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

Olivier Bochet† Simon Siegenthaler‡

June 22, 2021

Abstract

In markets with asymmetric information, where equilibria are often inefficient, bargaining can help promote welfare. We design an experiment to examine the impact of competition and price transparency in such settings. Consistent with the theoretical predictions, we find that competition promotes efficiency if bargainers cannot observe each other's price offers. Contrary to the predictions, however, the efficiency-enhancing effect of competition persists even when offers are observable. We explore different behavioral explanations for the absence of a detrimental effect of price transparency. Remarkably, implementing the strategy method improves subjects' conditional reasoning, delivering the predicted loss in efficiency when offers are observable.

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 four anonymous referees for helpful comments. We also thank seminar participants at NYU Abu Dhabi, Singapore Management University, University of Chile, the 2018 AEA meeting in Philadelphia and the 2019 Bargaining: Experiments, Empirics, and Theory (BEET) workshop in Oslo.

†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 they include 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 a risk that no better offers will be received 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; Fuchs, Öry and Skrzypacz, 2016) 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; Fuchs et al., 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

[1] It 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).

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 or may not be revealed.

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 *do* observe each other's offers. For each of the main treatments, we run a treatment with time frictions, in particular, 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 in 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 sellers with potentially high benefits from trade are less willing to postpone agreement than sellers with lower benefits, or equivalently, they are more willing to accept low prices. This allows the buyers to use specific price sequences to screen the different types of the seller. 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.

Hypotheses 2 and 3 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 relative to bilateral bargaining. Remarkably, 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; high offers would be observed by the other buyers and countered by even higher offers. Anticipating these counter-offers, 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 *qualitatively* (we are less interested in testing the exact quantitative predictions of the model). 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. 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. Is a seller better off when revealing information about past offers than when keeping offers private? Finally, this article contributes to the bargaining literature broadly, which has been successful at combining theoretical and experimental work to generate new insights into human behavior. We follow this tradition by experimentally exploring factors—time frictions, competition, and price transparency—that have been deemed relevant by 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 show that competitive bargaining leads to high levels of efficiency. This result holds irrespective of whether offers are private or public. Hence, contrary to the theoretical predictions, 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 data do not bear out the predicted differences between the settings with public and private offers.

To that end, we present a range of robustness checks for the finding that the transparency of offers does not have the anticipated effects. The original theoretical result of Hörner and Vieille (2009a) assumed an infinite stream of buyers. We show that their theoretical results continue to hold in our setting with only three buyers, taking turns in making offers. However, it is instructive to see behavior if we bring the experimental setting closer to the original model. To explore this, we implement treatments with six buyers instead of three buyers. We find that behavior remains largely the same as in the three-buyer treatments, in particular, the transparency of price offers still doesn't affect 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.

We conclude that the absence of a transparency effect is a robust phenomenon, indicating that the theoretical predictions for public offers are too subtle to be picked up by humans (at least in a laboratory environment). 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, in an analogy-based expectation equilibrium subjects bunch together decision nodes that differ only in terms of previously observed price offers.[2] To provide a test for this conjecture, we ran a set of treatments in which subjects make their decisions via the strategy method.

The strategy method highlights that buyers' offers as well as a seller's acceptance decisions can be conditional on the history of play. In particular, instead of asking buyers to enter their offers stage by stage, all decisions are made at the start of a bargaining game and conditional on other players' choices. Note that with private offers subjects can condition their choices only on the period, while with public offers subjects can also condition behavior on offers made in 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. This finding is potentially informative for the development of new and more accurate theories of bargaining that take into account the strategic complexity of an environment.

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

[2]Analogy-based expectation equilibrium is different from Eyster and Rabin (2005)'s cursed equilibrium, but see Miettinen (2009) for a discussion of the link between the two.

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.[3]

There is a well-established experimental literature on bargaining with incomplete information.[4] A set of related studies are Rapoport, Erev and Zwick (1995), Reynolds (2000), Cason and Reynolds (2005) and Fanning and Kloosterman (2018) who 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. 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, unique to our study is the focus on competition and price transparency. Other related experiments are Forsythe, Kennan and Sopher (1991) and Camerer, Nave and Smith (2018) who test the explanatory power of truth-telling constraints in free-form bargaining.

There is another related 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.

Finally, our experiment contributes to the literature on price setting in decentralized markets with search costs. 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

[3]Discounting 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.

[4]See 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. For some recent insightful studies on credible claims in bargaining see Gächter and Riedl (2005), Karagözoğlu and Riedl (2014) and Bolton and Karagözoğlu (2016).

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. Siegenthaler (2017) shows that cheap-talk can be informative in such settings, thus representing an alternative to repeated offers in helping alleviate adverse selection.

