

Competition and Price Transparency in the Market for Lemons: Experimental Evidence*

Olivier Bochet[†] Simon Siegenthaler[‡]

October 10, 2018

Abstract

In markets with asymmetric information the price mechanism often fails to allocate goods efficiently. Such inefficiencies can be alleviated through bargaining, where agents engage in repeated interaction and in the process learn about others' types. We design an experiment to examine the efficiency of bargaining in markets that are impaired by adverse selection. In line with the theoretical predictions, we find that in the absence of price transparency (i.e., competitors cannot observe each other's price offers) subjects increase prices over time, leading to trade of high quality goods and high efficiency levels. Contrary to the predictions, the efficiency-enhancing effects of competition persist even when offers are observable. We explore different behavioral explanations for the absence of a transparency effect. Remarkably, asking subjects to make decisions via the strategy method improves their understanding of the strategic opportunities offered by the game: while the strategy method has no impact on behavior if offers are unobservable, it delivers the predicted loss in efficiency if offers are observable. We complement our study with some robustness checks on the intensity of competition, the role of time frictions and risk aversion.

JEL Classification: C92, C70, D82

Keywords: Adverse Selection, Bargaining, Information Transmission, Screening, Strategic Complexity

*We thank Jack Fanning, Philippe Jehiel, Alistair Wilson, John Wooders and three anonymous referees for helpful comments. We also thank seminar participants at NYU Abu Dhabi, Singapore Management University, University of Chile, and the 2018 AEA meeting in Philadelphia.

[†]Division of Social Science, New York University Abu Dhabi, PO Box 129188, United Arab Emirates. E-mail: olivier.bochet@nyu.edu. Phone: +971 2628 5470

[‡]Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080, USA. E-mail: simon.siegenthaler@utdallas.edu. Phone: +1 972-883-5871.

1 Introduction

The presence of informational asymmetries can have a devastating effect on the efficiency of markets. For instance, in situations of adverse selection the presence of low quality goods reduces the buyers' willingness to pay, thereby eliminating trading opportunities for sellers of high quality goods (Akerlof, 1970). The reason markets collapse under adverse selection is the inability of the price mechanism to convey sufficient information about the quality of the goods: in equilibrium a single market price prevails that cannot accurately reflect the different reservation values of the sellers.

Numerous institutions are potentially impaired by information asymmetries between buyers and sellers; among others trading platforms such as eBay, the market for used cars, the housing market, and asset markets more broadly. Many of these institutions do not fall under Akerlof's static set-up, because sellers typically have several chances of selling their goods. In a dynamic setting, the alternative to not trading today is to trade in the future. This can have a correcting effect on rates of trade and lead to a partial alleviation of the adverse selection effect. The reason is that delaying agreement is costly—either because of time frictions or because of the risk that no better offers will be made in the future—and hence by rejecting offers a seller can endogenously signal her type or the quality of the good that she has for sale.¹

The ability of repeated interactions to promote trade persists across a wide range of market institutions, including bilateral bargaining (Evans, 1989; Vincent, 1989; Deneckere and Liang, 2006), when several uninformed agents compete for a single seller (Hörner and Vieille, 2009a) or in large decentralized markets where agents meet in pairs and are rematched in each period (Blouin and Serrano, 2001; Moreno and Wooders, 2010, 2016; Virag, 2016). However, it turns out that the effectiveness of bargaining as a mechanism to transmit information about the seller's true type and to facilitate trade of high quality goods critically depends on the transparency of offers (Hörner and Vieille, 2009a; Kim, 2015, 2016). Offers are said to be transparent or *public* if they can be observed among the competing buyers. If offers are *private*, buyers only know for how long (i.e., for how many time periods) the good has been for sale. For example, online trading platforms often display previous offers but eBay has recently also introduced the possibility for a seller to privately negotiate with a buyer without any information transmitted to other prospective buyers. The housing market serves as another example where previous offers may be may not be revealed.

¹It should be noted that while a dynamic set-up tends to increase rates of trade of high quality goods, it also introduces delay as an additional source of inefficiency. Depending on the trading environment, the effect of repeated interactions on efficiency can thus be ambiguous. Samuelson (1984) shows that no mechanism can lead to first-best efficiency if adverse selection is sufficiently strong (moreover, static mechanisms are in fact constrained efficient).

In this article, we experimentally examine the efficiency of bargaining in different settings with adverse selection where buyers have incomplete information about the seller's type. We ask whether the possibility to make repeated offers helps promote trade and efficiency and how this depends on the degree of competition and transparency in a market. All offers are made by the buyers and offers can be accepted or rejected by the seller. Our design includes three main treatments: (i) exclusive bargaining where there is one buyer and one seller, (ii) competitive bargaining with private offers where three buyers compete to trade with one seller and buyers cannot observe each other's offers, and (iii) competitive bargaining with public offers where the three buyers can observe each other's offers. For each of the main treatments, we run a treatment with time frictions, that is, there is an exogenous breakdown probability after each rejected offer, and a treatment without time frictions in which case there is a commonly known number of bargaining stages.

We study these environments via a game-theoretic model and use the predictions to formulate three main hypotheses. *Hypothesis 1: time frictions promote trade of high quality goods.* In the presence of time frictions delaying an agreement is costly. Informed agents with potentially high benefits from trade are less willing to postpone agreement (or equivalently, more willing to accept low prices) than agents with lower benefits. This allows the uninformed agents to use specific price sequences to screen the different types of the informed agent. As a result, bargaining can lead to trade despite the presence of adverse selection. Intuitively, low quality goods are sold early at a low price and high quality goods are sold late at a higher price.

Given the benchmark case of exclusive bargaining, Hypotheses 2 and 3 then focus on the effect of competition and transparency. *Hypothesis 2: If offers are private, competition among buyers promotes rates of trade and efficiency compared to exclusive bargaining.* In other words, if the competing buyers cannot observe each other's offers, competition drives up prices which tends to speed up trade and therefore leads to higher welfare levels. However, this positive effect of competition is reversed if offers are observable. *Hypothesis 3: If offers are public, competition among buyers reduces rates of trade and efficiency compared to exclusive bargaining.* In the presence of publicly observable offers, buyers have no incentive to outbid their competitors. This is because such offers would be observed by future buyers and countered by even higher offers. As a result, in theory, offers stay low throughout the bargaining process, leading to low rates of trade and welfare levels. To sum up, depending on the transparency of price offers, competition has a diametrically opposed effect on market outcomes.

There are several reasons why these theoretical benchmarks warrant an empirical examination. They are based on sophisticated equilibrium reasoning and it is instructive to see if the predictions hold up at least qualitatively. Indeed, it is unclear if the model succeeds at capturing the main determinants of behavior in an actual bargaining situation. An experiment

allows us to explore this and can potentially inspire new theories if the predicted effects are not observed or need to be qualified. Note that we are primarily interested in the qualitative features of the equilibria and to a lesser extent in the quantitative predictions. Our experiment also helps us understand real-world markets such as the housing market where it is common that an informed seller faces a series of potential buyers. For instance, it addresses the question of how competition on the uninformed market side affects bargaining and whether sellers are better off when revealing information about past offers. Finally, this article contributes to the economics literature on bargaining, which has been successful at combining theoretical and experimental work to generate new insights into human behavior (see the literature review below). We follow this tradition by experimentally exploring factors—time frictions, competition, and price transparency—that have been deemed relevant in the recent theoretical literature.

The experimental results confirm some but not all of our hypotheses. First, we find that exclusive bargaining leads to screening of low-type sellers, much like the qualitative predictions of the model. In particular, rates of trade with high-type sellers are significantly boosted upwards in the presence of time frictions, while trade failures are common without time frictions. Second, the data confirms that competitive bargaining leads to high levels of efficiency. This result holds irrespective of whether offers are private or public. Hence, in contradiction to theory, the transparency of offers does not affect behavior in the main treatments of the experiment. On the one hand, this is good news, because it shows that in dynamic settings competition can alleviate adverse selection effects irrespective of the institutional details. On the other hand, we would like to understand why the predicted differences between the settings with public and private offers are not borne out in the data.

To this end, we present a range of robustness checks for the finding that the transparency of offers does not have the anticipated effect. The original theoretical result of Hörner and Vieille (2009a) assumed an infinite stream of buyers. We show that the theoretical results continue to hold in our setting with three buyers taking turns in making offers. Nevertheless, it is instructive to see what happens if we bring the experimental setting closer to the original model. To explore this we implemented treatments with six buyers. We find that behavior remains largely the same as in the three-buyer treatments, in particular, offer transparency still doesn't affect subjects' behavior. We also check whether behavior depends on the way we model time frictions and find that outcomes remain the same if we use discounting instead of a breakdown probability.

Finally, we ran a set of treatments in which subjects make their decisions via the strategy method. The motivation was to check if at least some buyers exhibit behavior consistent with the equilibrium predictions of the public offers case. Instead of asking buyers to enter their

offers stage by stage, all decisions are made at the start of a bargaining game. In the case of public offers, subjects also had to condition their offers on the possible offers in the previous stages. Strikingly, we find a strong effect of offer transparency in these treatments: The rate of trade with high-type sellers in competitive bargaining with public offers decreases to 11%, while it remains at 55% with private offers. It seems that the strategy method highlights that current offers can be conditional on past offers and, as it turns out, this is sufficient to bring behavior closer to the theoretical predictions. We link this insight to analogy-based expectation equilibrium proposed in Jehiel (2005). The concept captures the idea that agents often use simplified representations of the available strategic opportunities. Applied to our case, analogy-based expectation equilibrium is in line with the data if we assume that players' choices are mainly driven by the duration of the bargaining game while bunching together decision nodes that differ only in previously observed price offers.

It is instructive to compare our results with the findings in Bochet and Siegenthaler (2018). There, we focused on bilateral bargaining under adverse selection, in particular, how it fares relative to a take-it-or-leave-it-offer institution. We found that trade with high-type sellers occurred substantially later than predicted. The reason for this delay is related to the fact that time frictions take the form of discounting, implying that the stakes at play become smaller over time (gains in later stages are smaller, but so are losses). Risk averse bargainers welcome the lower variability in earnings and are thus more willing to delay agreement. In the present article, time frictions take the form of a breakdown probability. As a consequence, there is no delay beyond the risk-neutral predictions in the exclusive bargaining treatment. Because the focus of the present article is on competition and price transparency, the fact that the outcome in exclusive bargaining is in line with theory allows us to use it as a clean benchmark against which we can evaluate the competitive bargaining treatments.²

There is a well-established experimental literature on bargaining with incomplete information.³ Rapoport, Erev and Zwick (1995), Reynolds (2000), and Fanning and Kloosterman (2018) study bargaining games with a focus on the Coase Conjecture. Price sequences in their experiments resemble the theoretically predicted shape, although the authors also find deviations from the comparative statics implied by their models (e.g., the effect of a change in

²Discounting and exogenous breakdowns both intend to model agents' impatience. Which approach is preferable depends on the context. The goal in this paper is to study competition and offer transparency and thus we wanted to keep the stakes at play constant across bargaining stages (i.e., no discounting). Another reason we chose to model impatience through a breakdown probability is that we don't need to specify a maximum number of offers that can be made.

³See Roth and Malouf (1979), Roth and Murnighan (1982) and Roth and Schoumaker (1983) for seminal contributions and Mitzkewitz and Nagel (1993), Straub and Murnighan (1995), Croson (1996), Rapoport, Sundali and Seale (1996), Güth and Van Damme (1998), and Nagel and Harstad (2004) for studies exploring experimental ultimatum games. Cason and Reynolds (2005) study the impact of bounded rationality in bargaining environments.

the discount factor). Valuations are independent in these studies, while we focus on the case of interdependent valuations where theory predicts trade failures due to adverse selection. Moreover, we focus on competition and the transparency of offers rather than the bilateral setting. Another related experiment is provided by Forsythe, Kennan and Sopher (1991) who test the explanatory power of truth-telling constraints in free-form bargaining (subjects can send hand-written messages). The experiment shows that truth-telling constraints indeed drive the occurrence of trade failures. In our experiment, subjects always face a situation in which, theoretically, inefficiencies are unavoidable due to adverse selection. We explore to what extent bargaining can help improve outcomes and how the answer depends on the specifics of the trading institution.

Our experiment also contributes to the experimental literature on price setting in decentralized markets with search or discounting costs. Siegenthaler (2017) shows that cheap-talk can be informative in such settings, thus representing an alternative to repeated offers in helping alleviate adverse selection. Cason and Friedman (2003) and Cason and Noussair (2007) examine price dispersion for different matching procedures, e.g., they test the predictions that price offers correspond to the monopoly price if buyer-seller meetings are bilateral and tend to approach Bertrand pricing if a seller meets more than one buyer simultaneously. They find evidence in favor of the theoretical predictions, while the results in Davis and Holt (1996) and Abrams, Sefton and Yavas (2000) suggests that theory fails to predict outcomes. In our setting, the Bertrand price is predicted to occur with private offers, while the monopoly price should be observed when offers are public.

