is it not profitable for buyer n to submit a serious offer? First of all, by definition of $x_K = f(x_{K-1})$, it would be unprofitable for this buyer to submit an offer x equal to or larger than x_{K-1}, because the corresponding

[11] Not surprisingly, this sequence also plays a key role in Deneckere and Liang's analysis, where it is denoted by q_n^*.

price would need to be at least as large as $c(x)$. Therefore, only offers in (x_K, x_{K-1}) need to be considered. Consider Figure 1(b). If future buyers were submitting losing offers, the price that buyer n would need to submit would, indeed, be $c(x)$, which is less than $v(x)$, so that such an offer would be profitable. But buyer $n+1$ does not submit a losing offer. Given his response, the price that buyer n must offer solves

$$p(x) - c(x) = \delta\left(c(x_{K-1}) - c(x)\right),$$

that is, $p(x) = c(x_{K-1}) - (1-\delta)\left(c(x_{K-1}) - c(x)\right)$. This function is nearly 'flat' when the buyer is sufficiently patient. Thus, the price is almost insensitive to the offer, so that the payoff of buyer n is strictly convex in his offer over this range. Either it is profitable to make the offer x_{K-1}, or no serious offer is profitable. But by definition of x_{K-1}, such an offer is not profitable.

This reasoning explains why it is optimal for buyers to follow the equilibrium strategies and submit losing offers, but it does not explain why this is the unique equilibrium outcome. This requires more work, and relies on backward induction over types, starting from the highest type. The actual method is described at the beginning of the following proof.

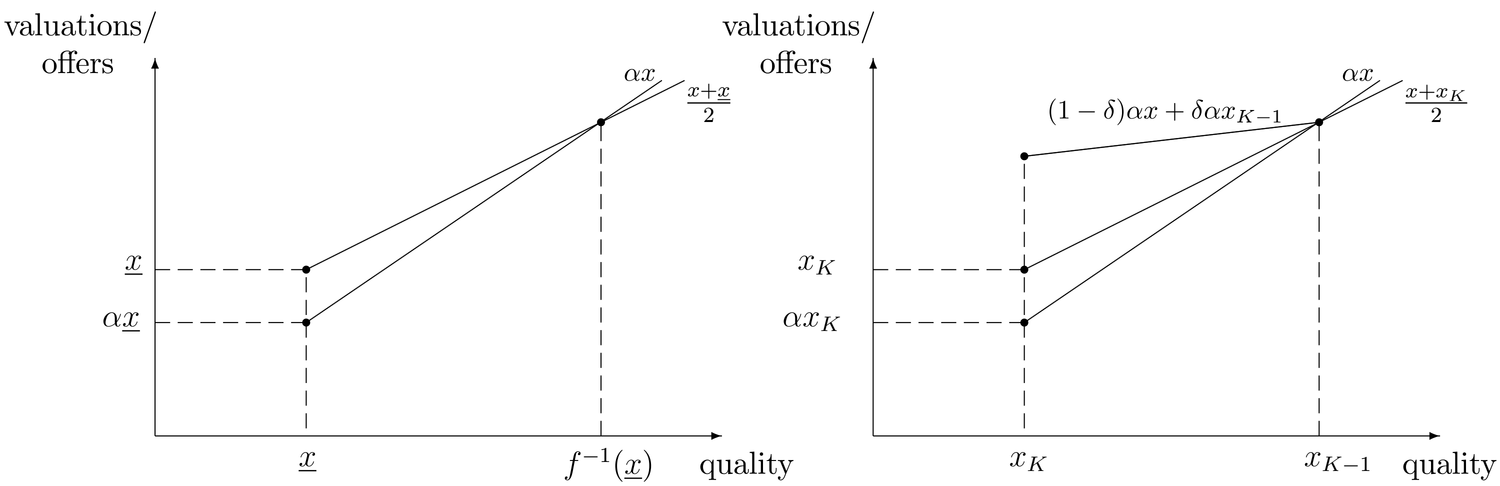

Figure 1(a) $(v(x) = x, c(x) = \alpha x)$

Figure 1(b) $(v(x) = x, c(x) = \alpha x)$

Proof of (4.1): We first show that, if there exists some equilibrium, it must have a simple structure. 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. Fix an equilibrium and assume that, for every $n \geq 1$ and after any history h^{n-1} on the equilibrium path such that $\underline{x}_n = x$, buyer n submits the price $c(x_l)$ whenever $x \in (x_{l+1}, x_l]$ for some $l < k$. We now prove that the same conclusion holds for $l = k$. The proof is broken into the following four steps:

- whenever $\underline{x}_n = x \in (x_{k+1}, x_k)$, no equilibrium offer of buyer n is accepted by some type $s > x_k$;
- whenever $\underline{x}_n = x \in (x_{k+1}, x_k]$, if an equilibrium offer of buyer n is accepted by $s = x_k$, then all subsequent offers are losing ones; in addition, if $\underline{x}_n = x_k$, the unique equilibrium price of buyer n is $c(x_k)$;
- whenever $\underline{x}_n = x \in (x_{k+1}, x_k]$ is close enough to x_k, the unique possible equilibrium price of buyer n is $c(x_k)$, which the seller accepts if and only his type is at most x_k;
- whenever $\underline{x}_n = x \in (x_{k+1}, x_k]$, the unique equilibrium price of buyer n is $c(x_k)$, which the seller accepts if and only his type is at most x_k.

Step 1: If buyer n submits a price $p(s)$ with marginal type $s \in (x_{l+1}, x_l]$ for some $l < k$, the following price is $c(x_l)$ by the induction hypothesis. Hence, $p(s)$ must solve

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

so that buyer n's payoff is

$$\frac{1}{1-x} \int_x^s (v(t) - \delta c(x_l) - (1-\delta)c(s))\, dt.$$

As a function of s, the integral is twice differentiable over the interval $(x_{l+1}, x_l]$, with first and second derivatives given by

$$v(s) - \delta c(x_l) - (1-\delta)c(s) - (1-\delta)c'(s)(s-x),$$

and

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

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

Step 2: We argue by contradiction. We thus assume that, for some n and h^{n-1} with $\underline{x}_n = x \in (x_{k+1}, x_k)$, there is a positive probability that an equilibrium price p_n by buyer n with marginal type x_k is eventually followed by a serious offer. This implies $p_n > c(x_k)$. Let $\bar{p}$ be the supremum of all such prices (with marginal type x_k), where the supremum is taken over all n and h^{n-1}.

Given any buyer n and history h^{n-1}, note that the price $p_n(x)$ with marginal type $x \leq x_k$ does not exceed $\bar{p}$. Indeed, denoting the supremum over all such prices by p^*, and denoting the first buyer submitting an offer that type x accepts by $\tau(x)$, one has $p_n(x) - c(x) = \mathbf{E}\left[\delta^{\tau(x)}(p_{\tau(x)} - c(x)) | \tau(x) > n\right] \leq \delta\left(\max\left(\bar{p}, p^*\right) - c(x)\right)$, hence $p^* \leq \bar{p}$.

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

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

By choosing s close enough to x_k, buyer n's payoff, $\dfrac{1}{1-x}\displaystyle\int_x^s \{v(t) - p(s)\}\, dt$, is thus higher than the equilibrium payoff, $\dfrac{1}{1-x}\displaystyle\int_x^{x_k} \{v(t) - p_n\}\, dt$, which is a contradiction.

We turn to the second assertion, and let n and $h^{n-1} \in H^{n-1}$ be given, with $\underline{x}_n = x_k$. Since $\underline{x} < x_K$, there must exist, along h^{n-1}, a buyer who submitted a serious offer with marginal type x_k. As we just proved, any such offer is necessarily followed by losing offers. In particular, buyer n's equilibrium price is $c(x_k)$.

Step 3: Let n and $h^{n-1} \in H^{n-1}$ be given, with $\underline{x}_n = x < x_k$. Consider a potential price $p(s)$, with marginal 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 x_k. Hence, buyer n's payoff, $\dfrac{1}{1-x}\displaystyle\int_x^s \{v(t) - p(s)\}\, dt$, is at most $\dfrac{1}{1-x}\displaystyle\int_x^s \{v(t)-c(s)\}\,dt$, and the difference converges to zero, as s increases to x_k. The latter integral, as a function of s, is differentiable, with derivative $v(s)-c(s)-c'(s)(s-x)$, which is positive as soon as $s-x<\frac{m}{M_{c'}}$. Thus, for x close enough to x_k, the upper bound, $\dfrac{1}{1-x}\displaystyle\int_x^s \{v(t)-c(s)\}\,dt$, is increasing over $[x,x_k)$. Hence, for such x, buyer n's only possible equilibrium price is $c(x_k)$ (assuming that some equilibrium exists).