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 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 $c_H > 0$. A buyer's valuation is 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 $q < (c_H - v_L)/(v_H - v_L)$. This condition 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 rejected, it is buyer 3's turn. If buyer 3's offer is also 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 bargaining process can also break 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 are interested in situations where players are relatively 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

Here, we 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 $(p_1, p_2, \ldots, p_{\bar{T}}) = (r^{\bar{T}-1}c_H, r^{\bar{T}-2}c_H, \ldots, c_H)$ 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 $t = 1, \ldots, \bar{T} - 2$ such that the buyer is indifferent between offering p_t and p_{t+1}. In stage $\bar{T} - 1$, the low-type seller accepts $p_{\bar{T}-1} = rc_H$ with probability 1.

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.[5,6] It is worth noting that the equilibrium in exclusive bargaining implies a positive probability of trade with high-type sellers. It is also the case, however, that bargaining introduces a risk of trade failure with low-type sellers. 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 (Bochet and Siegenthaler, 2018).

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

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

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

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

[5]For 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.

[6]The 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.

For stage $\ell = 2$, the probability λ_2 on the winning offer solves $v_L = rc_H\lambda_2 + [r^2 c_H \lambda^*(1-\lambda_2)]/[1-r(1-\lambda^*)]$.

The proof of Proposition 2 can be found in Appendix I (available at the journal website). 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 that the seller is a high type equals μ^*. At this belief, the winning offer c_H yields 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, 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.[7]

An important consequence of Proposition 2 is that with private offers trade with the high-type seller occurs with positive probability. If buyers can observe each other's offers (public offers), the equilibrium looks remarkably different: in particular, high-type sellers never trade.

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

[7] In Appendix I, we show that the equilibrium in Proposition 2 is unique for n sufficiently large relative to r (the incentive to screen are then sufficiently small). It is worth noting that the proof does not exclude the possibility that for some combinations of n and r the equilibrium in Proposition 2 is one among several.

The proof can be found in Appendix I. 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 the other buyers' expected offers. As a result, the seller is willing to accept 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 be supported in equilibrium, as 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 do not trade.

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

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

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

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 a 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 given in Deneckere and Liang (2006). Intuitively, as r approaches 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 efficiency 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.

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 $\{0, 0.01, 0.02, ..., 23\}$. 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]

[8] For 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.

Table 1: Experimental Design

Treatment	Subjects[a]	Competition	Transparency	Time Friction	Strategy Method
Exclusive	48 (8)	1 Buyer	–	$r = 0.9$	No
Exclusive T	48 (8)	1 Buyer	–	Known T	No
Private	84 (7)	3 Buyers	No	$r = 0.9$	No
Private T	36 (3)	3 Buyers	No	Known T	No
Private 6B	63 (3)	6 Buyers	No	$r = 0.9$	No
Private Strategy	36 (3)	3 Buyers	No	$r = 0.9$	Yes
Public	84 (7)	3 Buyers	Yes	$r = 0.9$	No
Public T	36 (3)	3 Buyers	Yes	Known T	No
Public 6B	63 (3)	6 Buyers	Yes	$r = 0.9$	No
Public Strategy	36 (3)	3 Buyers	Yes	$r = 0.9$	Yes

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.

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.

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 less likely that a buyer will be able to make more than one offer than in the treatments with three buyers.

Strategy Method: In treatments *Private Strategy* and *Public Strategy* we elicited participants' choices via the so-called strategy method. 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 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. Given these decisions, trade occurs as soon as a buyer's offer for a given stage exceeds the seller's minimum acceptable offer for that stage. Typically this would not happen in the first stage, and hence bargaining breakdowns are still possible. A similar procedure was used in treatment *Public Strategy*. There, however, past offers are observable and thus buyers and sellers could choose their offers and minimum acceptable offers for stage t conditional on the offer observed in stage $t-1$.[9] Notice that the offer in stage 1 is always made unconditionally, as there is no previous offer.

The rationale for the strategy method treatments is that it allows subjects to clearly see the strategic opportunities offered by the different bargaining environments. Given that the derived equilibrium strategies are based on sophisticated contingent reasoning, using the strategy method may increase the probability of observing behavior that is close to the theoretical predictions.

3.2 Hypotheses

For our experimental hypotheses we focus on 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 over all market participants.