Finally, there is a literature starting with Abreu and Gul (2000) that examines the effects of obstinate or behavioral types in bargaining. Obstinate types commit to a certain behavior (e.g., rejecting any offer below a certain price) at the start of the bargaining process. The presence of such types has interesting implications. In particular, rational players have an incentive to behave as if they were obstinate. Embrey, Fréchette and Lehrer (2015) confirm the existence of such effects experimentally. Fanning (2016) looks at the interaction of deadlines and obstinate types, while Fanning (2014) provides a discussion in the context of bargaining under incomplete information. The literature is relevant for us, as it can explain the type of price sequences we observe in our exclusive bargaining treatment.

The remainder of the paper is organized as follows. The next section presents the model and characterizes the equilibria for the different trading environments. Section 3 presents the experimental design and derives a set of hypotheses that will guide our data analysis. Section 4 provides the main results on exclusive bargaining, competition and offer transparency. It also presents a discussion on risk aversion and several robustness checks. Finally, Section 5 concludes.

2 Exclusive and Competitive Bargaining

2.1 Model

A seller and $n \geq 1$ buyers bargain over the price at which a single, indivisible good is traded. The seller can be of two types $\theta = \{L, H\}$. That is, the good is either of low (L) or high (H) quality. The reservation costs of the low-type seller are normalized to $c_L = 0$ and the reservation costs of the high-type seller are $c_H > 0$. The buyers' valuations are v_L for a low-quality and v_H for a high-quality good. We assume positive gains from trade for both qualities, i.e., $v_L > c_L$ and $v_H > c_H$. The seller's type is private information. The probability that the seller is a high type is $q < (c_H - v_L)/(v_H - v_L)$. The condition on q implies that offering c_H (which in equilibrium is accepted by both seller types) at belief q yields a negative expected payoff. In other words, buyers are willing to offer c_H only after some belief updating has taken place.

Bargaining takes the following form. All offers are made by the buyers. Buyers queue to sequentially make offers to the seller. For instance, with three buyers, the game starts with buyer 1 offering a price to the seller. If the seller rejects the offer, buyer 1 joins the end of the queue and buyer 2 is called to make the next offer. If buyer 2's offer is also rejected, it is buyer 3's turn. If buyer 3's offer is rejected, buyer 1 returns to make another offer and so on. The game ends if the seller accepts an offer. If a price offer p is accepted by a θ -type seller, the buyer who made the offer earns $v_\theta - p$ and the seller earns $p - c_\theta$. Buyers who do not trade earn 0. The game also ends if the bargaining process breaks down before the seller accepts an offer. Specifically, whenever an offer is rejected, there is a continuation probability $r \in [0, 1)$ that the next stage is entered. With probability $1 - r$ the bargaining process ends, in which case everyone earns 0. We assume throughout that $rc_H > v_L$, implying that we're interested in situations where players are patient.

The model reduces to bilateral or *exclusive* bargaining if $n = 1$. If $n > 1$, bargaining is said to be *competitive*. In the presence of competition, the transparency of offers turns out to be important. Offers are said to be *private* if buyers only know their own past offers. Offers are *public* if the buyer in stage t can observe the full price sequence that has been offered up to stage $t - 1$.

2.2 Equilibria

We next present the equilibria of the bargaining model. The equilibrium concept is perfect Bayesian equilibrium as defined in Definition 8.2 in Fudenberg and Tirole (1991). We begin

with the (on-path) equilibrium predictions for exclusive bargaining.

Proposition 1 (Deneckere and Liang, 2006): In the bargaining game with a single buyer ($n = 1$), there exists a unique equilibrium. The equilibrium price sequence is given by $(p_1, p_2, \dots, p_{\bar{T}}) = (r^{\bar{T}-1}c_H, r^{\bar{T}-2}c_H, \dots, c_H)$ where $\bar{T} > 1$ is the finite number of stages it takes until the offer of c_H is made. The high-type seller rejects all offers up to and including stage $\bar{T} - 1$ and accepts the final equilibrium offer $p_{\bar{T}} = c_H$. The low-type seller mixes between accepting and rejecting offers in each stage $t = 1, \dots, \bar{T} - 2$ such that the buyer is indifferent between offering p_t and p_{t+1} . In stage $\bar{T} - 1$, the low-type seller accepts $p_{\bar{T}-1} = rc_H$.

A full statement of the off-equilibrium behavior can be found in proposition 1 of Deneckere and Liang (2006) and Bochet and Siegenthaler (2018), including a proof that the equilibrium is unique. Note that the buyer uses an increasing price sequence such that the low-type seller is indifferent between accepting and delaying acceptance at each point in time. The low-type seller accepts in each stage with positive probability until the buyer is certain to face a high-type seller and offers $p_{\bar{T}} = c_H$. In particular, the low-type seller's behavior renders the buyer indifferent between offering p_t and p_{t+1} in stage t . This ensures that the buyer cannot gain by accelerating or slowing down trade with an off-equilibrium move.^{4,5} It is worth noting that the equilibrium in exclusive bargaining implies a positive probability of trade with high-type sellers. However, it is also clear that bargaining introduces a risk of trade failure with low-type sellers, since they don't accept for sure in stage 1. For instance, whether the above equilibrium is more or less efficient than a simple take-it-or-leave-it offer depends on the value of r (for more on this, see Bochet and Siegenthaler, 2018).

The analysis for the model with multiple buyers ($n > 1$) builds upon the analysis in Hörner and Vieille (2009a). Their model assumes an infinite stream of buyers. It turns out that while the details of the equilibrium are different, the main equilibrium properties remain unchanged for finite n , as long as n is sufficiently large. Proposition 2 states the (on-path) equilibrium behavior if offers are unobservable.

Proposition 2: Fix some $r \in (v_L/c_H, 1)$ and let $n > 1$ be sufficiently large. Then, in the bargaining game with n buyers and private offers, the high-type seller trades with strictly positive probability. Moreover, there exists an essentially unique equilibrium with the following features:

⁴For instance, suppose the buyer were to deviate to a price of $p'_t \in (p_t, p_{t+1}]$. Then the equilibrium prescribes that the buyer mixes between p_{t+1} and the higher offer p_{t+2} in stage $t + 1$, such that the expected offer in $t + 1$ is p'_t/r . This implies that the low-type seller is indifferent between accepting and rejecting the off-equilibrium offer p'_t . She accepts the offer p'_t with the same probability as p_t and hence the deviation is not profitable. All other deviations are deterred in the same way.

⁵The length of the price sequence \bar{T} is given by the maximum number of stages such that the cumulative ex-ante acceptance probability of the low-type seller does not exceed 1, working back from stage $\bar{T} - 1$ and applying the requirement that the buyer has to be indifferent between offering p_t and p_{t+1} at each step.

- The buyer in stage 1 offers $p_1 = v_L$. The low-type seller accepts p_1 with probability $a_1 = (\mu^* - q)/(\mu^*(1 - q))$ such that the buyers' belief after observing a rejection equals μ^* , solving $\mu^*v_H + (1 - \mu^*)v_L - c_H = 0$.
- For each stage $\ell \geq 2$, the corresponding buyer i_ℓ randomizes between the winning offer $p_\ell = c_H$ (accepted with probability 1) and a losing offer $p_\ell \leq v_L$ (rejected with probability 1). For stages $\ell \geq 3$, the probability λ^* on the winning offer solves $(1 - \mu^*)(v_L - \tilde{p}) + \mu^*r^n(1 - \lambda^*)^{n-1}(v_H - c_H) = 0$.⁶ For stage $\ell = 2$, the probability λ_2 on the winning offer solves $v_L = rc_H\lambda_2 + [r^2c_H\lambda^*(1 - \lambda_2)]/[1 - r(1 - \lambda^*)]$.

A proof of Proposition 2 can be found in the Appendix. Let us explain the equilibrium construction. The first buyer offers v_L . This offer is accepted with a probability such that the buyers' posterior belief to face a high-type seller equals μ^* . At this belief, a buyer who makes the winning offer c_H has an expected payoff of 0. From stage 2 onwards buyers are thus indifferent between offering c_H and a losing offer. In contrast to the case with an infinite stream of buyers, with finite n , a buyer potentially wants to offer $\tilde{p} \in (v_L, c_H)$ to screen out the low-type seller, followed by an offer of c_H when it is his turn to make another offer. The offer \tilde{p} must thus be the minimum offer that is accepted for sure by the low-type seller who anticipates that the buyer's offer n stages down the road will be c_H and the other buyers offer c_H with probability λ^* . Notice that the probability of a buyer to be able to make another offer is decreasing in λ^* . In equilibrium, λ^* is just high enough to deter such screening attempts. Finally, it turns out that if all buyers $\ell \geq 2$ were to mix according to λ^* , the expected payoff of the low-type seller when rejecting an offer in stage 1 would exceed v_L . But since v_L must be accepted with positive probability in stage 1, in equilibrium the buyer in stage 2 offers c_H with probability $\lambda_2 < \lambda^*$ such that the expected payoff of the low-type seller in stage 1 is exactly v_L . This is possible, because the buyer in stage 1 has belief $q < \mu^*$ and thus a lower incentives to screen.

The proof in the appendix shows that the behavior in Proposition 2 is the unique equilibrium for n sufficiently large relative to r . For large n , the incentive to screen is sufficiently small, because of the high probability that bargaining breaks down before a buyer comes back to make another offer. For our experimental parameters, the equilibrium indeed exists.⁷

⁶This expression corresponds to a buyer's expected payoff when trying to screen out the low-type seller with an offer of \tilde{p} followed by an offer of c_H . The price \tilde{p} is the lowest price accepted in such a situation, i.e.,

$$\tilde{p} = r^n(1 - \lambda^*)^{n-1}c_H + r\lambda^*c_H \frac{1 - r^{n-1}(1 - \lambda^*)^{n-1}}{1 - r(1 - \lambda^*)}.$$

⁷It is worth noting that our proof does not exclude the possibility that for some combinations of n and r the equilibrium of Proposition 2 is one among several. However, as it turns out, the data from the experiment is in fact in line with the described equilibrium, showing that it provides a useful benchmark.

Finally, if we allow buyers to observe each other's offers, the equilibrium looks remarkably different from the case of private offers.

Proposition 3: In the bargaining game with $n > 1$ buyers and public offers, the high-type seller never trades. Moreover, there exists a unique equilibrium with the following features:

- Buyers offer $p_\ell = 0$ in all stages $\ell = 1, 2, \dots$
- The low-type seller accepts the first offer $p_1 = 0$ with probability a_1 (as given in Proposition 2) and rejects all subsequent offers.

The proof can be found in the Appendix. As with private offers, the low-type seller in stage 1 must be indifferent between accepting and rejecting the first offer and accepts with a probability such that the buyers' belief moves to μ^* in stage 2. However, the initial offer is now 0 instead of v_L . The reason is that for any positive equilibrium offer, the buyer making the offer would have an incentive to deviate and make a slightly lower offer. In particular, the seller cannot deter the lower offer by lowering the acceptance probability. All other buyers would be able to observe the lower offer and realize that the acceptance probability has been lowered, too. In turn, this would lower their expected offers. The seller is therefore equally well off accepting the offer from the deviating buyer. The only offer where such a deviation is not possible is when buyers offer 0. An offer of 0 can indeed be supported in equilibrium. Any deviation to a higher offer in some stage l would trigger an even higher expected offer in stage $l + 1$. Hence, the observability of offers coupled with competition from future buyers keeps offers low in the public offer setting and as a result high-type sellers don't trade.

Given the equilibria presented in Propositions 1-3, we can compute the ex-ante efficiency levels defined as the sum of ex-ante expected payoffs for all players.

Corollary 1: The efficiency levels in the competitive environments are independent of r and are given by

$$e^{\text{Private}} = (1 - q)v_L$$

$$e^{\text{Public}} = (1 - q)v_L - qv_L \frac{v_H - c_H}{c_H - v_L}.$$

The efficiency level for exclusive bargaining as r goes to 1 is

$$e^{\text{Exclusive}} = (1 - q)v_L - qv_L \frac{v_H - c_H}{c_H}.$$

It follows that $e^{\text{Private}} > e^{\text{Exclusive}} > e^{\text{Public}}$.

The efficiency levels for the competitive bargaining environments follow directly from Propositions 2 and 3. If offers are private, the low-type seller in stage 1 is indifferent between accepting and rejecting the offer v_L . Hence, her expected payoff must be v_L . Since the high-type seller as well as all buyers have an expected payoff of 0, ex-ante efficiency equals the expected payoff of the low-type seller times the probability $(1 - q)$ that the seller is indeed of the low type. If offers are public, trade with the low-type seller occurs in stage 1 with probability a_1 and occurs with probability 0 thereafter. The efficiency level is thus $a_1(1 - q)v_L$. Plugging in a_1 yields the result.

The efficiency level for the case of exclusive bargaining is derived in section 6 of Deneckere and Liang (2006). Intuitively, as r gets closer to 1, the buyer's payoff approaches 0 since he loses almost all of his bargaining power. Indeed, the buyer's preferred bargaining institution would be a take-it-or-leave-it offer (see Samuelson, 1984). Hence, the efficiency level is approximately equal to the low-type seller's payoff. Further, because the increasing price sequence used by the buyer to screen out the low-type seller starts below v_L , the low-type seller's expected payoff is below v_L implying that ex-ante efficiency is below $(1 - q)v_L$. In Bochet and Siegenthaler (2018) we show that the efficiency level is non-monotonic in r and maximized at an interior value of r . However, as long as r is sufficiently large, the efficiency ranking of the bargaining institutions remains as given in Corollary 1 and, importantly, is preserved for our experimental parameters.