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

Consider now any n and h^{n-1} with $\underline{x} = x < \tilde{x}$. By definition of $\tilde{x}$, any price $p(s)$ with marginal type $s > \tilde{x}$ is followed by a price $c(x_k)$ from the next buyer, so that $p(s)$ must satisfy

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

and buyer n's payoff writes

$$\frac{1}{1-x}\int_x^s \{v(t)-\delta c(x_k)-(1-\delta)c(s)\}\,dt.$$

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

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

It is now possible to draw a comparison between the dynamic version with public offers and the static version (at least in the generic case in which $x_K > \underline{x}$). Observe that, depending on the exact value of $\underline{x}$, x_K could be anywhere in the interval $(\underline{x}, x_{K-1})$, so that both $x_K > x^*$ and $x_K \leq x^*$ may occur. 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 $\underline{x}$. However, a benchmark that is arguably most natural is the static version with two or more buyers, because the dynamic version involves more than one buyer. The comparison is then immediate, as the offer in the static offer must be at least as large as x_K. 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, i.e. 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 $c(x)$, where x is the largest root of $f(x) = \underline{x}$.

In light of the existing literature on bargaining with correlated values, the result that trade does not necessarily occur is rather surprising. The most recent contribution to this literature, Deneckere and Liang (2006), finds that such trade occurs with delay whenever (in our notation) $\underline{x} < x_1$, but that it occurs nevertheless with probability 1. The formal difference is that Deneckere and Liang consider the case of a single long-run buyer, rather than a sequence of short-run buyers. It is worth pointing out that Proposition 4.1 remains valid in the case of a single long-run buyer, provided he is sufficiently impatient. 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 that was already hinted at by Evans (1989) in the case of binary values.

Just as Proposition 4.1 remains valid with a single impatient buyer, 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. In either case, because the buyer discounts sufficiently the possible surplus from meeting the seller again, he either makes an aggressive offer, or makes none at all. Because the seller is patient, and a serious offer submitted by all but the first buyer triggers another serious offer in turn, such aggressive offers turn out to be unprofitable.

The positive 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 the underlying uncertainty is progressively eroded. Our negative result points to another familiar force in dynamic games; namely, the absence of commitment. Indeed, if the horizon were finite, as in the two-period example of the previous section, 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. It is then natural to wonder whether the result hinges on the buyer, rather than the seller, making the offer. As we shall discuss in Section 6, the real issue is not who has the initiative, but rather who has the last word in each period. We interpret the inability to commit as meaning that the seller has the last move in each period, deciding *in fine* whether to accept the outstanding offer, however it came to be, or to turn to another buyer.

To conclude this section, we comment briefly on the knife-edge case in which $\underline{x} = x_K$. Then, as long as the marginal type is $\underline{x}$, any randomization over the offers $\{x_K, x_{K-1}\}$ is optimal, the payoff of either offer being zero. Because $\underline{x} = x_K$, equilibrium considerations do not uniquely 'pin down' the mixture, as is done in the proof above for the case $\underline{x} < x_K$ in which the marginal type is x_k, $k \leq K$, after an equilibrium offer that is serious. Indeed, the only reason why the equilibrium (as opposed to the equilibrium outcome) for the case $\underline{x} < x_K$ is not unique is that nothing pins down the behavior when the marginal type is x_k, $k \leq K$, *following an out-of-equilibrium offer*. Beyond this indeterminacy, the case $\underline{x} = x_K$ is identical to the case $\underline{x} < x_K$. In particular, along the equilibrium path, the seller rejects all offers provided $t \geq x_{K-1}$.

5 General case: Hidden Offers

As in the previous section, we consider the general set-up of Section 2, in which there are countably many periods. Unlike in the previous section, we now assume that offers are unobservable.

5.1 General Properties

As we are unable to explicitly construct equilibria in general, we first argue that an equilibrium exists. If no later buyer sets a price exceeding $\bar{c}$, 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 mixed (or behavior) strategies to the set $\mathcal{M}([0, \bar{c}])$ of probability distributions over the interval $[0, \bar{c}]$, endowed with the weak-$\star$ topology. The set of strategy profiles is thus the countable product $\mathcal{M}([0, \bar{c}])^{\mathbb{N}}$. 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 multilinear extension of the payoff induced by pure profiles. We, however, follow the standard proof. Let any buyer n be given, and let $x(p, \mu)$ denote the marginal type for the offer p, given a strategy profile μ. It is jointly continuous in p and μ. Hence, the set $B_n(\mu) \subset \mathcal{M}([0, \bar{c}])$ of best replies of buyer n to the strategy profile μ is convex-valued and upper hemi-continuous in μ. 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' private histories are on the equilibrium path until trade occurs with probability 1, if ever.

While an equilibrium always exists, it need not be unique. To give a concrete example, if $c(x) = \alpha v(x) = \alpha x$, with $\alpha = 3/4$ and $\delta = 3/4$, and $\underline{x} = .4249$, it can be shown that the following two equilibria, and probably more, exist (details available from the authors).

- The first buyer offers $\beta = 1/2$ for sure, at a price $p_1 \simeq .46$. The second buyer makes a losing offer. Buyer 3 and later buyers randomize between a winning and a losing offer, offering the winning price 3/4 with probability 3/20, and a losing offer with complementary probability.

- The first buyer randomizes between a losing offer and the offer $x_1 \simeq .51$, at a price $p_1 \simeq .47$, assigning a probability $\mu_1 \simeq .40$ to the losing offer. The second buyer randomizes between a losing offer and the offer $x_2 \simeq .62$, at a price $p_2 \simeq .54$, assigning a probability $\mu_2 \simeq .76$ to the losing offer. Buyer 3 and later buyers randomize 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 5.1 *An equilibrium exists. For some parameters, the equilibrium is not unique.*

In fact, as we shall see, the equilibrium is unique if and only if $\underline{x} > x_1$, as defined in Section 4.

The two equilibria described above are quantitatively very different. They differ in terms of expected delay, revenue, and buyers' profit (the first buyer's profit is positive in the first equilibrium). Nevertheless, there are some qualitative similarities. Both equilibria involve mixed strategies. More importantly, in both equilibria, trade occurs with probability 1. This is no coincidence: as the following proposition shows, all equilibria necessarily involve eventual agreement.

Proposition 5.2 *In all equilibria, trade occurs with probability 1.*

Proof: Fix some equilibrium. Given $x \in [\underline{x}, 1]$, let $F_n(x)$ denote the (unconditional) probability that the seller is of type $t \leq x$ and has rejected all offers submitted by buyers $i = 1, \ldots, n-1$. Suppose, for the sake of contradiction, that trade does not occur with probability 1 eventually, i.e. $\lim_{n\to\infty} F_n(x) \neq 0$ for some $x < 1$. In particular, the probability that the seller will accept buyer n's offer, conditional on having rejected the previous ones, converges to zero as n increases. Hence, the successive buyers' equilibrium payoffs also converge to zero.

Let $F = \lim_{n\to\infty} F_n$. Choose x such that $F(x) > 0$ and

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

Note that $\dfrac{F(x) - F_n(x)}{F_n(x)}$ is the probability that type x will accept an offer from some buyer beyond n (conditional on having rejected all previous offers). Since $F(x) > 0$, this probability converges to zero, and the offer $p_n(s)$ with marginal type s thus converges to $c(x)$. As a result, $p_n(s) \leq c(s) + \frac{\nu}{2}$ for all n large and, using (4), buyer n's equilibrium payoff is bounded away from zero, which is a contradiction. ■

Observe that Proposition 5.2 holds independently of δ. This proposition establishes that offers arbitrarily close to 1 are eventually submitted. We shall show later that, if $\delta > \bar{\delta}$, a stronger result holds: a winning offer is submitted with probability 1 in finite time.