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

[9] In principle, buyers and sellers could base their choice for stage t on all previous offers in stages $1, 2, \ldots, t-1$. 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 grouped 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	$Pr(Trade\|\theta = L)$	$Pr(Trade\|\theta = H)$	*Efficiency*[a]
Exclusive	0.74	0.48	6.05
Exclusive T	1	0	6.67
Private	0.78	0.63	6.67
Private T	1	0.63	8.13
Public	0.42	0	2.78
Public T	1	0	6.67

(a) Efficiency equals the sum of ex-ante expected payoffs over 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*.

reached. The efficiency level is 6.04, falling short of first-best efficiency $(1-q)v_L+q(v_H-c_H) = 9$ due to the delay before trade occurs.

In the Appendix I, we characterize the equilibrium for treatment *Exclusive T*. The equilibrium is essentially unique. The buyer offers 0 in all stages. The low-type seller accepts for sure in stage T. In contrast to treatment *Exclusive*, high-type sellers do not 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 in stage 1 offers 10, and 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. In particular, the buyer in stage 2 offers 16 with probability 0.04 and all subsequent buyers offer 16 with probability 0.24. Both seller types reject all offers below 16 from stage 2 onwards. Notice how competition drives up prices (buyers 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:

[10]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.

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 a trading environment. 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 prices from increasing in treatment *Public*.[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 replicate and better explain the results from the first experiments. The second wave was done at the University of Valencia with a total of 270 participants.[12] 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 Appendix II (available at the journal website), we show that behavior does not differ between the two locations, looking at average opening offers, accepted prices, rates

[11]In Appendix I 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. Behaviorally, however, the presence of time frictions may be important even with competition. For this reason we chose to also run treatments *Private T* and *Public T*.

[12]All data can be found in Bochet and Siegenthaler (2020).

of trade with both seller types as well as efficiency. We are therefore able to establish our main results on offer transparency for two separate subjects pools. In our analysis, we will pool 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 Appendix II. Each subject participated in only one treatment. 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 comprehension questions. There were ten rounds, that is, the bargaining games described above were played ten times. 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.[13]

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 did not 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 risk aversion. Finally, Section 4.3 presents robustness checks and the results from the sessions implementing the strategy method.

[13]The data show 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 (no repeated game effects). 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 Appendix II.

All non-parametric tests are based on matching group averages as the unit of observation. We use Wilcoxon-Mann-Whitney tests for between-treatment comparisons and Wilcoxon signed-rank tests for within-treatment comparisons. 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 of 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 consistent with Hypothesis 1. □

The experimental outcomes in treatments *Exclusive* and *Exclusive T* are close to the theoretical predictions both in terms of efficiency and rates of trade. 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.*

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	*Overall*	0.61 (0.65)	5.63 (6.05)	11.46 (12.64)	5.06 (5.42)	4.97 (7.65)	1.75 (0.94)	3.95 (5.10)
	High	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)
	Low	0.70 (0.74)	6.95 (7.39)	8.25 (10.95)	3.93 (4.14)	4.95 (7.65)	1.21 (-0.26)	5.74 (7.65)
Exclusive T	*Overall*	0.65 (0.67)	6.34 (6.67)	10.29 (–)	8.23 (–)	4.03 (0.00)	2.68 (6.67)	3.75 (0.00)
	High	0.20 (0.00)	1.38 (0.00)	17.96 (–)	9.67 (–)	3.95 (0.00)	1.78 (0.00)	-0.13 (0.00)
	Low	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	*Overall*	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)
	High	0.73 (0.63)	5.09 (4.38)	17.63 (16.00)	7.00 (6.40)	7.25 (10.00)	1.49 (1.46)	0.71 (0.00)
	Low	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	*Overall*	0.83 (0.88)	7.83 (8.13)	12.92 (12.88)	7.14 (9.48)	6.77 (10.00)	0.43 (0.19)	6.72 (6.67)
	High	0.50 (0.63)	3.50 (4.38)	18.32 (16.00)	9.64 (10.00)	6.24 (10.00)	1.45 (1.46)	-0.31 (0.00)
	Low	1.00 (1.00)	10.00 (10.00)	10.23 (11.31)	5.89 (9.22)	7.03 (10.00)	-0.07 (-0.44)	10.23 (10.00)
Public	*Overall*	0.76 (0.28)	6.89 (2.78)	14.09 (–)	3.94 (–)	7.21 (0.00)	0.10 (0.93)	6.62 (0.00)
	High	0.75 (0.00)	5.22 (0.00)	18.51 (–)	5.63 (–)	7.41 (0.00)	1.26 (0.00)	1.52 (0.00)
	Low	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	*Overall*	0.74 (0.67)	7.11 (6.67)	12.15 (–)	7.97 (–)	6.23 (0.00)	0.40 (2.22)	6.03 (0.00)
	High	0.25 (0.00)	1.75 (0.00)	17.67 (–)	11.50 (–)	6.03 (0.00)	0.81 (0.00)	-0.27 (0.00)
	Low	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)