3 Experimental Design

The experiment implements the model presented in the previous section. The treatments vary the degree of competition and whether offers are private or public (price transparency). We also study the case without time frictions ($r = 1$) but a finite, commonly-known number of bargaining stages. The remaining parameters are constant across all treatments: the buyers' valuations are $v_H = 23$ and $v_L = 10$, the seller's reservation costs are $c_H = 16$ and $c_L = 0$, and the seller's probability of being a high type is $q = 1/3$.

3.1 Treatments

Table 1 presents the different treatments. The upper half shows our main treatments. In addition, we ran several control treatments to provide further evidence which will help us understand better the results of the initial data set. The robustness checks are shown in the lower half of table 1.

Table 1: Experimental Design

Treatment	Subjects ^a	Competition	Transparency	Time Friction
Main Experiment				
Exclusive	48 (8)	1 Buyer	–	$r = 0.9$
Exclusive T	48 (8)	1 Buyer	–	Known T
Private	84 (7)	3 Buyers	No	$r = 0.9$
Private T	36 (3)	3 Buyers	No	Known T
Public	84 (7)	3 Buyers	Yes	$r = 0.9$
Public T	36 (3)	3 Buyers	Yes	Known T
Robustness Checks				
Private 6B	63 (3)	6 Buyers	No	$r = 0.9$
Private Strategy ^b	36 (3)	3 Buyers	No	$r = 0.9$
Public 6B	63 (3)	6 Buyers	Yes	$r = 0.9$
Public Strategy ^c	36 (3)	3 Buyers	Yes	$r = 0.9$

Sessions were run at the labs of the University of Bern (first wave: 264 subjects) and Valencia (second wave: 270 subjects). (a) Number of independent observations (matching groups) in parentheses. (b) Subjects chose offers for bargaining stage t conditional on reaching t . (c) Subjects chose offers for bargaining stage t conditional on reaching t and the offer in stage $t - 1$.

Exclusive Bargaining: In treatment *Exclusive*, we set $n = 1$ and $r = 0.9$. Hence, there is a single buyer making a sequence of offers to the seller. Offers can be made from the discrete grid $\{0, 0.01, 0.02, \dots, 23\}$. We do not require offers to be increasing over time. If an offer is rejected, bargaining ends with a probability of 10%; with a probability of 90% bargaining continues and the buyer can make another offer. Treatment *Exclusive T* is identical to treatment *Exclusive*, except that the stage in which bargaining ends is commonly known. That is, instead of a random breakdown due to r , there is a pre-announced stage T after which the bargaining process ends. The number of available bargaining stages follows the same distribution as the realized breakdown stages in treatment *Exclusive*.⁸

Competitive Bargaining with Private Offers: In treatment *Private*, the number of buyers is $n = 3$, that is, three buyers take turns in making offers. Whether a buyer is the first, second, or third to make an offer was randomly chosen at the start of each bargaining game. The continuation probability is $r = 0.9$. The three buyers do not observe each other's offers. Treatment *Private T* removes the time frictions from treatment *Private*, i.e., there is a pre-announced final stage T . Treatment *Private 6B* is identical to treatment *Private* except that there are six buyers.

⁸For instance, after session 1 of treatment *Exclusive*, the stages in which bargaining breakdowns occurred (or would have occurred) were used to determine the pre-announced number of stages T in each bargaining game of session 1 in treatment *Exclusive T*. This was done separately for each session. The same procedure was used for the competitive bargaining treatments.

Competitive Bargaining with Public Offers: In treatment *Public*, the number of buyers is $n = 3$ and the continuation probability is $r = 0.9$. In contrast to treatment *Private*, offers are observable, in particular, the three buyers observe each other’s previous offers. In treatment *Public T* there is a pre-announced final stage T . Treatment *Public 6B* is identical to treatment *Public* except that there are six buyers.

Comparing treatments *Exclusive*, *Private*, and *Public* will allow us to examine the effects of competition and offer transparency on efficiency. Comparing each treatment to its counterpart with a known breakdown stage T will allow us to study the impact of time frictions on bargaining outcomes. The treatments with six buyers highlight that competition is sequential or intertemporal, as it is much less likely that a buyer will be able to make more than one offer than in the treatments with three buyers.

In a final set of treatments *Private Strategy* and *Public Strategy* we elicited participants’ choices via the “strategy method.” For instance, buyers in the first position chose their offers for stages 1, 4, 7, etc. at the start of a bargaining game. Buyers in the second position chose their offers for stages 2, 5, 8, etc., and buyers in the third position chose their offers for stages 3, 6, 9, etc. This procedure allows us to observe the buyers’ offers in stages that are not reached in the actual bargaining game. Sellers chose whether to accept/reject offers by entering a minimum acceptable offer for each stage. If the buyer’s offer for a given stage exceeds a seller’s minimum acceptable offer, the offer is accepted. Otherwise, the offer is rejected and the game continues to the next stage (or there is a bargaining breakdown). The same procedure was used in treatment *Public Strategy*. However, because past offers are observable in this treatment, buyers and sellers chose their offers and minimum acceptable offers for stage t conditional not only on the stage but also on the offer in stage $t - 1$.⁹ Notice that the offer in stage 1 is always made unconditionally and thus determines which future offers will be relevant.

3.2 Hypotheses

For our experimental hypotheses we are less interested in the exact quantitative predictions but rather the qualitative predictions of our model. The equilibrium rates of trade and efficiency levels are summarized in table 2. Recall that efficiency is measured as the sum of ex-ante expected payoffs of all market participants.

⁹In principle, buyers and sellers could base their choice for stage t on all previous offers in stages 1, 2, ..., $t - 1$. In order to be able to implement the strategy method in an comprehensible way, we asked subjects to choose offers and minimum acceptable offers conditional only on the offer in $t - 1$. Specifically, we group offers in $t - 1$ into 11 bins 0-2, 2.01-4, 6.01-8, 8.01-10, 10.01-12, 12.01-14, 14.01-16, 16.01-18, 18.01-20, and 20.01-23. Importantly, the theoretical predictions on offer transparency continue to hold if the buyer in stage t can only observe the offer of the buyer in stage $t - 1$.

Table 2: Equilibrium Predictions

Treatment	$Pr(\text{Trade} \theta = L)$	$Pr(\text{Trade} \theta = H)$	Efficiency ^a
Exclusive	0.74	0.48	6.04
Exclusive T	1	0	6.66
Private	0.78	0.63	6.66
Private T	1	0.63	8.13
Public	0.42	0	2.78
Public T	1	0	6.66

(a) Efficiency equals the sum of ex-ante expected payoffs of all market participants. The predicted efficiency level for *Private 6B* and *Private Strategy* is the same as for *Private* and likewise the efficiency for *Public 6B* and *Public Strategy* is the same as for *Public*.

There is a unique equilibrium for treatment *Exclusive*. Given the parameters used in the experiment, the buyer's optimal price sequence is (7.7, 8.5, 9.4, 10.5, 11.7, 13, 14.4, 16). The increasing price offers exhaust the low-type seller's patience before trade with a high-type seller takes place. The corresponding acceptance probabilities of the low-type seller are (0.22, 0.19, 0.19, 0.20, 0.24, 0.32, 1.00, -). The high-type seller accepts in stage 8. Trade is reached with both seller types unless there is an exogenous breakdown before the offer of 16 is made. The efficiency level is 6.04, falling short of the first-best efficiency $(1 - q)v_L + q(v_H - c_H) = 9$ due to the delay before trade occurs.

In the Appendix, we characterize the equilibrium for treatment *Exclusive T*. The equilibrium is essentially unique. Important for us to note is that the buyer offers 0 in all stages. The low-type seller accepts for sure in stage *T*. However, in contrast to treatment *Exclusive*, high-type sellers don't trade in *Exclusive T*, showing that time frictions are essential for trade with high-type sellers.¹⁰ These observations lead to our first hypothesis:

Hypothesis 1: *In treatment Exclusive, the buyer uses an increasing price sequence to screen the seller. High-type sellers trade with positive probability. In treatment Exclusive T, only low-type sellers trade.*

The equilibrium for treatment *Private* predicts that the buyer's offer in stage 1 is 10, which is accepted by the low-type seller with probability 0.42. When entering stage 2, all buyers update their belief that the seller is a high type from 0.33 to 0.46. At this belief, a buyer is indifferent between offering 16 and a losing offer. All subsequent buyers offer 16 with positive probability, ensuring that the low-type seller was indifferent between accepting and rejecting the offer of 10 in stage 1. Both seller types reject all offers below 16 from stage 2 onwards.

¹⁰Interestingly, efficiency is predicted to be higher in treatment *Exclusive T* than in treatment *Exclusive*. This does not imply that *Exclusive T* is invariably preferable from a social perspective. In some circumstances, one may want to maximize market liquidity (trading rates) rather than efficiency.

Notice how competition drives up prices (buyers now have an expected profit of 0) such that, as shown in table 2, high-type sellers trade with a higher probability than under exclusive bargaining. Hence, competition promotes efficiency:

Hypothesis 2: *Competitive bargaining with private offers increases efficiency and rates of trade with both seller types compared to exclusive bargaining.*

In treatment *Public*, the equilibrium is such that buyers offer 0 in all stages. As with private offers, the offer in stage 1 is accepted by the low-type seller with probability 0.42 and the buyers update their beliefs from 0.33 to 0.46. But then all future offers are rejected until the bargaining process breaks down. Hence, high-type sellers do not trade and only some low-type sellers trade. As shown in table 2, relative to treatment *Exclusive*, competition lowers efficiency if offers are public:

Hypothesis 3: *Competitive bargaining with public offers decreases efficiency and rates of trade with both seller types compared to exclusive bargaining.*

Remarkably, whether offers are private or public—seemingly an institutional detail—is a key factor determining the performance of our trading environments. Why is this so? In the case of private offers, as long as profits are positive, buyers have an incentive to slightly increase their offer. Such a deviation would not be observed by the other buyers and hence the low-type seller would accept the offer. In contrast, when offers are public, any deviation to an offer above 0 would be observed and used by the seller to solicit an even higher offer by the next buyer. Anticipating this, buyers refrain from making offers above 0. Stated differently, it is the inability of the seller to commit to accepting a positive offer that prevents price offers to increase above 0.¹¹

3.3 Procedures

The first wave of data collection took place at the experimental laboratory of the University of Bern with a total of 264 participants. We decided to collect additional data in a second wave to support and better explain the results of the first experiments. The second wave was done at the University of Valencia with a total of 270 participants. The data from Valencia includes three matching groups for treatments *Private* and *Public* (72 participants). This allows us to check if there are any differences in behavior between Bern and Valencia. In the

¹¹In the Appendix we show that the effect of competition and offer transparency are independent of time frictions. In particular, *Private T* leads to the same rate of trade with high-type sellers as *Private*. Similarly, *Public T* leads to the same rate of trade with high-type sellers as *Public*. This highlights that, in contrast to exclusive bargaining, screening and trade with high-type sellers in treatment *Private* does not occur due to time frictions. Instead screening is caused by the possibility to make repeated offers coupled with competition.

Online Appendix we show that behavior does not differ between the two locations, looking at average opening, accepted prices, rates of trade with both seller types as well as efficiency. We are therefore able to establish our main results on offer transparency in two separate subjects pools. In the following we will pool this data across locations. The sessions in Valencia also included new treatments. Specifically, we implemented sessions with six instead of three buyers as well as treatments where decisions were made using the strategy method (see section 3.1). In total 534 students from various fields participated in the experiment. The experiment was programmed in z-Tree (Fischbacher, 2007). Sessions lasted less than 100 minutes and earnings averaged €25.66 (€28.40 in Bern and €23.00 in Valencia), including a show-up fee of €10.

At the start of a session we distributed the instructions of the respective treatment (available in the Online Appendix). Once the participants finished reading the instructions, a member of the experimenter team provided a brief verbal summary. Participants were also asked to answer a set of control questions. There were ten rounds, that is, the bargaining games described above were played ten times where each subjects participated in only one treatment. At the start of each round, subjects were randomly assigned the role of a buyer or a seller. Sellers' types (high or low) were drawn according to the probability $q = 1/3$. In each bargaining round, subjects were randomly rematched and this was commonly known. The goal was to minimize repeated game effects.¹²

At the end of each session we elicited subjects' risk preferences. Subjects were presented six lotteries, each of which they could accept or decline. A lottery gave a 50-50 chance between winning €6 or losing an amount of either €2, €3, €4, €5, €6, or €7. One lottery was randomly selected for payment. If the selected lottery was accepted, the earnings/losses were realized. The earnings didn't change for subjects who declined the selected lottery. The fact that the lotteries may result in a loss is consistent with the bargaining game. In general, if subjects made losses, these were subtracted from the show-up fee of €10.