It may be instructive to understand why the unique equilibrium outcome in the game with observable offers is no longer an equilibrium outcome with unobservable offers. After all, given that in the former equilibrium buyers use pure strategies, it may seem that whether their actions are observed or simply inferred in equilibrium should make no difference. The difference, of course, lies in what happens after a deviation. With public offers, buyer $n+1$ only submits a losing offer if buyer n does so as well. If, instead, buyer n deviates and submits a serious offer, then so would buyer $n+1$. With unobservable offers, it is no longer possible for buyer $n+1$ to submit different offers as a function of buyer n's offer. In turn, this implies that, if buyer $n+1$ and later buyers were only submitting losing offers in the game with unobservable offers, buyer n could make a profit by offering a price just above the lowest possible seller's cost, because the seller could not reject it for the sole reason of initiating an aggressive offer from the next buyer.

According to Proposition 5.2, agreement is always reached in finite time. It is natural to wonder how fast agreement is reached, because this is directly related to the efficiency of the bargaining outcome. 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 is necessarily reached.

Proposition 5.3 *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 result, and the proofs of all remaining results in the section, can be found in the appendix. Delay ($c_2 < 1$) should not come as a surprise. Since the seller can always wait until the first winning offer is submitted, and serious offers until then must yield a profit to the buyers submitting them, real delay must make waiting until the winning offer is made a costly alternative to the sellers's lower types. Slightly less obvious is the second conclusion; namely, that the cost of delay remains finite ($c_1 > 0$).

In the first example of an equilibrium given in this section, the first buyer enjoys a profit, but all other buyers make zero profit. More complicated examples of equilibria can be constructed in which more than one buyer makes a profit, although it may be true that there always exists some equilibrium in which all buyers' profits are zero. However, in all equilibria, all buyers' profits are small, as formalized in the following proposition.

Proposition 5.4 *There exists a constant $M > 0$ such that, for every $\delta > \bar{\delta}$ and every equilibrium, the profit of buyer n is at most*

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

Observe that Proposition 5.4 implies that the buyers' *aggregate* profits also converge to zero as the friction disappears, as each buyer's profit converges to zero at the rate $(1-\delta)^2$.

According to the next proposition, buyers that make a profit are infrequent, in the sense that two buyers who make a profit must be sufficiently far apart in the sequence of buyers.

Proposition 5.5 *There exists a constant $M > 0$ such that, for every $\delta \geq \bar{\delta}$ and every equilibrium, the following holds. If buyers n_1 and buyers n_2 have a positive equilibrium profit, then*

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

In fact, as we will see, there is an upper bound on the number of buyers with positive profit that is independent of both the equilibrium and the discount factor.

5.2 Equilibrium Strategies

As mentioned at the beginning of this section, we do not provide an explicit characterization of an equilibrium for general parameters. Nevertheless, as the examples described above suggest, all equilibrium strategies share common features. To describe these features, further notation must be introduced. Given some equilibrium, let $F_n(x)$ denote the (unconditional) probability that the seller's type t is less than or equal to x and that the seller has rejected all offers submitted by buyers $i = 1, \ldots, n-1$. Set $\underline{x}_n = \inf\{x : F_n(x) > 0\}$. Buyer n's strategy is a probability distribution over prices in $[c(\underline{x}), \bar{c}]$ that he offers. We denote by P_n the support[12]of this distribution and by T_n the corresponding (closed) set of marginal types. That is, if buyer n's strategy has finite support, $x \in T_n$ if and only if it is an equilibrium action for buyer n to submit some price p_n that a seller with type t accepts if and only if $t \leq x$.

The following proposition complements Proposition 5.2, because it shows that a winning offer is eventually submitted. It also complements Proposition 5.5, because together they imply that the number of buyers that enjoy a profit is bounded above, uniformly in the discount factor and the equilibrium.

Proposition 5.6 *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 $M > 0$ such that, in all equilibria, $N_0 \leq M/(1-\delta)$. Further, given some equilibrium, $T_n \subset \{\underline{x}_{N_0}, 1\}$ for all $n \geq N_0$. For all $n > N_0$, buyer n's equilibrium payoff is zero.*

Thus, from period N_0 on, buyers only make winning or losing offers, and all but the first of these make zero profit. In fact, it follows readily from the proof that buyer N_0 makes zero profit as well, as long as $\underline{x} \leq x_1$. There may be several values of N_0 that are consistent with the statement of Proposition 5.6. We choose N_0 to be the smallest of these. Because buyer n does not submit an offer above x, the largest root of $\underline{x}_n = f(x)$, it follows that $N_0 \geq K - 1$, where K is defined by x_K, as in Section 4.

It is convenient to discuss here the case in which $\underline{x} > x_1$. In this case, observe that, provided that he is called upon to submit an offer, any buyer is guaranteed a profit, because he can always

[12] That is, the smallest closed set with probability one.

offer the price $\bar{c}$. Hence, it follows from Proposition 5.6 that the unique equilibrium outcome of the game is such that the first buyer offers $\bar{c}$, which the seller accepts. Therefore, if $\underline{x} > x_1$, the equilibrium is unique. The next proposition shows that, in all other cases, there exists 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 finite but unbounded time.

Proposition 5.7 *Assume that $\delta > \bar{\delta}$ and $\underline{x} \leq x_1$. For every equilibrium σ:*

- *There is an equilibrium $\tilde{\sigma}$, with $\tilde{\sigma}_B^n = \sigma_B^n$ for all $n < N_0$, and $\underline{x}_n = 1$ for some $n \geq N_0$;*
- *There is an equilibrium $\tilde{\sigma}$, with $\tilde{\sigma}_B^n = \sigma_B^n$ for all $n < N_0$, and $\underline{x}_n = \underline{x}_{N_0}$ for all $n \geq N_0$.*

Moreover, infinitely many such equilibria exist. While the equilibria exhibited in the proof of the previous proposition are payoff-equivalent, for the seller and the buyers alike, the first example in this section shows that this need not be true in general.

For $\underline{x} < x_1$, the next proposition formalizes the idea that all equilibria involve mixed strategies.

Proposition 5.8 *Assume that $\delta > \bar{\delta}$ and $\underline{x} \leq x_1$. 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, under suitable convexity assumptions, each buyer's strategy is a distribution with finite support, so that each buyer randomizes over finitely many offers only.

Proposition 5.9 *Assume that $\delta > \bar{\delta}$ and $\underline{x} \leq x_1$. Further, assume that v is concave and c is convex over $(\underline{x}, 1)$, with either v or c being strictly so. Then, for any equilibrium σ, the support of the probability distribution σ_B^n is a finite set, for every buyer $n \geq 1$.*

Together, Propositions 5.6 to 5.9 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 profit, while all others have zero profit; in fact, it can be shown that 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 make zero profit, and randomize between the winning offer and a losing offer, with relative probabilities that are to a large extent free variables. Thus, as long as equilibrium offers are rejected, the expected value of the unit strictly increases over time until period N_0, and remains constant thereafter.

Because we have not ruled out the existence of an equilibrium in which all buyers make zero profit, it is tempting to investigate the existence of such equilibria, in which all buyers' strategies, except possibly the first buyer's strategy, assign positive probability to exactly two offers; a losing offer, and a serious, and eventually winning, offer.[13] The second example of an equilibrium given at the beginning of this section belongs to this family. Unfortunately, numerical examination suggests that even in the special case in which $c(x) = \alpha v(x) = \alpha x$, such equilibria only exist if $\underline{x}$ is close enough to x_1. This suggests 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, so that the lower end of the support of the buyers' belief increases over time. It is numerically possible to construct equilibria of the second kind for some parameter configurations, but showing whether such equilibria always exist appears to be an intractable problem.

6 Discussion and Extensions

Comparison between observable and hidden offers

[13] This can be shown to be equivalent to the seemingly weaker statement that each buyer's strategy assigns a positive probability to a losing offer and that the expected offer is increasing over time.

The striking difference between the two scenarios lies in the probability of agreement. With observable offers, this probability may be arbitrarily small (if $\underline{x}$ is close to x_K), and falls short of 1 whenever $\underline{x} < x_1$. (The set of equilibria coincide in the case $\underline{x} \geq x_1$, a case that we disregard in what follows). On the other hand, agreement always obtains eventually when offers are unobservable.