4 Results

The discussion of the experimental results is separated into three sections. In section 4.1 we discuss the main hypotheses on competitive bargaining and offer transparency. Section 4.2 examines the data on the individual level by identifying price sequences and the effects of

¹²The data shows no trends in behavior over the ten bargaining rounds and no endgame effects. This indicates that the random matching successfully induced subjects to treat each bargaining round as a separate game. Moreover, subjects' behavior was not affected by how many times they were assigned the role of a buyer or seller. The corresponding analysis can be found in the Online Appendix.

risk aversion. Finally, section 4.3 presents robustness checks and our results on the strategy method. All non-parametric tests are based on matching group averages as the unit of observation (table 1 shows the number of matching groups per treatment). We use Wilcoxon-Mann-Whitney tests for between-treatment comparisons and Wilcoxon signed-rank tests for comparisons within treatment. Unless stated otherwise, the analysis will be based on bargaining rounds 3 to 10. We drop the first two rounds to account for the possibility that some subjects were experimenting and learning about the environment early on.

4.1 Time Frictions, Competition, and Offer Transparency

Table 3 summarizes the experimental results. The theoretical predictions are given in parentheses. We begin with the exclusive bargaining setting:

Result 1: *In exclusive bargaining, time frictions promote trade with high-type sellers. That is, the rate of trade with high-type sellers is significantly higher in *Exclusive* than in *Exclusive T*. However, because time frictions lower the rate of trade with low-type sellers, there is no significant difference in efficiency between the two treatments.*

Support: The rate of trade with high-type sellers is 43% in *Exclusive* and 20% in *Exclusive T* ($p = 0.006$). The average trading price in *Exclusive* is 8.25 with low-type sellers and 17.87 with high-type sellers ($p = 0.011$), implying that buyers were able to separate the two seller types. Screening was successful in *Exclusive*: 79% of the trades with low-type sellers occurred at an ex-post individually rational price for the buyer, i.e., at a price of 10 or below, significantly more than the 49% predicted in theory ($p = 0.011$). The rate of trade for low-type sellers is higher in *Exclusive T* than in *Exclusive* ($p = 0.004$). The efficiency level is 5.63 in *Exclusive* (versus 6.05 in theory) and 6.43 in *Exclusive T*, a difference that is not significant ($p = 0.291$). Taken together, result 1 is fully in line with hypothesis 1. \square

Observed outcomes in treatments *Exclusive* and *Exclusive T* are close to the theoretical predictions both in terms of efficiency and rates of trade.¹³ We use the behavior observed in treatment *Exclusive* as a benchmark of comparison for the treatments with competition. We next state our main results on competition and offer transparency:

Result 2: *Competitive bargaining with private offers promotes efficiency and increases rates of trade compared to exclusive bargaining.*

¹³The prices accepted by low-type sellers in *Exclusive T* are substantially larger than the predicted offer of 0. This is not surprising given the literature on ultimatum games showing that people are reluctant to accept offers which allocate less than half of the gains of trade to them (see footnote 3).

Table 3: Summary of Outcomes in Main Treatments

Treatment	Seller Type	Trading Rate	Efficiency	Trading Price	Trading Stage	Opening Offer	Buyer Profit	Seller Profit
Exclusive	<i>Overall</i>	0.61 (0.65)	5.63 (6.05)	11.46 (12.64)	5.06 (5.42)	4.97 (7.65)	1.75 (1.05)	3.95 (5.00)
	<i>High</i>	0.43 (0.48)	2.98 (3.35)	17.87 (16.00)	7.31 (8)	5.02 (7.65)	2.81 (3.35)	0.37 (0.00)
	<i>Low</i>	0.70 (0.74)	6.95 (7.39)	8.25 (10.95)	3.93 (4.14)	4.95 (7.65)	1.21 (-0.10)	5.74 (7.49)
Exclusive T	<i>Overall</i>	0.65 (0.67)	6.34 (6.67)	10.29 (-)	8.23 (-)	4.03 (0.00)	2.68 (6.67)	3.75 (0.00)
	<i>High</i>	0.20 (0.00)	1.38 (0.00)	17.96 (-)	9.67 (-)	3.95 (0.00)	1.78 (0.00)	-0.13 (0.00)
	<i>Low</i>	0.88 (1.00)	8.83 (10)	6.45 (0.00)	7.51 (10.00)	4.07 (0.00)	3.13 (10.00)	5.70 (0.00)
Private	<i>Overall</i>	0.77 (0.73)	6.97 (6.67)	13.20 (14.33)	4.10 (4.90)	7.23 (10.00)	0.32 (0.00)	6.03 (6.67)
	<i>High</i>	0.73 (0.63)	5.09 (4.38)	17.63 (16.00)	7.00 (6.40)	7.25 (10.00)	1.49 (1.47)	0.71 (0.00)
	<i>Low</i>	0.79 (0.78)	7.91 (7.81)	10.99 (13.5)	2.66 (4.15)	7.22 (10.00)	-0.26 (-0.73)	8.69 (10.00)
Private T	<i>Overall</i>	0.83 (0.88)	7.83 (8.13)	12.92 (12.88)	7.14 (9.48)	6.77 (10.00)	0.43 (0.49)	6.72 (6.67)
	<i>High</i>	0.50 (0.63)	3.50 (4.38)	18.32 (16.00)	9.64 (10.00)	6.24 (10.00)	1.45 (1.47)	-0.31 (0.00)
	<i>Low</i>	1.00 (1.00)	10.00 (10.00)	10.23 (11.31)	5.89 (9.22)	7.03 (10.00)	-0.07 (-0.73)	10.23 (10.02)
Public	<i>Overall</i>	0.76 (0.28)	6.89 (2.78)	14.09 (-)	3.94 (-)	7.21 (0.00)	0.10 (0.93)	6.62 (0.00)
	<i>High</i>	0.75 (0.00)	5.22 (0.00)	18.51 (-)	5.63 (-)	7.41 (0.00)	1.26 (0.00)	1.52 (0.00)
	<i>Low</i>	0.77 (0.42)	7.73 (4.17)	11.87 (0.00)	3.09 (1.00)	7.11 (0.00)	-0.48 (1.39)	9.17 (0.00)
Public T	<i>Overall</i>	0.74 (0.67)	7.11 (6.67)	12.15 (-)	7.97 (-)	6.23 (0.00)	0.40 (2.22)	6.03 (0.00)
	<i>High</i>	0.25 (0.00)	1.75 (0.00)	17.67 (-)	11.50 (-)	6.03 (0.00)	0.81 (0.00)	-0.27 (0.00)
	<i>Low</i>	0.98 (1.00)	9.78 (10.00)	10.15 (0.00)	6.20 (9.22)	6.33 (0.00)	0.20 (3.33)	9.19 (0.00)

The table shows averages across bargaining games by seller type. Overall (high- and low-type sellers) averages are calculated by weighing the observed averages by the theoretical probability of the seller types. Data includes observations from bargaining rounds 3-10 (bargaining games where high-type sellers accept prices below 16 are excluded). Theoretical predictions are given in parentheses.

Support: The rate of trade with high-type sellers is 73% in *Private*, which is significantly larger than in *Exclusive* ($p = 0.017$). The rate of trade with low-type sellers is also larger in *Private* (79% versus 70%) although the difference is not significant ($p = 0.220$). In terms of efficiency, these observations imply that *Private* outperforms *Exclusive* ($p = 0.063$), which is confirmed by the results of the regression model in the last column of table 5. The efficiency level in *Private* is similar to the theoretical predictions (6.97 versus 6.67). The increased trade frequency with high-type sellers in *Private* comes at the cost of a higher average trading price with low-type sellers, 10.99 in *Private* versus 8.25 in *Exclusive* ($p = 0.003$). However, trading prices in *Private* are still different for trades with low- and high-type sellers ($p = 0.018$), indicating that screening is an important component of subjects' behavior: the fraction of trades with low-type sellers occurring at a price below 10 is 69%, higher than the predicted 42% ($p = 0.017$). \square

Result 2 on private offers is in line with the theoretical prediction that competition promotes efficiency if offers are unobservable and confirms hypothesis 2. Let us now turn towards the public offers case.

Result 3: *The transparency of offers does not affect rates of trade and efficiency in competitive bargaining. Observed outcomes in treatments Public and Private are not significantly different from each other. By extension, efficiency and rates of trade are significantly higher in treatment Public than in treatment Exclusive.*

Support: The rate of trade for high-type sellers in *Public* is 75%, significantly larger than in *Exclusive* ($p = 0.010$) and not significantly different from treatment *Private* ($p = 0.796$). Trading rates with low-type sellers are similar across the three treatments. As a consequence, efficiency in *Public* is significantly higher than in *Exclusive* ($p = 0.081$), while there is no significant difference between *Private* and *Public* ($p = 0.654$). Interestingly, the percentage of trades with low-type sellers that occurred at an ex-post individually rational price (i.e. at a price of 10 or below) is 59% in *Public*. Since this percentage was 69% in *Private*, this suggest that buyers in *Public* were more likely to raise prices quickly. We will return to this in section 4.2. These observations are confirmed by the results of the regression model in the last column of table 5. The observed trading prices also don't differ between *Public* and *Private* ($p > 0.277$ for both seller types). Since non-parametric tests are conservative, we also performed multilevel regressions (with individual and matching group random intercepts) confirming that there are no significant differences between *Public* and *Private*, neither for opening offers ($p = 0.625$), trade with high-type sellers ($p = 0.598$), or efficiency ($p = 0.967$). Note that the efficiency level of 6.89 in *Public* is substantially higher than the predicted level of 2.78 ($p = 0.017$). \square

Result 3 rejects hypothesis 3. The transparency of offers does not affect behavior, contrary

to the theoretical predictions. While we did not expect public offers to fully eliminate trading opportunities for high-type sellers, the complete absence of an effect of transparency is puzzling. In section 4.3, we will examine this finding in more detail. We conclude the current section with a result on the effect of time frictions in competitive bargaining:

Result 4: *In competitive bargaining, time frictions increase the rate of trade with high-type sellers. Because time frictions also lower rates of trade with low-type sellers, efficiency is not significantly different between Private and Private T, or between Public and Public T.*

Support: The difference in the rate of trade for high-type sellers in *Public* (75%) and *Public T* (25%) is significant ($p = 0.016$), while the difference between *Private* (73%) and *Private T* (50%) is not significant ($p = 0.168$). However, the latter difference is highly significant in a multilevel regression with individual and session random intercepts (not reported). Pooling the competitive bargaining sessions with time frictions and comparing them to the sessions without time frictions shows that the probability of trade for high-type sellers is significantly higher in the former ($p = 0.003$). On the other hand, rates of trade for low-type sellers are higher in the treatments without time frictions ($p < 0.001$). Taken together, time frictions do not lead to significantly different efficiency levels in competitive bargaining ($p = 0.247$).

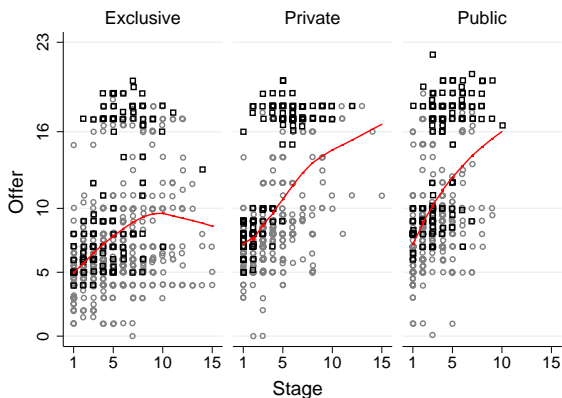
□

Two main conclusions follow from the above results. First, time frictions enable buyers to use price offers to extract information from sellers, which in turn promotes trade with high-type sellers. This conclusion holds independently of whether bargaining is exclusive or competitive. Recall that the only difference between the treatments with and without time frictions is that the breakdown stage is random when there are time frictions and commonly known otherwise. Buyers seem to understand that the breakdown probability allows for effective screening. The second main insight is that competition promotes rates of trade and efficiency compared to exclusive bargaining, irrespective of the transparency of offers. This was expected for private offers. If offers are public, however, the finding is the opposite of what the model predicts.

4.2 Price Sequences and Risk Aversion

Figure 1 presents all offers made by the buyers (grey circles) and the accepted offers (black squares). Price offers follow a steep increase in *Private* and *Public* and increase at a lower rate in *Exclusive*. Only few offers are between 10 and 16, even in *Exclusive* where theory predicts that such offers are common. The typical price sequence in all three treatments corresponds to *threshold screening*: A sequence of offers acceptable only for the low-type seller (i.e., offers between 0 and 10) is followed by a jump to an offer above the high-type seller's reservation

Figure 1: Rejected & Accepted Offers



Rejected offers depicted as grey circles. Accepted offers depicted as black squares. Graphs include smoothed values from locally weighted regressions.

Table 4: Threshold Screening

	Exclusive	Private	Public
Component 1			
Mean	6.38	8.33	9.33
Variance	2.91	1.69	2.06
Frequency	0.45	0.26	0.21
Component 2			
Mean	17.34	17.63	18.38
Variance	2.55	0.72	1.76
Frequency	0.55	0.74	0.79

Two-component finite mixture model for the maximum offer in a bargaining game with standard errors clustered on subjects. Components are assumed to be normal distributions.

value of 16.

Threshold screening can essentially lead to one of two outcomes. Either the probability of an offer above 16 is sufficiently high such that the majority of price sequences end with a high offer. Alternatively, if buyers are reluctant to make high offers, bargaining breaks down before such an offer is made. Let y denote the maximum offer in a given bargaining game and assume that the distribution of y is a mixture of two normal distributions. Given that threshold screening was common, we expect that one component of the mixture model has a mean below 10 and another a mean above 16. It is interesting to see how frequent each outcome occurred in each treatment. Table 4 shows the estimates of the corresponding mixture model. In treatment *Exclusive* 55% of the bargaining games reach a maximum offer centered at 17.34, while 45% have maximum offers centered at 6.38. Competition increases the probability that a bargaining game reaches a maximum offer above 16 to 74% in *Private* and 79% in *Public*. The maximum offers in the lower component are centered at 8.33 in *Private* and 9.33 in *Public*.¹⁴

Threshold screening is in line with the equilibrium predictions in treatment *Private*. In treatment *Exclusive*, however, the unique equilibrium is characterized by gradual screening, involving a number of offers between 10 and 16. Can we reconcile threshold screening with

¹⁴We can also look at the price sequences in the treatments without time frictions: In *Exclusive T* the two components of the mixture model are centered around 5.26 and 11.52, i.e., there is no component for offers above 16, in line with the theoretical prediction that time frictions are essential for screening in bilateral bargaining. The component means are centered at 7.04 and 18.31 in *Private T* and at 8.55 and 17.26 in *Public T* (in both cases the probability of each component is around 50%).

Table 5: The Role of Risk Aversion for Buyer Behavior

<i>Dependent Variable:</i>	<i>Offers</i>	<i>Trade $\theta = L$</i>	<i>Trade $\theta = H$</i>	<i>Profit</i>	<i>Efficiency</i>	
Stage	0.807*** (0.142)					
Private	2.095*** (0.471)	-0.440*** (0.063)	-0.331*** (0.103)	-0.691* (0.381)	0.991 (0.660)	1.307** (0.523)
Public	3.240*** (0.517)	-0.423*** (0.074)	-0.388*** (0.104)	-1.21*** (0.396)	0.919 (0.703)	1.227** (0.603)
Risk Averse (RA)	-1.112*** (0.383)	-0.007 (0.057)	-0.249*** (0.104)	0.778** (0.351)	-0.567 (0.646)	
Private \times RA	1.057** (0.453)	0.017 (0.073)	0.238* (0.136)	-1.069*** (0.366)	0.557 (0.683)	
Public \times RA	0.898** (0.457)	-0.030 (0.101)	0.375*** (0.139)	-0.446 (0.371)	0.539 (0.747)	
Constant	4.442*** (0.456)	0.737*** (0.064)	0.613*** (0.108)	0.939** (0.384)	5.710 (1.001)	5.269*** (0.920)
Period Dummies	✓	✓	✓	✓	✓	✓
Observations	4623	788	403	1191	1191	519
Individuals	216	216	190	216	216	190

Linear multilevel models with individual and session random intercept. Standard errors clustered on independent observations (22 matching groups) in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The reference group is treatment *Exclusive* and $RA = 0$ (low risk aversion).

equilibrium reasoning? An interesting possibility is the existence of obstinate types (Abreu and Gul, 2000). Obstinate types are committed to never raise offers above a certain threshold. Interestingly, rational players may want to mimic obstinate types, a prediction that has been confirmed in the lab (Embrey et al., 2015). In our setting, assuming the presence of obstinate buyers, there exists an equilibrium in which a rational buyer uses threshold screening. He starts by pretending to be an obstinate type who only makes offers below 10 and switches to a high offer when the belief to be matched with a high-type seller is sufficiently high.¹⁵

Table 5 presents multilevel regressions for the following dependent variables: buyers' offer choices, their probability of trade with low-type and high-type sellers, buyers' profits, and efficiency. The main explanatory variables are the treatments and buyers' risk aversion. Risk aversion is captured by the dummy *Risk Averse (RA)*, which is constructed based on the risk elicitation task described in section 3.3.¹⁶ Remarkably, risk aversion is an important

¹⁵The full equilibrium characterization is intricate and beyond the scope of this paper. See Fanning (2014) for a discussion of obstinate types in the presence of incomplete information. An alternative explanation for the flat price sequences in *Exclusive* is that some subjects fail to correctly update their beliefs towards the high-type seller following a rejection, as argued in the literature on cursed beliefs (e.g., Eyster and Rabin, 2005; Esponda, 2008).

¹⁶We find that most subjects are risk averse. Almost all subjects (97%) have a unique switching point from

determinant of behavior in exclusive bargaining but the effects of risk aversion disappear in the presence of competition:

Result 5: *In treatment Exclusive, risk aversion reduces buyers' offers and slows down trade. In treatments Private and Public, risk aversion neither reduces price offers nor trade frequency.*

Support: Column 1 in table 5 shows that offers are higher under competitive bargaining than under exclusive bargaining. Moreover, risk aversion lowers offers in *Exclusive*. This is intuitive. Postponing agreement allows for better screening, thereby reducing the risk of offering a high price to a low-type seller. As a consequence, as shown in column 3, risk averse buyers are less likely to trade with high-type sellers. In contrast, risk aversion does not lead to lower offers or efficiency in competitive bargaining. In fact, in treatment *Public* the probability of trade with a high-type seller is *larger* for risk averse buyers than their less risk averse counterparts ($p = 0.067$). \square

Interestingly, price offers tend to be higher when offers are observable, whereas the opposite should occur in theory. Column 1 in table 5 shows that offers in treatment *Public* are higher than in treatment *Private*, namely by 1.145 points for individuals with a relatively low level of risk aversion ($p = 0.012$) and 0.986 points for the more risk averse group ($p = 0.037$). The fact that offers are higher when they are observable does not lead to a significant difference in the trading rate with high-type sellers between *Public* and *Private* ($p > 0.170$). Nevertheless, it shows that, if anything, observable offers cement the seller's monopoly position.

This begs the question of whether buyers in treatment *Public* who frequently raised prices above 16 made losses. Table 3 presents buyers' payoffs. It shows that on average buyers made a loss of -0.48 points when the seller was a low type. Since this is an average across three buyers, the loss conditional on trading with a low-type seller was $-0.48 * 3 = -1.44$. This loss is offset by the gain in case the seller was a high type such that the overall expected payoff is close to 0. We conclude that for buyers in treatment *Public* increasing offers above 16 (after some initial low offers) or keeping offers below 10 both yielded an expected payoff of approximately 0. In principle, this is in line with theory where the belief after stage 1 should be such that buyers are indifferent between an offer of 16 and a losing offer. However, given that most subjects are risk averse according to our risk elicitation task, in terms of expected utility, a losing offer is more attractive than offering a price of 16 which involves the risk of making a loss. Once again, in line with result 5, the competitive environments seemingly

accepting less risky to rejecting more risky lotteries. Thus, we use the switching point as the risk aversion measure. The risk aversion dummy is created by splitting subjects into two equally large groups. The regression results are robust to using as our risk measure the switching point in the lottery task (i.e., integers 0 to 6) instead of a dummy variable.

eliminate the role of risk aversion.

4.3 Price Transparency: Robustness Checks and An Explanation

The main deviation from theory in our data is the failure of hypothesis 3. High-type sellers trade with a high probability in treatment *Public* while theory predicts a rate of trade of 0. Indeed, in table 3 we have seen that trading rates, efficiency, trading prices, trading stages, opening offers, and profits of buyers and sellers are virtually identical in *Public* and *Private*.¹⁷ In this section, we explore several factors that may help explain this result.

In Hörner and Vieille (2009a) there is an infinite stream of buyers who sequentially meet the seller to make an offer. If an offer is rejected, the respective buyer leaves the market and never comes back to make another offer. In our environment, competition is also sequential but the three buyers take turns in making offers. Theoretically, the detrimental effect of public offers persists in our environment. Behaviorally, however, the possibility that a buyer may come back to make another offer could make a difference. For instance, buyers may try to screen the low-type seller by offering prices above 0 and then, in case of a rejection, trade with the high-type seller three stages down the road. Or, alternatively, buyers may interpret competition as being simultaneous rather than sequential when there are only few buyers, i.e., they may treat the situation as if all buyers make offers at the same time. The lab doesn't allow us to implement an infinite stream of buyers. But we can look at the case with six buyers, where it is much less likely that a given buyer will be able to make multiple offers. Table 6 reports the results for these treatments *Private 6B* and *Public 6B*. In line with the above reasoning, the probability for the first buyer to come back to make a second offer is 49.4% in *Private* and 48.2% in *Public*, while it drops to 22.2% and 9.7% in the respective six-buyers treatments. In spite of this, we find that overall rates of trade and efficiency levels are similar in *Private 6B* and *Public 6B* and also similar to the respective three-buyers treatments. Note that competition is again slightly stronger when offers are observable. This is best seen when looking at the trading prices with low-type sellers, which are higher in *Public 6B* than in *Private 6B* ($p = 0.049$). We conclude that the number of buyers doesn't explain the discrepancies between behavior and theory.

Another possible concern is that offer transparency interacts with how one models time frictions. In three sessions of treatments *Private*, *Private 6B*, *Public*, and *Public 6B*, after the main experiment had been completed, we asked subjects to play five additional bargaining

¹⁷We can exclude the possibility that the absence of an effect of offer transparency is due to subjects not paying enough attention to the history of offers. The history of offers was explicitly discussed in the experimental instructions and prominently displayed on the subjects' computer screens during the experiment.

Table 6: Summary of Outcomes in *6-Buyers* and *Strategy Method* Treatments

Treatment	Seller Type	Trading Rate	Efficiency	Trading Price	Trading Stage	Opening Offer
Private 6B	<i>Overall</i>	0.75 (0.73)	6.89 (6.67)	13.49 (14.33)	3.93 (4.90)	8.13 (10.00)
	<i>High</i>	0.59 (0.63)	4.13 (4.38)	18.65 (16.00)	6.54 (6.40)	8.12 (10.00)
	<i>Low</i>	0.83 (0.78)	8.26 (7.81)	10.91 (13.5)	2.63 (4.15)	8.14 (10.00)
Private Strategy	<i>Overall</i>	0.72 (0.73)	6.64 (6.67)	14.08 (14.33)	3.12 (4.90)	8.88 (10.00)
	<i>High</i>	0.55 (0.63)	3.82 (4.38)	18.09 (16.00)	5.25 (6.40)	8.17 (10.00)
	<i>Low</i>	0.81 (0.78)	8.04 (7.81)	12.07 (13.5)	2.05 (4.15)	9.23 (10.00)
Public 6B	<i>Overall</i>	0.78 (0.28)	7.04 (2.78)	15.56 (-)	3.56 (-)	8.66 (0.00)
	<i>High</i>	0.72 (0.00)	5.04 (0.00)	18.54 (-)	5.28 (-)	8.69 (0.00)
	<i>Low</i>	0.80 (0.42)	8.04 (4.17)	14.07 (0.00)	2.70 (1.00)	8.64 (0.00)
Public Strategy	<i>Overall</i>	0.54 (0.28)	5.33 (2.78)	13.18 (-)	4.76 (-)	7.36 (0.00)
	<i>High</i>	0.11 (0.00)	0.78 (0.00)	19.25 (-)	8.50 (-)	7.14 (0.00)
	<i>Low</i>	0.76 (0.42)	7.61 (4.17)	10.15 (0.00)	2.89 (1.00)	7.46 (0.00)

The table shows averages across bargaining games by seller type. Overall (high- and low-type sellers) averages are calculated by weighing the observed averages by the theoretical probability of the seller types. Data includes observations from bargaining rounds 3-10. Theoretical predictions are given in parentheses.

rounds in which we used discounting instead of a breakdown probability.¹⁸ Efficiency levels in the rounds with discounting are 6.09 in *Private*, 6.66 in *Private 6B*, 6.22 in *Public*, and 6.63 in *Public 6B*. The differences are insignificant. It is possible that participants' behavior in the rounds with discounting was affected by the initial ten rounds of the main experiment. We cannot exclude this possibility. Apart from this limitation, the data from the last five rounds suggests that behavior in competitive bargaining is unaffected by whether time frictions take the form of a breakdown probability or discounting.

Our final set of treatments *Private Strategy* and *Public Strategy* employ the strategy method to learn more about buyers' choices. In *Private Strategy*, we ask buyers to make conditional offers of the form "suppose that stage t is reached and it is your turn to make an offer, what will your offer be?" In *Public Strategy*, buyers made offers conditional not only on the stage t but also on the offer made in stage $t - 1$. Because there is a large number of possible histories, we asked buyers to choose an offer conditional on the offer observed in stage $t - 1$ but not conditional on offers made in all previous stages.¹⁹ The motivation for using the strategy method was to check if at least some buyers in *Public Strategy* play according to

¹⁸In the main part of the experiment, subjects knew that there will be a second part but they didn't know the details of it.