Other comparisons are less clear-cut. Since it is possible that $x^* = x_K$, it follows from Samuelson's (1984) Proposition 1 that the equilibrium outcome with public offers is the preferred one among the outcomes of all bargaining procedures from the point of view of the first buyer. In particular, since eventual agreement in the unobservable case implies that serious (but not winning) offers necessarily involve prices higher than the cost of the corresponding marginal type, the first buyer strictly prefers the outcome of the game with public offers to the outcome in the game with private outcome whenever x^* happens to be sufficiently close to x_K. The same argument applies to the *aggregate* profit of the buyers. As for buyers $n \geq 2$, they weakly prefer the outcome with hidden offers, although any difference disappears as frictions vanish (see Proposition 5.4).

From the seller's point of view, our first example of an equilibrium in Section 5 is clearly preferred to the corresponding outcome with observable offers by all types of the seller, so that this equilibrium outcome with unobservable offers is *ex ante* more efficient than the unique outcome with observable offers. We have not found any example in which this conclusion would be reversed. However, it is straightforward to show that no equilibrium is second-best efficient (the proof is available from the authors).

In terms of interim efficiency, the comparison can go either way. Considering the second example in Section 5, 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.

Voluntary information disclosure

These results suggest that the seller may prefer that the offers remain hidden. In many economic environments, this may be difficult to achieve (for instance, many types of takeover bids are public by design). Aside from feasibility, there is the issue of commitment. If the seller cannot commit to not revealing offers, the equilibrium outcome with public offers remains an equilibrium in the game in which the seller chooses, in every period, whether to make public the bid that was just submitted, as a function of this bid. As the argument is standard, we only provide a sketch. Whenever the seller does not reveal the offer in period n, later buyers assign probability 1 to buyer n having submitted the lowest offer within the support of his equilibrium strategy. Therefore, buyers' beliefs are always uniform over some interval of types, as in the public case, and the equilibrium play then proceeds as with observable offers. Given this, the seller has no choice but to reveal the offer he has received, as he can only (weakly) gain from doing so. On the other hand, without further refinement, this game also admits as equilibria all the equilibria from the game with hidden offers. For this, it suffices to specify buyers' beliefs such that, if the seller deviates and reveals an offer, say offer x, then all future buyers assign probability 1 to the seller being of type x.

Messages by the seller

Failure to reach agreement is often due to lack of communication. Since bargaining with public offers reaches an impasse almost immediately, it is then natural to ask whether allowing the seller to send messages (signalling his valuation, suggesting an agreeable price, etc.) would help to break the deadlock. While an analysis of all possible bargaining procedures in this framework lies outside of the scope of this paper, we comment here on probably the simplest such extension. In each period n, the seller makes an announcement from some set of messages. After the announcement, buyer n submits an offer, which the seller then accepts or rejects.

Given the focus on observability, it is natural to assume that, when offers are observable, so are the seller's messages: that is, all future buyers observe all the messages sent by the seller in all past periods. We restrict attention to the case in which buyers follow pure strategies (such an equilibrium exists). Fixing any such equilibrium, we can construct a realization-equivalent equilibrium in which offers are independent of all messages. Indeed, we claim that any serious offer submitted in period n is independent of the history of messages along the equilibrium path. If two different serious offers were submitted in period n for two distinct histories, then any seller's type whose strategy assigns positive probability to the sequence of messages along one of these histories and who accepts the resulting offer would be strictly better off if he were to send, instead, the other sequence of messages. This implies that messages do not matter. The equilibrium outcome is the same in this game as in the original one without messages.

When offers are private, assume that the seller's messages are private: that is, buyer n only observes the message that was sent in period n (of course, he observes all offers submitted previously). In that case as well, messages are irrelevant: all (pure or mixed strategy) equilibria of this game are realization-equivalent to some equilibrium without messages. The proof of this claim is more involved than the previous one and is available from the authors. However, the intuition is straightforward. Since messages and offers are not observed, they do not affect continuation payoffs. Consider the seller's type that is indifferent between accepting and rejecting the highest price that buyer n may submit given (one of) this seller's type equilibrium message. Then no other message can induce the buyer to assign positive probability to a higher offer, for otherwise the seller would prefer such a message. More generally, if there are two equilibrium messages inducing the buyer to offer prices that some seller's type would accept, then the distribution over these acceptable prices must be the same. Consequently, the seller could just as well send no message at all.

Offers by the seller

The previous extension allows for a first move by the seller in each period, but also preserves his position as a last mover. As we have argued in Section 4, this is because we view the deadlock in the case of public offers as a consequence of the inability of the seller to commit to not soliciting another offer. Assigning this position to the buyer would introduce some limited commitment power to the seller, because this would allow him to 'tie his hands'. To see this, consider the game in which, in every period, the seller moves first and makes an offer. The buyer moves then second and last, by either accepting or rejecting. Because the informed party makes offers, many equilibria can be supported by a suitable choice of beliefs. However, we claim that there exists an equilibrium in which agreement is reached in finite time. Here as well, we only sketch the construction (details are available upon request). In every period, the seller makes one of two offers, depending on his type. A lower offer is made by all the seller's types below some threshold, which the buyer accepts. A higher, constant offer is made by all the types above this threshold, which the buyer rejects. The lower offer is such that the buyer is indifferent between accepting it and rejecting it, given the threshold. Out-of equilibrium offers below (above) the lower offer are accepted (rejected) by the buyer, because his belief given such offers is identical to his belief given the lower equilibrium offer.

This construction extends the common wisdom that giving bargaining power to the informed party promotes efficiency to environments with interdependent values and provides a partial answer to the open question raised in Deneckere and Liang's (2006) conclusion. However, we view this power more as the commitment power that befalls the party that is not the last to move, as opposed to the bargaining power usually identified with the role as a proposer.

What if buyers do not know calendar time?

In our analysis of the hidden case, we have maintained the assumption that buyers can infer from calendar time the number of offers that have been submitted before. In applications, this assumption is rarely satisfied: the duration of unemployment, or the time-on-market of a house, is only an imperfect proxy of the number of offers previously rejected. Alternatively, one might consider the case in which buyers do not know their position in the sequence and update their beliefs about this position given the event that they are called upon to submit an offer. One possible, mathematically rigorous formalization of such an environment consists in assuming that the total number of buyers is finite and unknown, geometrically distributed with parameter δ. Thus, discounting appears naturally as a consequence of this uncertainty and buyers face a stationary problem. In the case of linear valuation and cost function, the unique equilibrium can be solved for. In the more interesting case of large δ and $\underline{x} < x_1$, buyers randomize between a losing and a serious offer, and their payoff is consequently zero. Expected delay is independent of the discount factor.

Multiple buyers in each period

As mentioned in the introduction, the logic of our results is similar to the mechanism at work in Nöldeke and van Damme (1990) and Swinkels (1999) in the context of Spence's signaling model. In contrast to these papers, we have considered so far the case in which there is only one offer made in every period. This modeling choice is consistent with the interpretation of the sequence of buyers as a unique, but impatient buyer. If there are two or more buyers submitting offers in every period, there is no longer a unique equilibrium in the game with public offers. In some equilibria, agreement is never reached. In others, agreement is reached in finite time. An example of such an equilibrium is available from the authors. Given this intriguing result, it would be interesting to know what would happen in the Spencian model with only one offer in every period. Taken together, these findings delineate fascinating strategic patterns whose understanding awaits further research.

References

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

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

[3] Blouin, M., 2003. Equilibrium in a decentralized market with adverse selection, *Economic Theory*, **22**, 245-262.

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

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

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

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

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

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

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

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

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

[13] 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.

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

[15] 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.

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

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

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

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

[20] Vincent, D., 1990. Dynamic Auctions, *Review of Economic Studies*, **57**, 49-61.

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

Appendix

A Preliminaries

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 5.3 through 5.9, though in a different order. We will start with Propositions 5.4 and 5.5. We then need to prove Proposition 5.6 – with the exception of the upper bound on N_0. Indeed, it is a logical preliminary to Propositions 5.7 and 5.8, which we prove next, and its proof is instrumental in the proof of Proposition 5.3.

A strategy of buyer n is a probability distribution σ_B^n over price offers. We denote $\tilde{p}_n$ the random price offered by buyer n. Any profile σ of such distributions induces a probability distribution over sequence of prices, which we denote $\mathbf{P}_\sigma$. Expectation w.r.t. $\mathbf{P}_\sigma$ is denoted by $\mathbf{E}_\sigma$.