¹⁹We refer to the experimental design section in this article and the instructions in the Online Appendix for more details on how we implemented the strategy method. Notice that subjects choices are incentivised as in the original treatments: In stage 1, offers cannot be conditional (there is no history) and thus the first offer determines which choices will be relevant for payment in all future stages.

the equilibrium predictions (i.e., to see if they choose low offers conditional on observing low previous offers). Strikingly, asking participants to make offers via the strategy method had more far-reaching effects than we anticipated:

Result 6: *Price offers and the rate of trade for high-type sellers are substantially lower in treatment Public Strategy than in treatment Private Strategy.*

Support: Behavior in *Private Strategy* is similar to the one observed in treatment *Private* and hence still in line with the theoretical predictions. However, the transparency of offers now has a large effect. Opening offers as well as the rate of trade with high-type sellers are significantly lower in treatment *Public Strategy* than in treatment *Private Strategy* (non-parametric tests $p = 0.049$, multilevel regressions $p < 0.016$). The rate of trade for high-type sellers is reduced to just 11% in *Public Strategy* and thus close to the theoretical prediction of 0%. Two-component mixture models (similar to the ones presented in table 4) show that only 27% of all price sequences involve an offer above 16 in *Public Strategy* while the majority of the price sequences are centered around a maximum offer of 8.94. In *Private Strategy*, 69% of all price sequences reach an offer which exceeds 16. \square

Implementing the bargaining game via the strategy method thus reveals results that are in line with hypothesis 3. That is, compared to the bilateral bargaining setting, adverse selection effects are mitigated by competition if offers are private but reinforced if offers are public. This begs the question why subjects behave differently in treatment *Public* than in treatment *Public Strategy*. We offer the following explanation.

The equilibrium construction in the competitive bargaining setting with public offers requires buyers to use conditional reasoning of the form: “If I keep my current offer low, the next buyer will do so as well; but if I raise my current offer, the next buyer will react by offering an even higher expected offer.” The results in treatment *Public* show that buyers typically do not behave like this. However, the strategy method makes explicit the importance of conditional reasoning. As a result, observed behavior is closer to the equilibrium predictions. Jehiel (2005) introduces an equilibrium concept—analogy-based expectation equilibrium—motivated by the observation that in complex situations agents often seem to use simplified representations to learn about their environment and other players’ reactions. In particular, players pool together several contingencies in which the other players must move. Applied to our setting, a buyer may believe that other buyers take into account the duration of an ongoing bargaining process (i.e., in which stage they are) but not, or not to a sufficiently large degree, what the previous offers were. In our main treatments *Private* and *Public*, the consequence is that the environment with public offers essentially reduces to the one with private offers. Using the strategy method, on the other hand, helps buyers realize the strategic richness of the environment with public offers.

5 Conclusion

We conduct an experiment examining the efficiency properties of different bargaining institutions in the presence of adverse selection. Our choice of institutions is rooted in the theoretical literature, which shows that whether or not bargaining—defined as the possibility to make repeated offers—can successfully transmit information and thus promote trade depends on the presence of time frictions, the level of competition, and the transparency of offers where transparency refers to whether or not offers are observable among competitors.

The experimental results are qualitatively in line with the theoretical predictions for most treatments. The possibility to bargain leads to increased rates of trade for high-type sellers and, in line with theory, time frictions are shown to be important for this result. Competition between buyers (the uninformed parties) leads to consistently high efficiency levels in our markets, independently of whether offers are private or public. The latter observation that the transparency of offers does not affect efficiency is not in line with theory. We explore this finding in more detail. Remarkably, when we implement the bargaining games via the strategy method, we do observe the predicted negative effects of public offers. That is, making explicit the strategic opportunities linked to public offers allows us to reconcile behavior in the experiment with the general features of the theoretical predictions.

Other factors may help explain our results. Hörner and Vieille (2009b) theoretically explore the consequences of assuming that a seller is myopic with some probability. A myopic seller accepts any offer that exceeds her production costs. As a result, buyers gradually update their beliefs until at some point trade with high-type sellers must occur. This model is consistent with our data in some respects. However, the presence of myopic sellers cannot explain the complete absence of a difference in behavior between treatments *Private* and *Public*. In fact, Hörner and Vieille (2009b) show that if the probability of a seller to be myopic is small the equilibrium is close to the fully rational one.

Intention-based fairness models in the spirit of Rabin (1993) can also explain why trade with high-type sellers may occur even when offers are observable. The reason is that low-type sellers reject offers below, say, a price of 5, but receive a positive utility for offers slightly above 5. If low-type sellers obtain a positive utility in some stage, the bargaining game must eventually lead to offers above 16. However, fairness models are not in line with the observation that opening offers are substantially above 5 and that the seller extracts almost 100% of the gains from trade.

Diamond (1971) has observed that a small search cost can suffice to give the offering party (the buyers in our case) substantial monopoly power. Intuitively, at each possible price level a

buyer could offer slightly less such that the seller still prefers this offer to incurring additional search costs. This unravels until the price is at 0. Previous experiments (e.g. Davis and Holt, 1996) have shown that this result is not typically observed in the lab. However, it is important to realize that in our environment Diamond's result doesn't apply. Indeed, in treatment *Private* there are time frictions (search costs) and yet the price is never at 0. Instead, due to incomplete information about the seller's type, price offers start at the buyers' valuation for the low quality good. Furthermore, in treatment *Public T* there are no search costs and yet the price is at 0. This illustrates that the theoretical predictions for the competitive bargaining settings are not driven by search costs. Instead, the prediction of a price of 0 when offers are observable follows from the inability of the seller to commit to not using positive price offers to solicit even higher offers in the future.

The theoretical literature on adverse selection has demonstrated the importance of time frictions (for bilateral bargaining), competition, and price transparency. Our experiment confirms that these factors are important determinants of market outcomes. At the same time, our results challenge the existing theories to better explain the interaction of risk aversion and competition and calls for more precise modeling of how agents' perceive trading environments, that is, which strategic opportunities they focus on and which parts they overlook. The literature following Jehiel (2005) is an important step in this direction. More generally, the continued interplay between theory and experiments is a promising avenue to improve our understanding of markets in the presence of adverse selection.

References

- Abrams, Eric, Martin Sefton, and Abdullah Yavas**, “An experimental comparison of two search models,” *Economic Theory*, 2000, 16 (3), 735–749.
- Abreu, Dilip and Faruk Gul**, “Bargaining and reputation,” *Econometrica*, 2000, 68 (1), 85–117.
- Akerlof, George A.**, “The market for lemons: Quality uncertainty and the market mechanism,” *The Quarterly Journal of Economics*, 1970, 84 (3), 488–500.
- Blouin, Max R. and Roberto Serrano**, “A decentralized market with common values uncertainty: Non-steady states,” *The Review of Economic Studies*, 2001, 68 (2), 323–346.
- Bochet, Olivier and Simon Siegenthaler**, “Better later than never? An experiment on bargaining under adverse selection,” *International Economic Review*, 2018, 59 (2).
- Cason, Timothy N. and Charles Noussair**, “A market with frictions in the matching process: An experimental study,” *International Economic Review*, 2007, 48 (2), 665–691.
- **and Daniel Friedman**, “Buyer search and price dispersion: a laboratory study,” *Journal of Economic Theory*, 2003, 112 (2), 232–260.
- **and Stanley S. Reynolds**, “Bounded rationality in laboratory bargaining with asymmetric information,” *Economic Theory*, 2005, 25 (3), 553–574.
- Croson, Rachel T.**, “Information in ultimatum games: An experimental study,” *Journal of Economic Behavior & Organization*, 1996, 30 (2), 197–212.
- Davis, Douglas D. and Charles A. Holt**, “Consumer search costs and market performance,” *Economic Inquiry*, 1996, 34 (1), 133–151.
- Deneckere, Raymond and Meng-Yu Liang**, “Bargaining with interdependent values,” *Econometrica*, 2006, 74 (5), 1309–1364.
- Diamond, Peter A.**, “A model of price adjustment,” *Journal of Economic Theory*, 1971, 3 (2), 156–168.
- Embrey, Matthew, Guillaume R. Fréchette, and Steven F. Lehrer**, “Bargaining and Reputation: An Experiment on Bargaining in the Presence of Behavioural Types,” *The Review of Economic Studies*, 2015, 82 (2), 608–631.
- Esponda, Ignacio**, “Behavioral equilibrium in economies with adverse selection,” *American Economic Review*, 2008, 98 (4), 1269–1291.

- Evans, Robert**, “Sequential bargaining with correlated values,” *The Review of Economic Studies*, 1989, *56* (4), 499–510.
- Eyster, Erik and Matthew Rabin**, “Cursed equilibrium,” *Econometrica*, 2005, *73* (5), 1623–1672.
- Fanning, Jack**, “Bargaining, fairness and commitment,” *Unpublished manuscript*, 2014.
- , “Reputational bargaining and deadlines,” *Econometrica*, 2016, *84* (3), 1131–1179.
- **and Andrew Kloosterman**, “A Simple Experimental Test of the Coase Conjecture: Fairness in Dynamic Bargaining,” *Unpublished manuscript*, 2018.
- Fischbacher, U.**, “z-Tree: Zürich toolbox for ready-made economic experiments,” *Experimental Economics*, 2007, *10* (2), 171–178.
- Forsythe, Robert, John Kennan, and Barry Sopher**, “An experimental analysis of strikes in bargaining games with one-sided private information,” *American Economic Review*, 1991, *81* (1), 253–278.
- Fudenberg, Drew and Jean Tirole**, “Game theory, 1991,” *Cambridge, Massachusetts*, 1991, *393* (12), 80.
- Güth, W. and E. Van Damme**, “Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study,” *Journal of Mathematical Psychology*, 1998, *42* (2), 227–247.
- Hörner, Johannes and Nicolas Vieille**, “Public vs. private offers in the market for lemons,” *Econometrica*, 2009, *77* (1), 29–69.
- **and** – , “Public vs. private offers: The two-type case to supplement “Public vs. private offers in the market for lemons”,” *Econometrica*, 2009, *77* (1), 29–69.
- Jehiel, Philippe**, “Analogy-based expectation equilibrium,” *Journal of Economic Theory*, 2005, *123* (2), 81–104.
- Kim, Kyungmin**, “Public offers in the market for lemons with large discounting,” *Economics Letters*, 2015, *130*, 63–65.
- , “Information about sellers past behavior in the market for lemons,” *Journal of Economic Theory* (*forthcoming*), 2016.
- Mitzkewitz, Michael and Rosemarie Nagel**, “Experimental results on ultimatum games with incomplete information,” *International Journal of Game Theory*, 1993, *22* (2), 171–198.

- Moreno, Diego and John Wooders**, “Decentralized Trade Mitigates the Lemons Problem,” *International Economic Review*, 2010, *51* (2), 383–399.
- **and** – , “Dynamic markets for lemons: Performance, liquidity, and policy intervention,” *Theoretical Economics*, 2016, *11* (2), 601–639.
- Nagel, Rosemarie and Ron Harstad**, “Ultimatum Games with Incomplete Information on the Side of the Prosper: An experimental Study,” *Cuadernos de economía: Spanish Journal of Economics and Finance*, 2004, *27* (75), 37–74.
- Rabin, Matthew**, “Incorporating fairness into game theory and economics,” *The American Economic Review*, 1993, pp. 1281–1302.
- Rapoport, Amnon, Ido Erev, and Rami Zwick**, “An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting,” *Management Science*, 1995, *41* (3), 377–394.
- , **James A. Sundali, and Darryl A. Seale**, “Ultimatums in two-person bargaining with one-sided uncertainty: Demand games,” *Journal of Economic Behavior & Organization*, 1996, *30* (2), 173–196.
- Reynolds, Stanley S.**, “Durable-goods monopoly: laboratory market and bargaining experiments,” *The RAND Journal of Economics*, 2000, *31* (2), 375–394.
- Roth, Alvin E. and Francoise Schoumaker**, “Expectations and reputations in bargaining: An experimental study,” *American Economic Review*, 1983, *73* (3), 362–372.
- **and J. Keith Murnighan**, “The role of information in bargaining: An experimental study,” *Econometrica*, 1982, *50* (5), 1123–1142.
- **and Michael W. Malouf**, “Game-theoretic models and the role of information in bargaining,” *Psychological review*, 1979, *86* (6), 574–594.
- Samuelson, William**, “Bargaining under asymmetric information,” *Econometrica*, 1984, *52* (4), 995–1005.
- Siegenthaler, Simon**, “Meet the lemons: An experiment on how cheap-talk overcomes adverse selection in decentralized markets,” *Games and Economic Behavior*, 2017, *102*, 147–161.
- Straub, Paul G. and J. Keith Murnighan**, “An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers,” *Journal of Economic Behavior & Organization*, 1995, *27* (3), 345–364.

Vincent, Daniel R., “Bargaining with common values,” *Journal of Economic Theory*, 1989, 48 (1), 47–62.

Virag, Gabor, “Dynamic adverse selection with many types,” *Working paper*, 2016.

A Appendix: Proofs

Let $\mu^* = (c_H - v_L)/(v_H - v_L)$ be the buyers' belief at which an offer of c_H yields an expected profit of 0, i.e., $v_H\mu^* + v_L(1 - \mu^*) - c_H = 0$. Further, let $\mu_T = c_H/v_H$ be the buyers' belief at which a take-it-or-leave-it offers (occurring in stage T in the finite horizon setting) of 0 and c_H yield the same expected payoff, i.e., $(1 - \mu_T)v_L = \mu_Tv_H + (1 - \mu_T)v_L - c_H$.

A.1 Exclusive Bargaining

The proof of Proposition 1 (treatment *Exclusive*) is given in Deneckere and Liang (2006) Proposition 1, including a proof that the equilibrium is unique.

A.2 Exclusive Bargaining with T stages

In the exclusive bargaining game with T stages (treatment *Exclusive T*), there exists an essentially unique perfect Bayesian equilibrium with the following features: (i) the buyer offers 0 in all stages $t = 1, \dots, T$. (ii) The low-type seller's acceptance probabilities up to stage $T - 1$ are such that the buyer's belief satisfies $\mu_{T-1} \leq \mu_T$, while the buyer's belief must be μ_T in the final stage T . The low-type seller accepts with probability 1 in stage T . (iii) The high-type seller rejects all offers.

Note that there cannot be a stage t at which the buyer strictly prefers to offer c_H . If this were the case, the low-type seller would reject all previous offer below c_H and hence offering c_H in t would yield a negative expected payoff to the buyer (since $v_Hq + v_L(1 - q) - c_H < 0$). Thus, if an offer of c_H is made with positive probability in some stage t , there must be at least one alternative optimal offer below c_H in t . In addition, an offer of c_H in t will only be made if the low-type seller has accepted with positive probability in a previous stage $t' < t$ (i.e., the buyer has updated his belief) where the offer in t' must be strictly larger than 0 (otherwise the low-type seller would reject). Suppose now that there is an equilibrium in which an offer above c_H occurs with positive probability. We have shown that in each stage, it must also be an optimal strategy to offer below c_H . So, the expected payoff of the buyer in such an equilibrium can be expressed as the payoff in case the buyer always chooses an optimal offer below c_H . These offers are rejected by the high-type seller. Since at least one offer (the one in t') exceeds 0, the buyer trades with the low-type seller at a price strictly above 0 with strictly positive probability. It is then obvious that the implied expected payoff for the buyer is strictly below the expected payoff of the strategy to offer 0 in all stages (the low-type accepts the offer of 0 in the last stage with probability 1). Hence, there is no

equilibrium at which c_H is offered with positive probability and the unique optimal strategy for the buyer must be to offer 0 in all stages. This is not surprising as Samuelson (1984) has shown that the buyer's optimal trading mechanism is to make a take-it-or-leave-it offer to the seller, which in the absence of time frictions is equivalent to offering 0 in all stages. The acceptance probabilities of the low-type seller support the buyer's equilibrium offers. In particular, they are such that the buyer's belief is exactly μ_T in stage T such that the buyer is willing to mix between 0 and c_H . The latter is necessary to deter deviations above an offer of 0 in the previous stages: the low-type seller is willing to reject such offers, knowing that in the final stage the buyer will put a positive probability on the offer c_H , and hence the buyer has no incentive to deviate from the zero offer sequence to begin with.

A.3 Competitive Bargaining with Private Offers

We prove Proposition 2 (treatment *Private*). To see why trade with the high-type seller occurs with positive probability, note that the low-type seller will eventually accept an offer with probability 1. If not, then only the losing offer of $c_L = 0$ would be offered, in which case any buyer could profitably deviate by offering just slightly more than 0. It follows that the buyers' belief to face a high-type seller increases over time, converging to 1, and thus there exists a buyer in some stage ℓ with belief μ_ℓ sufficiently large such that offering c_H yields a strictly positive profit. At the same time, delaying such an offer indefinitely implies that the continuation equilibrium payoff converges to 0. This is a contradiction.

We next show that the behavior described in Proposition 2 indeed constitutes a perfect Bayesian equilibrium. Let \tilde{p} be the price a buyer needs to offer such that a low-type seller is indifferent between accepting and rejecting \tilde{p} , anticipating that the same buyer will offer c_H as his next offer n stages in the future and all buyers in between offer c_H with probability λ^* and a losing offer otherwise. We get

$$\tilde{p} = r^n(1 - \lambda^*)^{n-1}c_H + r\lambda^*c_H \sum_{l=0}^{n-2} r^l(1 - \lambda^*)^l = r^n(1 - \lambda^*)^{n-1}c_H + r\lambda^*c_H \frac{1 - r^{n-1}(1 - \lambda^*)^{n-1}}{1 - r(1 - \lambda^*)}.$$

Choose λ^* such that a buyer with belief μ^* who offers \tilde{p} (accepted with probability 1 by the low-type seller) and offers c_H the next time he is called to make an offer has an expected profit of 0. That is, λ^* solves

$$(1 - \mu^*)(v_L - \tilde{p}) + r^n\mu^*(1 - \lambda^*)^{n-1}(v_H - c_H) = 0. \quad (1)$$

Plugging μ^* into (1) shows that when $\lambda^* = 0$, the left-hand side strictly exceeds 0. The

left-hand side is strictly below 0 when $\lambda^* = 1$ (since $\tilde{p} > r^n c_H > v_L$). Hence, there exists $\lambda^* \in (0, 1)$ at which (1) holds.

Finally, choose the probability λ_2 with which the buyer in stage 2 offers c_H such that the expected payoff of the low-type seller is v_L when rejecting the offer she receives in stage 1. Hence, λ_2 solves

$$v_L = c_H (r\lambda_2 + r^2(1 - \lambda_2)\lambda^* + r^3(1 - \lambda_2)(1 - \lambda^*)\lambda^* + r^4(1 - \lambda_2)(1 - \lambda^*)^2\lambda^* + \dots)$$

which simplifies to

$$v_L = rc_H \lambda_2 + \frac{r^2 c_H \lambda^* (1 - \lambda_2)}{1 - r(1 - \lambda^*)}. \quad (2)$$

One can show that $\lambda_2 < \lambda^*$. In particular, if one assumes that $n - 1$ buyers mix according to λ^* and the remaining buyer only makes losing offers (i.e., $n - 1$ buyers choose c_H with probability λ^* and a losing offer otherwise and the remaining buyer chooses an offer of v_L or less), we get an expected future payoff of exactly v_L for the low-type seller who rejects the offer in stage 1. Thus, if all n buyers (instead of just $n - 1$) mix according to λ^* , the expected future payoff of the low-type seller exceeds v_L . It follows that in order for (2) to be satisfied, we need $\lambda_2 < \lambda^*$. But it also follows that, as n grows large, the difference between the case when $n - 1$ and n buyers mix according to λ^* becomes negligible and in fact the expected payoff of the low-type seller is very close to v_L even if n buyers mix according to λ^* . For n sufficiently large, the two probabilities λ_2 and λ^* are therefore arbitrarily close, and since $\lambda^* \in (0, 1)$, there also exists $\lambda_2 \in (0, \lambda^*)$.

Consider now the behavior stated in Proposition 2. Clearly, λ_2 guarantees that the expected payoff of the low-type seller when rejecting in stage 1 equals v_L . It is thus optimal for the low-type seller to mix between accepting and rejecting the offer of v_L in stage 1, and it is also optimal to reject any offer of v_L or below in the next stages. The high-type seller also behaves optimally, given that only an offer above c_H gives her a positive payoff. The buyers in stages $\ell \geq 2$ would need to offer at least \tilde{p} to be accepted by the low-type seller (lower offers would be rejected, followed by the deviating buyer putting a higher probability on c_H in his next turn). But λ^* ensures that \tilde{p} yields an expected profit of 0 and hence is not a profitable deviation. So, for these buyers it is indeed optimal to mix between c_H (with an expected payoff of 0) and a losing offer. The buyer in stage 1 would only need to make an offer of slightly below \tilde{p} to screen out the low-type seller (since $\lambda_2 < \lambda^*$). But we have shown that for n sufficiently large λ_2 is close to λ^* and hence $q < \mu^*$ implies that buyer 1 also has no profitable deviation. We have thus identified an equilibrium. For the parameters in the experiment, we obtain $\lambda^* = 0.236$ and $\lambda_2 = 0.042$. The low-type seller's probability of accepting offer $p_1 = v_L$ is $a_1 = 5/12$ such that $\mu_2 = \mu^* = 6/13$.

It remains to show that the equilibrium is essentially unique. We use a series of steps to make this point.

Step 1: In any equilibrium, the upper bound on beliefs is μ^* .

Suppose to the contrary that there exists a buyer i_ℓ for whom $\mu_\ell > \mu^*$. Let $\bar{\mu}$ be the limit of μ_ℓ , assume that $\bar{\mu} > \mu^*$, and choose a history h_ℓ where μ_ℓ is close to the limit $\bar{\mu}$. We claim that buyer i_ℓ makes a winning offer c_H with probability 1. Because $\bar{\mu} > \mu^*$ buyer i_ℓ 's equilibrium payoff is larger than 0 and hence he will not make a losing offer. Hence, an alternative offer $p_\ell < c_H$ must be accepted with positive probability by the low-type seller. Because μ_ℓ is close to the limit $\bar{\mu}$, offer p_ℓ must be offered with a small probability for the belief of buyer $i_{\ell+1}$ not to exceed the limit $\bar{\mu}$. Since the same is true for buyer $i_{\ell+1}$, the low-type seller expects to receive a winning offer of c_H with a probability close to 1. So, p_ℓ must be close to c_H in order to be accepted and in fact above v_L . But $p_\ell \in (v_L, c_H)$ cannot be optimal, because if the offer is rejected the game ends in the next stage with high probability, and the offer leads to a negative payoff if it is accepted. Hence, buyer i_ℓ makes a winning offer c_H . Notice that the above argument is adapted from Hörner and Vieille (2009b) for the case of an infinite stream of buyers. However, for any r , as long as we choose n large enough, the probability of a buyer to return to make another offer is sufficiently close to 0 such that we can use the same reasoning for finite but large n .

Consider now the last buyer i_t that makes an offer p_t different from c_H with positive probability. By construction this implies that $\mu_{t+1} = \bar{\mu}$. As before, $p_t \in (v_L, c_H)$ must hold since the next offer is c_H . But such an offer yields a negative payoff to buyer i_t . Hence, buyer i_t submits only losing and winning offers such that $\mu_{t+1} = \mu_t$. Hence, $\mu_t = \bar{\mu} > \mu^*$ and thus buyer i_t 's expected payoff is positive, which implies that he makes the winning offer c_H with probability 1, a contradiction with the way we defined buyer i_t . We conclude that in any equilibrium, there is no stage ℓ with $\mu_\ell > \mu^*$.

Step 1 implies that in any equilibrium there exists a buyer $i_{\bar{\ell}}$ with $\bar{\ell} > 1$ such that $\mu_\ell = \mu^*$ for all $\ell \geq \bar{\ell}$, where, once the earliest stage $\bar{\ell}$ for which this is true has been reached, the buyers randomize between the winning offer of c_H and a losing offer.

Step 2: We have $\bar{\ell} = 2$.

Suppose that $\bar{\ell} > 2$. Consider buyer $i_{\bar{\ell}-1}$. By definition, $\mu_{\bar{\ell}-1} < \mu^*$. It must be that buyer $i_{\bar{\ell}-1}$ puts positive probability on an offer that is both acceptable only to the low-type seller and that is accepted by the low type seller with positive probability such that $\mu_{\bar{\ell}} = \mu^*$. Since the expected payoff of any buyer in stages $\ell \geq \bar{\ell}$ is equal to 0 (by the optimality of offering c_H and the definition of belief μ^*), the buyer at $\bar{\ell} - 1$ cannot make a serious offer that exceeds v_L (in case of acceptance, his payoff is negative; in case of a rejection his expected payoff is

0). Hence, in stage $\bar{\ell} - 2$, due to the breakdown probability $1 - r$, the low-type seller must be willing to accept an offer strictly less than v_L . So, the expected equilibrium payoff of the buyer at stage $\bar{\ell} - 2$ is positive. Therefore, the buyer at stage $\bar{\ell} - 2$ does not make a losing offer. He also does not make a winning offer of c_H since $\mu_{\bar{\ell}-2} < \mu^*$. Thus, the low-type seller must be indifferent between accepting and rejecting buyer $i_{\bar{\ell}-2}$'s offer, which is strictly below v_L . But then buyer $i_{\bar{\ell}-2}$ has a profitable deviation to slightly raise his offer to be accepted with probability 1. To prevent such deviations, it must be that $\bar{\ell} = 2$, i.e., there is no stage $\bar{\ell} - 2$.

Step 3: The buyer in stage 1 randomizes between a losing offer in $(0, v_L)$ and an offer of v_L and the low-type seller accepts the offer of v_L such that $\mu_2 = \mu^*$.

We have already shown that $\mu_2 = \mu^*$. Thus, the low-type seller accepts an offer from the first buyer with positive probability smaller than 1. This offer must be v_L , otherwise, if the offer is below v_L , the first buyer could profitably deviate by offering a slightly higher price than the acceptable equilibrium price and force the low-type seller to accept with probability 1. Buyer i_1 's expected payoff from offering a price in the interval (v_L, c_H) is negative, again because the expected payoff is 0 once μ^* has been reached. Hence in any equilibrium, the offer that is accepted in stage 1 must be v_L . Note that this implies that the equilibrium outcome is unique, but the belief $\mu_2 = \mu^*$ can be reached by different combinations of buyer i_1 randomizing between v_L and a losing offer and the low-type seller's acceptance probability of the offer v_L . The acceptance probability a_1 in the proposition is for the case when the first buyer offers v_L with probability 1.