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

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

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

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

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

It follows that

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

with equality if and only if buyer $n+1$ makes no offer above x: competition between successive buyers prevents p_n from being much below p_{n+1}.[14] Using a version of the envelope theorem, the function p_n has a left-derivative everywhere, given by

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

Note that $\mathbf{E}_\sigma\left[1 - \delta^{\tau_n(x)-n}\right]$ is non-increasing in x, 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 x, and the discounted time at which a seller with type x 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 $F_n(x)$ is the (unconditional) probability that the seller is of type $t \le x$, and rejects offers from buyers 1 through $n-1$. Letting $f_n(x) = \dfrac{1}{1-\underline{x}} \prod_{k=1}^{n-1} \mathbf{P}_\sigma\left(\tilde{p}_k < p_k(x)\right)$ denote the (normalized) probability that a seller with type x rejects the first $n-1$ offers, one has

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

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

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

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

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

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

[14] 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.

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

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

We stress that the quantities introduced so far, p_n, F_n, f_n, $\underline{x}_n$, π_n, $\overline{q}_n$ 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.

Since π_n is continuous, the equilibrium payoff π_n^* of buyer n is equal to $\max_{[\underline{x},1]} \pi_n$, and one has $\pi_n(x) = \pi_n^*$ for every $x \in T_n$, where T_n is the support of the random offer of buyer n.[15]

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

Lemma A.1 *Assume $\underline{x}_{n+2} > \underline{x}_{n+1}$, for some $n \in \mathbb{N}$. Then buyer n submits no offer in $(\underline{x}_n, \underline{x}_{n+2})$.*[16] *In particular, $\underline{x}_{n+1} = \underline{x}_n$, buyer n submits a losing offer with positive probability, and $\pi_n^* = 0$.*

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

Proof. Let a type $x \in (\underline{x}_n, \underline{x}_{n+2})$ be given. By assumption, a seller with type x 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 $\delta\left(\mathbf{E}_{\sigma^*}\left[\tilde{p}_{n+1} - c(x)\right]\right)$, and therefore, $p_n(x) = (1-\delta)c(x) + \delta\mathbf{E}_{\sigma^*}\left[\tilde{p}_{n+1}\right]$. Since $\delta > \overline{\delta}$, this implies that $v(x) - p_n(x)$ is increasing.

[15] That is, T_n is the smallest closed set of types that is assigned probability one by σ_B^n.
[16] That is, $\sigma_B^{*,n}$ assigns probability zero to $(\underline{x}_n, \underline{x}_{n+2})$.

Set $z := \inf\{x \in [\underline{x}_n, 1] : v(x) \geq p_n(x)\}$ (with $\inf \emptyset = \underline{x}_n$). Note that $D^-\pi_n(x) = (v(x) - p_n(x)) f_n(x) - c'(x)F_n(x) < 0$ on $(\underline{x}_n, z]$. On the other hand, on the interval $\left(z, \underline{x}_{n+2}\right]$, $D^-\pi_n$ 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

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

thus,

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

where the first inequality holds since $F_n(x) \leq f_n(x)$ 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 $\left(z, \underline{x}_{n+2}\right]$, hence π_n is strictly convex over $\left[z, \underline{x}_{n+2}\right]$.

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

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

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

B Proof of Proposition 5.4

We here prove that equilibrium payoffs are very small. Proposition B.1 below implies Proposition 5.4.

Proposition B.1 *The equilibrium payoff of each buyer n is at most*

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

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

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

We bound below each of the two terms.

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

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

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

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

Recall finally (5):

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

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

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

hence

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

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

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

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

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

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

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

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

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

Therefore,

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

Collecting (12) and (14) then yields

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

as desired. ■

C Proof of Proposition 5.5

The intuition for the proof is as follows. A seller with type $\underline{x}_{n_1+1} < \underline{x}_{n_2+1}$ 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 $\underline{x}_{n_1+1}$ expects to trade. In lemma C.1, we first provide a lower bound on $D^- p_n(x)$. We will then rely on the relation between $D^- p_n\left(\underline{x}_{n_1+1}\right)$ and the discounted time at which type $\underline{x}_{n_1+1}$ trades.

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

The proof of Lemma C.1 is a simple consequence of the following technical inequality.

Lemma C.2 *Let $h : [\underline{x}, 1] \to \mathbb{R}_+$ be non-decreasing. Then, for any $[a, b] \subseteq [\underline{x}, 1]$, 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 $\frac{0}{0} = 0$).

The proof of Lemma C.2 is postponed to the end of the section.

Proof of Lemma C.1. Since π_n is maximal at x, one has $D^-\pi_n(x) \leq 0$, that is,

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

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

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

The result then follows by applying Lemma C.2. ■

Proof of Proposition 5.5. Recall from (6) that $D^- p_n(x) = c'(x)(1 - \mathbf{E}_{\sigma^*}\left[\delta^{\tau_n(x)-n}\right])$, where $\tau_n(x) = \inf\{k > n : \tilde{p}_k \geq p_k(x)\}$. Consider buyer $n = n_1$, and $x = \underline{x}_{n_1+1} > \underline{x}_{n_1}$. Since $x \in T_n$, and by Lemma C.1, one has $\mathbf{E}_{\sigma^*}\left[\delta^{\tau_n(x)-n}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. On the other hand, $\tau_n(x) \leq n_2$, $\mathbf{P}_\sigma$-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 $x \in T_1$, ensures that $\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. This yields the upper bound in Proposition 5.3.

Proof of Lemma C.2. It is sufficient to prove that the inequality holds for step functions, as the result then follows using a limit argument. Since the inequality is homogenous w.r.t. $h(b)$, we may assume $h(b) = 1$.

We argue by induction over the cardinality of the range of h. If $h(t) = 1$ for all $t \in [\underline{x}, 1]$, then the desired inequality follows from $v(t) \leq v(b) + m_{v'}(t - b)$ ($t \in [a, b]$). Assume now that (15) holds for every $a < b$, and every step function that assumes at most n different positive values, and let $h = \sum_{i=0}^{n} \lambda_i 1_{[x_i, x_{i+1})}(\cdot)$ be a step function with possibly $n+1$ different positive values: $0 \leq \lambda_0 \leq \cdots \leq \lambda_n = 1$, and $a = x_0 < x_1 < \cdots < x_{n+1} = b$.

For such an h, we view the left-hand side of (15) as a function of λ_0,

$$\psi(\lambda_0) = \frac{1}{\sum_{i=0}^{n} \lambda_i (x_{i+1} - x_i)} \sum_{i=0}^{n} \lambda_i \int_{x_i}^{x_{i+1}} v(t)dt + \frac{m_{v'}}{2} \sum_{i=0}^{n} \lambda_i (x_{i+1} - x_i),$$

defined for $\lambda_0 \in [0, 1]$.

After simplification, the derivative of ψ is seen to be

$$\psi'(\lambda_0) = \frac{m_{v'}}{2}(x_1 - x_0) + \frac{1}{\left(\sum_{i=0}^{n} \lambda_i (x_{i+1} - x_i)\right)^2} \times \left\{ \sum_{i=1}^{n} \lambda_i \left((x_{i+1} - x_i) \int_{x_0}^{x_1} v(t)dt - (x_1 - x_0) \int_{x_i}^{x_{i+1}} v(t)dt \right) \right\}.$$

Since v is increasing, the summation between braces is negative, hence ψ' is non-decreasing, so that ψ is convex. As a result,

$$\psi(\lambda_0) \leq \left(1 - \frac{\lambda_0}{\lambda_1}\right) \psi(0) + \frac{\lambda_0}{\lambda_1} \psi(\lambda_1). \tag{17}$$

Note that if λ_0 is either equal to 0 or to λ_1, the function h assumes at most n different positive values. By the induction assumption, $\psi(0) \leq v(b)$ and $\psi(\lambda_1) \leq v(b)$, and the result follows from (17). ■

D Proof of Proposition 5.6

We here prove Proposition D.1 below. It corresponds to Proposition 5.6, except for the upper bound on N_0, which will be established in the proof of Proposition 5.3.

Proposition D.1 *There is a stage N_0 such that:*

P1 $T_n \subseteq \{\underline{x}_{N_0}, 1\}$, *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 $N := 1 + \max\{n : \max T_n < 1 \text{ and } F_n(1) \geq \frac{\nu}{M_{c'}}\}$. 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 $\pi_{N_0}(1) = F_{N_0}(1)\left(\overline{q}_{N_0}(1) - \overline{c}\right) \geq 0$, and thus $\pi_n(1) \geq 0$, for all $n \geq N_0$. We argue by contradiction, and assume that $T_n \cap (\underline{x}_{N_0}, 1) \neq \emptyset$, for some $n \geq N_0$, and we call n the first such stage. One thus has $\overline{q}_{n+1}(1) > \overline{q}_n(1)$, hence $\pi_{n+1}(1) > 0$ so that $\pi^*_{n+1} > 0$, and thus $\underline{x}_{n+2} > \underline{x}_{n+1}$. 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, $\underline{x}_{n+2} = 1$. By Lemma A.1, buyer n makes no serious offer in $(\underline{x}_n, \underline{x}_{n+2}) = (\underline{x}_{N_0}, 1)$. This is the desired contradiction. Therefore, $T_n \subseteq \{\underline{x}_{N_0}, 1\}$, for all $n \geq N_0$.

Assume now that $F_N(1) < \frac{\nu}{M_{c'}}$, and that $\overline{x} := \max_{n<N} \max T_n < 1$. We will prove that no buyer $n \geq N_0$ ever submits a serious offer in $(\overline{x}, 1)$. Since $\lim_n F_n(1) = 0$, this will imply 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 $T_n \subseteq \{\underline{x}_{N_0}, 1\}$, 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 an offer in $(\overline{x}, 1)$. Since this is also true for buyers $k < N$, one has $f_n(t) = \frac{1}{1-\underline{x}}$, for $t \in [\overline{x}, 1]$.

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

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

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, $\tilde{\pi}_n(x) \geq \pi_n(x)$, with equality if $x = 1$. The derivative of $\tilde{\pi}_n$ is given by

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

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

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 A.1, buyer $N_0 - 1$ would make no serious offer in $(\underline{x}_{N_0-1}, \underline{x}_{N_0+1}) = (\underline{x}_{N_0-1}, 1)$, which would stand in contradiction to the definition of N_0. ■

E Proof of Proposition 5.7

For convenience, we recall below the statement of Proposition 5.7. As stated in Proposition D.1, any offer is either winning or losing from stage N_0 on. However, the specific behavior of buyers $n \geq N_0$ is to a large extent indeterminate. If trade takes place in bounded time according to σ^*, then an irrelevant change in σ^* yields an equilibrium under which trade occurs in finite, but not bounded, time. And vice-versa.

Proposition E.1 *Let N_0 be as given in Proposition D.1. Then:*

A *There is an equilibrium σ, with $\sigma^n_B = \sigma^{*,n}_B$ for all $n < N_0$, and such that $T_n = \{1\}$, for some $n \geq N_0$;*

B *There is an equilibrium σ, with $\sigma^n_B = \sigma^{*,n}_B$ for all $n < N_0$, and such that $T_n = \{\underline{x}_{N_0}, 1\}$, for all $n \geq N_0$.*

Since the proof will involve two profiles, σ and σ^*, some of the notations will be starred when they refer to σ^*. Obviously, either $T^*_n = \{1\}$ for some $n \geq N_0$, or $T^*_n = \{\underline{x}_{N_0}, 1\}$ for all $n \geq N_0$. Hence, only **A** or **B** has to be proven, depending on σ^*.

Proof. Assume first that $T_n^* = \{\underline{x}_{N_0}, 1\}$ for all $n \geq N_0$, and let us prove that **A** holds. Define $n_* \in \mathbb{N}$ through the inequalities $\delta^{n_*+1} \leq \mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right] < \delta^{n_*}$. Note that $\tau(1) \geq N_0$ a.s., and $\tau(1) > N_0$ with positive probability, hence $n_* \geq N_0$.

We define the buyers profile σ_B as:

- $\sigma_B^n = \sigma_B^{*,n}$, for $n < N_0$;
- buyer n assigns probability one to the offer $c(\underline{x}_{N_0})$, for $n = N_0, \ldots, n_* - 1$;
- buyer n_* assigns probability $\dfrac{\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right] - \delta^{n_*+1}}{\delta^{n_*}(1-\delta)} \in [0,1)$ to a winning offer, and offers $c(\underline{x}_{N_0})$ otherwise.
- buyer $n_* + 1$ submits a winning offer (and later buyers as well).

We first claim that the optimal acceptance policy of the seller in any stage $n < N_0$ is the same, when facing either σ_B or σ_B^*. As a consequence, σ_B^n is a best reply to σ_B^{-n}, for such n's.

Let $n < N_0$ and a type x be given. When computing the optimal continuation payoff of type x, we may restrict our attention to those stopping times τ that either accept an offer before N_0, or wait until the first winning offer is received: $\tau \leq N_0 - 1$, or $\tau = \tau(1)$. For such a τ, one has

$$\begin{aligned}\mathbf{E}_\sigma\left[\delta^{\tau-n}(\tilde{p}_\tau - c(x))\right] &= \mathbf{E}_\sigma\left[\delta^{\tau-n}(\tilde{p}_\tau - c(x)1_{\tau<N_0}\right] + (\bar{c} - c(x))\mathbf{E}_\sigma\left[\delta^{\tau(1)-n}\right] \times \mathbf{P}_\sigma(\tau \geq N_0) \\ &= \mathbf{E}_{\sigma^*}\left[\delta^{\tau-n}(\tilde{p}_\tau - c(x))\right]\end{aligned}$$

since buyers strategies coincide up to stage N_0, and since $\mathbf{E}_\sigma\left[\delta^{\tau(1)}\right] = \mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right]$.

This readily implies that $p_n(x) = p_n^*(x)$ for all $n < N_0$ and $x \in [\underline{x}, 1]$, as desired.

We next prove that σ_B^n is a best reply to σ_B^{-n} for all $N_0 \leq n \leq n_* + 1$, by induction over n. Observe first that the distribution of types faced by buyer N_0 is the same under both profiles σ and σ^*: $F_{N_0}(\cdot) = F_{N_0}^*(\cdot)$. In particular, $c(\underline{x}_{N_0})$ is a losing offer. Assume that $\dfrac{F_{N_1}(\cdot)}{F_{N_1}(1)} = \dfrac{F_{N_1}^*(\cdot)}{F_{N_1}^*(1)}$ for some $N_0 \leq n \leq n_* + 1$.

By construction, $\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}|\tau(1) > n\right] \leq \mathbf{E}_\sigma\left[\delta^{\tau(1)}|\tau(1) > n\right]$. Therefore, one has $p_n(x) \geq p_n^*(x)$, with equality if $x = \underline{x}_n$, or $x = 1$. It follows that $\dfrac{\pi_n(\cdot)}{F_n(1)} \leq \dfrac{\pi_n^*(\cdot)}{F_n^*(1)}$, with equality if $x = \underline{x}_n$ or $x = 1$.

Therefore, both the offer $c(\underline{x}_{N_0})$ (a losing offer) and the winning offer are optimal for buyer n given σ^{-n}. The concludes the proof of **A**.

We now assume that $T_n^* = \{1\}$ for some $n \geq N_0$, and let N_1 be the first such n. Observe that, by Lemma A.1, $S_{N_1-1} \subseteq \{\underline{x}_{N_1-1}, 1\}$, hence $N_1 > N_0$. We define the profile σ as follows:

- $\sigma_B^n = \sigma_B^{*,n}$ for all buyers $n < N_1 - 1$;
- buyer $N_1 - 1$ assigns probability π to a winning offer, and probability $1 - \pi$ to the price offer $c(\underline{x}_{N_0})$;
- all buyers $n \geq N_1$ assign probability α to a winning offer, and probability $1 - \alpha$ to the price offer $c(\underline{x}_{N_0})$.

Given α, the probability π is chosen s.t.

$$\pi_* + (1 - \pi_*)\delta = \pi + (1 - \delta)\frac{\alpha\delta}{1 - \delta(1 - \alpha)}, \tag{18}$$

where $\pi_* < 1$ is the probability assigned by σ_N^{*,N_1-1} to a winning offer. Provided α is close to one, $\pi \in [0, 1)$. In addition, we assume that α is sufficiently close to one so that

$$1 - \delta(1 - \alpha) > (1 - \delta)\frac{2M_{c'} + M_{c''}}{m_{v'}}.$$

The choice of π in (18) ensures that $\mathbf{E}_\sigma\left[\delta^{\tau(1)}\right] = \mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right]$, and $\mathbf{E}_\sigma\left[\delta^{\tau(1)}|\tau(1) > n\right] = \mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}|\tau(1) > n\right]$, for all $n < N_1 - 1$. As in the proof of **A** above, this implies that $p_n(\cdot) = p_n^*(\cdot)$, hence σ_N^n is a best reply to σ_B^{-n}, for all buyers $n < N_1 - 1$, and the distribution of types faced by buyer $N_1 - 1$ is the same under the two profiles σ and σ^*. In particular, the offer $c(\underline{x}_{N_0})$ is a losing offer.