This concludes the proof of Proposition 2. Notice that we have not specified the equilibrium for the case when n is small relative to r . The derivation of the equilibrium in this case is beyond the scope of the present paper. We conjecture that the equilibrium will involve features of the exclusive bargaining equilibrium, including some form of screening. It is also worth stating that such an equilibrium is likely to have an efficiency level that lies in between the case of exclusive bargaining with no competition between buyers and our case where competition is strong with n being large relative to r (here efficiency is the same as with an infinite stream of buyers). Hence, qualitatively, the predictions in Corollary 1 would hold even when n is low relative to r .

A.4 Competitive Bargaining with Private Offers and T stages

In the bargaining game with private offers and T stages (treatment *Private T*), there exists an essentially unique perfect Bayesian equilibrium with the following features: (i) The buyer in stage T randomizes between the offers c_H and 0 with probability λ_T on the offer of c_H such

that $\lambda_T c_H = v_L$. (ii) The buyer in $T - 1$ offers v_L . (iii) The buyers in stages $t < T - 1$ offer v_L or below. (iv) The low-type seller's cumulative acceptance probability up to stage $T - 1$ is such that the buyer's belief satisfies $\mu_{T-1} \leq \mu^*$. (v) In stage $T - 1$ the low-type seller accepts the offer of v_L with a probability such that the buyer's posterior belief is $\mu_T = c_H/v_H$. (vi) In stage T the low-type seller accepts all offers. (vii) Finally, the high-type seller only accepts offers of c_H or higher.

Proof. Consider stage T . Only an offer of 0 (accepted by the low-type) or c_H (accepted by both types) can be optimal for the buyer. Moreover, if buyer i_T offered 0 for sure, buyer i_{T-1} would make an offer slightly above 0 to force acceptance with probability 1 by the low-type seller. This would imply that $\mu_T = 1$, which leads to a contradiction since the last buyer should offer c_H in this case. If buyer i_T offered c_H for sure, the low-type seller, anticipating that an offer of c_H will be made for sure in the last stage, would reject all previous offers below c_H and hence the belief in stage T would be below μ^* . This would contradict the optimality of offering c_H . Hence, the only remaining possibility is that the buyer in stage T randomizes between 0 and c_H . The belief must thus be μ_T such that $\mu_T(v_H - c_H) + (1 - \mu_T)(v_L - c_H) = (1 - \mu_T)(v_L - 0)$, i.e., $\mu_T = c_H/v_H$.

Next consider stage $T - 1$. It must be the case that $\mu_{T-1} < \mu_T$. Otherwise the buyer in stage $T - 1$ would offer c_H , which would imply $\mu_T = 1$ (if $\mu_{T-1} > \mu_T$ offering c_H is strictly optimal; if $\mu_{T-1} = \mu_T$ the low-type seller wouldn't accept an offer of 0 in $T - 1$ since we have shown that c_H follows with strictly positive probability). Hence, the low-type seller must be indifferent between accepting and rejecting the offer p_{T-1} so that the belief moves from a belief below μ_T to μ_T . This implies $p_{T-1} = v_L$. For any lower offer the buyer could slightly increase the offer and be accepted with probability 1. Any offer strictly between v_L and c_H is not profitable for buyer i_{T-1} . Hence, $\lambda_T c_H = v_L$ and $\mu_{T-1} \leq \mu^*$. Otherwise the offer of c_H would yield a positive expected payoff to buyer i_{T-1} , while the equilibrium offer of v_L yields a payoff of 0.

Finally, the behavior in stages $t < T - 1$ cannot involve offers above v_L (such an offer would be accepted with probability 1) and the low-type seller only accepts offers of v_L with a probability such that the belief never exceeds μ^* before stage $T - 1$ is reached. In other words, the equilibrium is essentially unique because the buyers' expected payoff is 0 in stages 1 to $T - 1$ in any equilibrium and in stage $T - 1$ the belief moves to $\mu_T = c_H/v_H$ for any $\mu_{T-1} \in [q, \mu^*]$, followed by an offer of c_H with probability $\lambda_T = v_L/c_H$. \square

A.5 Competitive Bargaining with Public Offers

The proof of Proposition 3 (treatment *Public*) is divided into a series of steps.

Step 1: If μ_l is close enough to 1, then independently of the history h_l , the price offered by buyer i_l is $p_l = c_H$, which is accepted for sure.

To see this, notice that the best-case scenario for buyer i_l when offering below c_H is to get $(1 - \mu_l)(v_L - 0) + r^n \mu_l(v_H - c_H)$, which for μ_l sufficiently large is less than the expected payoff from offering c_H of $(1 - \mu_l)(v_L - c_H) + \mu_l(v_H - c_H)$. Let $\hat{\mu}$ be the infimum over those beliefs μ_l close enough to 1.

Step 2: If $\mu_l \leq \hat{\mu}$, then offer p_l cannot lead to a posterior $\mu_{l+1} \in (\hat{\mu}, 1]$.

By step 1, if $\mu_{l+1} > \hat{\mu}$, buyer i_{l+1} will offer $p_{l+1} = c_H$ (note that all buyers must have the same beliefs if offers are public). Thus, if the posterior is $\mu_{l+1} \in (\hat{\mu}, 1]$, the low-type seller must be willing to accept p_l even when knowing that the next offer will be c_H . This implies $p_l \geq rc_H$. Since $rc_H > v_L$, the offer p_l yields a negative expected payoff and will never be offered, a contradiction.

Step 3: The threshold belief is $\hat{\mu} = \mu^*$.

By the definition of $\hat{\mu}$, for any $\epsilon > 0$, there exists a history h_l such that given $\mu_l = \hat{\mu} - \epsilon$, $p_l < c_H$ is optimal. By step 2, the offer p_l can only lead to $\mu_{l+1} \in [\hat{\mu} - \epsilon, \hat{\mu}]$. This implies that the probability of sale when offering p_l tends to 0 for small ϵ , and so does the expected payoff from offering p_l .²⁰ So, because the buyer must not strictly prefer to make offer $p_l = c_H$ at belief $\hat{\mu} - \epsilon$, we need $v_H(\hat{\mu} - \epsilon) + v_L(1 - \hat{\mu} + \epsilon) - c_H \leq 0$, which implies $v_H\hat{\mu} + v_L(1 - \hat{\mu}) - c_H \leq 0$ as ϵ goes to 0. Conversely, because offer c_H is optimal at belief $\hat{\mu} + \epsilon$, it must also be that $v_H(\hat{\mu} + \epsilon) + v_L(1 - \hat{\mu} - \epsilon) - c_H \geq 0$ for any $\epsilon > 0$ and thus $v_H\hat{\mu} + v_L(1 - \hat{\mu}) - c_H \geq 0$. Hence, $v_H\hat{\mu} + v_L(1 - \hat{\mu}) - c_H = 0$ which is the definition of μ^* .

Step 4: Suppose that given history h_l , we have $\mu_l < \mu^*$ and the equilibrium is such that p_l leads to $\mu_{l+1} = \mu^*$. Then all subsequent offers are equal to 0, i.e., $p_t = 0$ for all $t \geq l+1$.

Suppose to the contrary that not all future offers are equal to 0. Then the buyer in stage l must offer $p_l > 0$. Offering $p_l = 0$ leads to a rejection with probability 1, because the seller knows that an offer above 0 will be made with strictly positive probability in at least one of the subsequent stages. Let \bar{p} be the supremum over all such offers $p_l > 0$. Choose an ϵ

²⁰Notice that it could be that buyer i_l makes offer p_l to induce a finer screening in the hope of trading in stage $l+n$, or $l+2n$, etc. with a higher certainty to be facing a high-type seller. A simple way to show that this cannot be the case is to observe that because the probability of sale is very small, the increase in the belief toward the high-type seller in case of rejection does not justify the waiting cost due to r .

small enough such that $\bar{p} - 2\epsilon > 0$ and choose the history h_l such that $p_l > \bar{p} - \epsilon$. Suppose that buyer i_l deviates and offers $\bar{p} - 2\epsilon < p_l$. If $\mu_{l+1} \geq \mu^*$ (after the deviation), this is a profitable deviation since the price offered is lower than p_l and the probability of acceptance by the low-type seller has weakly increased. If $\mu_{l+1} < \mu^*$, the low-type seller gets $\bar{p} - 2\epsilon$ when accepting the offer and cannot hope for more than an expected payoff of $r\bar{p}$ when rejecting the offer. Since the low-type seller is supposed to be willing to reject, this is a contradiction for ϵ small enough.

Step 5: If $\mu_l < \mu^*$, then $p_l = 0$ and $\mu_{l+1} = \mu^*$.

By step 2 we have $\mu_{l+1} \leq \mu^*$. Suppose first that $\mu_{l+1} < \mu^*$. The offer p_l must be below c_H , as an offer of c_H would give a negative expected payoff to the buyer. The offer $p_l < c_H$ cannot be accepted with probability 1 by the low-type seller, because the belief cannot move beyond μ^* . Next, p_l can also not be rejected with probability 1. To see this, note that then there must be some offer $p_t > p_l$ with $t > l$ that occurs with positive probability along the equilibrium path and is accepted with a probability such that the buyers' belief moves to $\mu_{n+1} = \mu^*$ (clearly, if belief μ^* is never reached, any buyer could profitably offer $v_L - \epsilon$, which would be accepted by a low-type seller). Further, $p_t \leq v_L$ to ensure a non-negative expected payoff of buyer i_t . But then buyer i_l is not willing to make the losing offer p_l , because a deviation to $p_l = p_t - \epsilon > rp_t$ is accepted for sure by the low-type seller in stage l . Finally, it could be that the low-type seller is indifferent between accepting and rejecting p_l . But then the buyer could also deviate and slightly raise the offer to guarantee acceptance by the low-type seller at a negligible price increase. This exhausts all possibilities.

So it must be that $\mu_{l+1} = \mu^*$. In this case, step 4 implies that the low-type seller cannot hope for more than an offer of 0 in the future. Hence, the offer that moves the belief to μ^* must be 0 as well (any higher offer would be accepted with probability 1).

This proves Proposition 3: all offers are 0 and the belief jumps to μ^* in stage 1. It is clear that the low-type seller is indifferent between accepting and rejecting the zero offers and hence her behavior is optimal. The high-type seller always rejects, since the offers are below her reservation cost. For the buyers, the equilibrium prescribes that any deviation from the zero price sequence to an offer of $p' \in (0, v_L)$ is deterred by the next buyer who would observe the deviation and mix between 0 and c_H (at belief μ^* , c_H yields an expected payoff of 0) with probability x on the offer c_H such that $xrc_H \geq p'$ and hence the seller would reject p' .

A.6 Competitive Bargaining with Public Offers and T stages

In the bargaining game with public offers and T stages (treatment *Public T*), there exists an essentially unique perfect Bayesian equilibrium with the following features: (i) In each stage $t = 1, \dots, T$, the offer is $p_t = 0$. (ii) The low-type seller's acceptance probabilities up to stage $T - 1$ are such that the buyers' belief satisfies $\mu_{T-1} \leq \mu^*$. (iii) In stage $T - 1$ the low-type seller accepts the offer of 0 with a probability such that the buyer's posterior belief is $\mu_T = c_H/v_H$. (iv) In stage T the low-type seller accepts the offer of 0 with probability 1. Finally, the high-type seller rejects all offers of 0.

Proof. For the same reasons as in *Private T*, the buyer in stage T must have a belief of $\mu_T = c_H/v_H$ that makes him indifferent between offering 0 and c_H . Again as in *Private T*, the offer in $T - 1$ must be accepted with positive probability and the belief must satisfy $\mu_{T-1} \leq \mu^* < \mu_T$. That is, the low-type seller in $T - 1$ must be indifferent between accepting and rejecting. In contrast to *Private T*, the offer p_{T-1} cannot be v_L . Because offers are observed, *any* (on and off-equilibrium) offer in stage $T - 1$ must be accepted by the low-type seller such that $u_T = c_H/v_H$. Only then the buyer in stage T can mix between 0 and c_H , and he must be willing to do so even after an off-equilibrium move. Hence, the buyer in $T - 1$ offers 0, maximizing his payoff given that the acceptance probability is the same for all offers below rc_H . It follows that the offer in stage T is also equal to 0. But then the offer in stage $T - 2$ must be equal to 0 as well as otherwise it would be accepted with probability 1 by the low-type seller. Repeating this argument shows that all offers must be 0. Any deviation to an offer of $p' \in (0, v_L)$ is deterred by the last buyer putting positive probability on offer c_H . The equilibrium is essentially unique: trade can occur with any of the buyers in stages 1 to $T - 1$ at a price of 0 as long as $\mu_{T-1} \in [q, \mu^*]$, but in any equilibrium all offers are 0, the high-type seller doesn't trade, and the low-type seller trades with probability 1, where a significant portion of the acceptance probability occurs in stage T . \square