It remains to prove that no buyer $n \geq N_1 - 1$ would find it profitable to make a serious, non-winning offer. For such n, one has

$$\begin{aligned} p_n(x) - c(x) &= (\bar{c} - c(x))\mathbf{E}_\sigma\left[\delta^{\tau_n(1)}\right] \\ &= (\bar{c} - c(x))\left(\delta\alpha + (1 - \alpha)\alpha\delta^2 + \cdots\right) \\ &= (\bar{c} - c(x))\frac{\alpha\delta}{1 - \delta(1 - \alpha)}, \end{aligned}$$

hence

$$p_n(x) = \frac{1}{1-\delta(1-\alpha)}\left((1-\delta)\,c(x) + \alpha\delta\overline{c}\right). \tag{19}$$

We now replicate an argument used in the proof of Proposition D.1. Recall that $D^-\pi_n(x) = (v(x) - p_n(x))f_n(x) - D^-p_n(x)F_n(x)$, and let $z = \min\{x \in \underline{x}_{N_0} : v(x) \geq p_n(x)\}$. Then π_n is decreasing over $\left[\underline{x}_{N_0}, z\right]$, and convex on $[z,1]$. The result follows. ■

F Proof of Proposition 5.8

For convenience, we recall the statement of Proposition 5.8.

Proposition F.1 *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 $T_n = \{\underline{x}_n\}$, 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 $\overline{x}_{n_*} := \max T_{n_*} < 1$.

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

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

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

Since buyers $n \leq k < n_*$ only submit losing offers, the distribution of types faced by buyers n and n_* is the same, and $\overline{q}_n(\overline{x}_{n_*}) = \overline{q}_{n_*}(\overline{x}_{n_*})$. It follows that $\pi_n(\overline{x}_{n_*}) > \pi_{n_*}(\overline{x}_{n_*}) = \pi_{n_*}^* \geq 0$ – 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 offer, it must be to a serious offer $x_n > \underline{x}_n$, and then $\underline{x}_{n+1} = x_n$. On the other hand, $p_{n-1}(x_n) - c(x_n) = \delta(p_n(x_n) - c(x_n))$, hence $p_{n-1}(x_n) < p_n(x_n)$. By Lemma A.1, $\pi_{n-1}^* = 0$ and buyer $n-1$ makes no offer in $(\underline{x}_{n-1}, x_n)$, hence $\overline{q}_{n-1}(x_n) = \overline{q}_n(x_n)$. As above, this implies $\pi_{n-1}(x_n) > 0$ – a contradiction. ■

G Proof of Proposition 5.9

We here prove Proposition 5.9: in the presence of further concavity and convexity assumptions on v and c, the equilibrium distributions $\sigma_B^{*,n}$ have finite support.

Plainly, this is true for $n \geq N_0$, since $T_n \subseteq \{\underline{x}_{N_0}, 1\}$. For $n < N_0$, we argue by induction. Assume that the strategies $\sigma_B^{*,1}, \ldots, \sigma_B^{*,n-1}$ have finite support, for some $n \geq 1$. This implies that $f_n(t)$ is a step function: there exist types $x_0 = \underline{x} < x_1 < \cdots < x_K = 1$, and values $0 \leq \lambda_0 < \lambda_1 < \cdots < \lambda_{K-1}$, such that $f_n(t) = \lambda_k$ for $t \in [x_k, x_{k+1})$.

Claim: the map $x \mapsto \overline{q}_n(x)$ is concave over $[x_k, x_{k+1}]$, for each $0 \leq k < K$.

Proof of the claim. Fix k. For $x \in [x_k, x_{k+1}]$, the average quality $\overline{q}_n(x)$ is given by $\overline{q}_n(x) = Q_n(x)/F_n(x)$, with $Q_n(x) := \displaystyle\int_{\underline{x}_n}^{x} v(t) f_n(t) dt$. Both Q_n and F_n are twice differentiable on $[x_k, x_{k+1})$, with $F_n'(x) = \lambda_k$, $F_n''(x) = 0$, $Q_n'(x) = \lambda_k v(x)$ and $Q_n''(x) = \lambda_k v'(x)$. Thus,

$$\overline{q}_n''(x) = \frac{1}{F_n(x)^3} \left\{ F_n(x)^2 Q_n''(x) - 2F_n'(x) \left(Q_n'(x) F_n(x) - Q_n(x) F_n'(x) \right) \right\}$$

has the same sign as

$$N(x) := v'(x) F_n(x)^2 - 2\lambda_k v(x) F_n(x) + 2\lambda_k Q_n(x).$$

Note that $N'(x) = v''(x) F_n(x)^2 \leq 0$, hence

$$\begin{aligned} N(x) &\leq N(x_k) = v'(x_k) F_n(x_k)^2 - 2\lambda_k v(x_k) F_n(x_k) + 2\lambda_k Q_n(x_k) \\ &= v'(x) \left\{ \sum_{i=0}^{k-1} \lambda_i (x_{i+1} - x_i) \right\}^2 - 2\lambda_k v(x_k) \left\{ \sum_{i=0}^{k-1} \lambda_i (x_{i+1} - x_i) \right\} + 2\lambda_k \left\{ \sum_{i=0}^{k-1} \lambda_i \int_{x_i}^{x_{i+1}} v(t) dt \right\} \end{aligned}$$

For given $\lambda_0 < \lambda_1 < \cdots < \lambda_{K-1}$, we view this right-hand side as a function h of $x_1, x_2, \ldots, x_k$, defined over the domain $x_1 \leq \cdots \leq x_k$.

After simplification, the derivative of the map $x \mapsto h(x_1, x_2, \ldots, x_j, x, x, \ldots, x)$ is given by

$$v''(x) F_n(x)^2 + 2v'(x) F_n(x) (\lambda_j - \lambda_k),$$

and is therefore non-positive. Thus, $h(x_1, x_2, \ldots, x_j, x_{j+1}, x_{j+1}, \ldots, x_{j+1}) \leq h(x_1, x_2, \ldots, x_j, x_j, x_j, \ldots, x_j)$ for every j (with strict inequality if $x_j < x_{j+1}$) and therefore,

$$h(x_1, x_2, \ldots, x_k) \leq h(x_1, x_1, \ldots, x_1) = 0.$$

This concludes the proof of the claim. $\blacksquare$

We now complete the proof of the proposition. For any type x, one has $\pi_n(x) = F_n(x)(\overline{q}_n(x) - p_n(x))$. One thus has

$$\overline{q}_n(x) - p_n(x) \leq \frac{\pi_n^*}{F_n(x)}, \tag{20}$$

with equality if $x \in T_n$.

The left-hand side of (20) is strictly concave over each interval $[x_k, x_{k+1}]$, while the right-hand side is convex. Hence, there is at most one x in $[x_k, x_{k+1}]$, such that $\overline{q}_n(x) - p_n(x) = \pi_n^*/F_n(x)$: therefore, $T_n \cap [x_k, x_{k+1}]$ is either empty, or a singleton. This concludes the induction step.

H Equilibrium delay

Recall that $\tau(1) = \inf\{n : \tilde{p}_n = \overline{c}\}$ is the first buyer who submits a winning offer. From the proof of Proposition 5.5, we know that $\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right] \leq 1 - \frac{m_{v'}}{2M_{c'}}$. We here proceed to provide a lower bound for $\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right]$, which is independent of δ and of the equilibrium σ.

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

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

By **Steps 1–3**, one thus has $\mathbf{E}_{\sigma^*}\left[\delta^{\tau(1)}\right] \geq C_2 \times \left(\delta^{\frac{1}{1-\delta}}\right)^{C_0+C_1}$, 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, which links the discounted expectation of a random variable (hereafter, r.v.) to its tail distribution.

Lemma H.1 *Let τ be a random time with integer values, and such that*

$$\mathbf{E}\left[\delta^{\tau-n}|\tau > n\right] \geq a, \tag{21}$$

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 $b > 0$.

Proof. For simplicity, set $N_{a,b} := \frac{1}{ab(1-\delta)}$. Among the r.v.'s that satisfy (21), choose a r.v. τ_* for which $\mathbf{P}(\tau_* \geq N_{a,b})$ is minimal.[17] For $n \geq 1$, set $p_n^* := \mathbf{P}(\tau_* = n|\tau_* \geq n)$. Casual inspection shows that one must have $1 < \tau_* \leq N_{a,b} + 1$ $\mathbf{P}$-a.s., and that the constraint $\mathbf{E}\left[\delta^{\tau_*-n}|\tau_* > n\right] \geq a$ is binding whenever $p_n^* > 0$.

As a result, there is a stage N_* such that:

- $\mathbf{E}\left[\delta^{\tau_*-n}|\tau_* > n\right] = a$ for all $1 \leq n \leq N_*$;
- $p_n^* = 0$ for all $N_* < n \leq N_{a,b}$.

The value of N_* is determined by the condition $\delta^{N+1-(N_*+1)} > a \geq \delta^{N+1-N_*}$:

$$N_* = -\left\lfloor \frac{\ln a}{\ln \delta} \right\rfloor + 1 + N_{a,b}.$$

For $1 < n \leq N_*$, the value of p_n^* is pinned down by the condition $\mathbf{E}\left[\delta^{\tau_*-n}|\tau_* > n\right] = a$. One gets $p_n^* = \dfrac{a(1-\delta)}{1-\delta a}$ for $1 < n < N_*$, and $p_{N_*} \leq \dfrac{a(1-\delta)}{1-\delta a}$. Therefore,

$$\mathbf{P}(\tau_* \geq N_{a,b}) \leq \left(1 - \frac{a(1-\delta)}{1-\delta a}\right)^{N_*-1},$$

and the result follows by standard algebraic manipulations, which we omit. ■

We now proceed to **Steps 1–3**. Our computation of C_1 and C_0 will involve three parameters. We choose α, β s.t. $0 < \alpha < \beta < \nu$,[18] and $\eta \in (0, \frac{\beta-\alpha}{2v(1)})$. Next, set $K = 1 + \left\lceil \frac{M_{v'}(1-x}{2(\nu-\beta)} \right\rceil$, $\varepsilon = \frac{1}{2}\eta^K$, and

[17] Existence of such a r.v. follows from standard compactness arguments.

[18] Where $\nu = \min_n (v(x) - c(x))$ is a lower bound on the gains from trade.

$a = \frac{\alpha\delta}{\bar{c} - c(\underline{x})}$. We will express C_1 and C_0 as functions of these constants. We will make no attempt at optimizing the choice of α, β and η.

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

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

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 x_n, is bounded away from zero. That is, from buyer n's viewpoint, types below x_n have a significant probability.

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

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

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

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

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

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 $y_0 = \underline{x}$ and

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

until $y_J = 1$. In particular, $\mathbf{E}\left[v(t)|t \in [y_j, y_{j+1}]\right] = c(y_{j+1}) + \beta$ 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

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

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 $x_n \geq y_{j_n}$, which will yield $F_n(x_n) \geq F_n(y_{j_n}) \geq 2\varepsilon F_n(1)$, 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,[19]

$$\left|\mathbf{E}_{F_n}\left[v(t)|t \leq y_{j_n}\right] - \mathbf{E}_{F_n}\left[v(t)|t \in [y_{j_n-1}, y_{j_n}]\right]\right| \leq 2\eta v(1), \tag{22}$$

since $|\mathbf{E}(X) - \mathbf{E}(X 1_A)| \leq 2\mathbf{P}(\overline{A}) \sup |X|$ for every bounded r.v. X and every event A. The first expectation in (22) is $\overline{q}_n(y_{j_n})$, while the second one is at least $c(y_{j_n}) + \beta$. Therefore, $\overline{q}_n(y_{j_n}) \geq c(y_{j_n}) + \alpha$ and thus, $x_n \geq y_{j_n}$. ■

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

$$p_m(x) - c(x) \leq (\overline{c} - c(x))\mathbf{E}_{\sigma^*}\left[\delta^{\tau_n(x) - m}|\tau_n(x) > m\right]. \tag{23}$$

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

On the other hand, $p_m(x) = \overline{q}_m(x) \geq \overline{q}_n(x)$ if $\pi^*_m = 0$, and then $p_m(x) - c(x) \geq \alpha$, while $p_m(x) - c(x) \geq \delta(p_{m+1}(x) - c(x)) \geq \underline{\delta}\alpha$ if $\pi^*_m > 0$, since then, $\pi^*_{m+1} = 0$. Using (23), this implies

$$\mathbf{E}_{\sigma^*}\left[\delta^{\tau_n(x)-m} | \tau_n(x) > m\right] \geq \frac{\underline{\delta}\alpha}{\overline{c} - c(x)} \geq a.$$

Apply now Lemma H.1 to obtain $\mathbf{P}_\sigma(\tau(x) \geq n + N_{a,\varepsilon} | \tau(x) \geq n) \leq \varepsilon$.

Finally, observe that

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

whereas

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

Therefore,

$$\frac{F_{n+N_{a,\varepsilon}}(x)}{F_n(1)} \leq \mathbf{P}_\sigma(\tau(x) \geq n + N_{a,\varepsilon} | \tau(x) \geq n) \leq \varepsilon,$$

as desired. ■

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

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

On the other hand, $N_0 = \inf\{n : \overline{q}_n(1) = \overline{c}\}$. Indeed, $\overline{q}_{N_0}(1) = \overline{c}$ since $\pi^*_{N_0} = 0$, and $\overline{q}_{N_0}(1) > \overline{q}_{N_0-1}(1)$ 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 $\overline{q}_n(1)$ to increase from $c(x_*)$ to $\overline{c}$. Between stages N_1 and N_0, and using the proof of Proposition D.1, no buyer ever submits a serious offer above x_*. Hence, $\overline{q}_n(1)$ increases steadily with time, at a speed which is related to the probability with which successive buyers do trade. Lemma H.2 below provides a precise estimate of this relationship.

Lemma H.2 *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:*

$$\overline{q}_m(1) - \overline{q}_n(1) \geq \frac{m_{v'}(1-\underline{x})}{2} \times (1-x_*)\frac{\pi_{n,m}}{F_m(1)}. \tag{24}$$

The proof of Lemma H.2 is tedious and somewhat lengthy. It is postponed to the end of the section. By Lemma H.1, one has, as in **Step 1**, $\pi_{n,n+N_{a,\varepsilon}} \geq \varepsilon$. By Lemma H.2, one thus has

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

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

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

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

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

which yields

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

This concludes the proof.

Proof of Lemma H.2. Fix the distribution f_n of types faced by buyer n, and the value of $\pi_{n,m}$. We will minimize $\overline{q}_m(1)$ 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, $\overline{q}_m(1)$ is minimal when $\int_{\underline{x}}^1 g(t)f_n(t)v(t)dt$ 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 $[x_*, 1]$ (since there is no serious offer beyond x_*);

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

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

Thus,

$$\overline{q}_m(1) - \overline{q}_n(1) \geq \frac{1}{F_m(1)}\int_{\underline{x}}^1 v(t)g(t)f_n(t)dt - \frac{1}{F_n(1)}\int_{\underline{x}}^1 v(t)f_n(t)dt. \tag{25}$$

The rest of the proof consists in showing that the right-hand side of (25) is at least equal to the right-hand side in (24).

Plugging g into (25), one has

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

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

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

$$\overline{q}_m(1) - \overline{q}_n(1) \geq \frac{\pi_{n,m}}{F_m(1)F_n(x_*)}(1-x_*)m_{v'}\left\{F_n(x_*)\frac{1-x_*}{2} + \int_{x_*}^{x}(x_*-t)f_n(t)dt\right\}. \tag{26}$$

Consider finally the right-hand side of (26). For a given value of $F_n(x_*)$, the integral is minimized when f_n is constant over $[\underline{x}, x_*]$, and equal to $F_n(x_*)/(x_* - \underline{x})$. The integral is then equal to $\frac{1}{2}(x_* - \underline{x})F_n(x_*)$. Substituting into (26), this yields

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

as desired. ■