Data include observations from bargaining rounds 3-10 and exclude cases where high-type sellers accepted prices below 16. 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% ($p = 0.017$). □

Result 2 is in line with the theoretical prediction that competition promotes efficiency if offers are unobservable and confirms Hypothesis 2. Let us now turn toward 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. These observations are confirmed by the results of the regression model in the last column of Table 5. The observed trading prices also do not 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 ($p = 0.017$). □

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 price 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 significantly affect 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

This section stresses two results. First, price sequences look similar in *Private* and *Public*. Thus, the absence of a transparency effect is not just observed in terms of rate of trade and efficiency but also in terms of subjects' initial offers and how offers change over time. Second, risk aversion leads to lower offers and slower trade in *Exclusive* but this effect disappears when bargaining is competitive (independent of the observability of offers).

Figure 1: Rejected & Accepted Offers

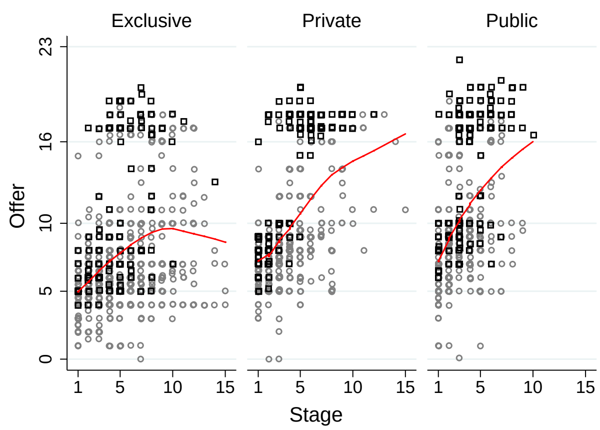

Rejected offers depicted as grey circles. Accepted offers depicted as black squares. Graphs include smoothed values from locally weighted regressions.

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.

Figure 1 presents all offers made by the buyers, rejected offers (grey circles) and 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 value of 16.[14]

Threshold screening can lead to two distinct outcomes. First, the probability that buyers offer prices above 16 is sufficiently high such that the majority of price sequences ends 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

[14]It is worth noting that in treatment *Private* many offers are strictly below 10 in the first two or three stages. This is best explained by social preferences and the vast literature on fairness emphasizing that many people deviate from standard rational behavior in an attempt to avoid outcomes in which one side extracts the lion's share of the trade surplus.

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*.[15]

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 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.[16]

Table 5 presents multilevel regressions exploring the effect of risk aversion 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.[17] We find that risk aversion is an important 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.*

[15]Price sequences in the treatments without time frictions are as follows. 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%.

[16]The 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).

[17]Almost all subjects (97%) have a unique switching point from accepting less risky to rejecting more risky lotteries. 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.

Table 5: Risk Aversion and Buyer Behavior

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

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, under competitive bargaining, risk aversion does not lead to lower offers, rates of trade, or efficiency. □

An interesting explanation why risk aversion affects behavior differently depending on the competitiveness of the environment is that it is not a fixed trait across contexts. For instance, Barseghyan, Prince and Teitelbaum (2011) find that most people exhibit greater risk aversion in their home deductible choices than their auto deductible choices; for a general discussion, see Cox and Harrison (2008). A more formal explanation can be given for the special case where risk preferences are constant across buyers. Recall that the equilibrium in *Private* is such that the low-type seller accepts the offer of the first buyer with a probability such that the buyers' belief after observing a rejection solves $\mu(v_H - c_H) + (1 - \mu)(v_L - c_H) = 0$. This belief ensures that all future buyers are willing to mix between offering c_H and a losing offer. With risk averse buyers, the same equilibrium construction applies, except that the buyers' belief needs to satisfy $\mu u(v_H - c_H) + (1-\mu)u(v_L - c_H) = 0$, where $u(\cdot)$ is the buyers' utility function. The equilibrium value of μ must be larger under risk aversion than under risk-neutrality, thus requiring a larger acceptance probability of the low-type seller in the first period. The behavior of the buyers, however, does not depend on risk aversion: at equilibrium the relative reluctance of risk-averse buyers to make high offers is offset by their higher belief to face a high-type seller.

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* ($p > 0.295$). 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.

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*.[18] In this section, we explore several factors that may help explain this result.

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

Data include observations from bargaining rounds 3-10 and exclude cases where high-type sellers accepted prices below 16. Theoretical predictions are given in parentheses.

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, three buyers take turns in making offers. Theoretically, the detrimental effect of public offers persists but, behaviorally, the possibility that a buyer may come back to make another offer could make a difference. For instance, buyers may interpret competition as being simultaneous rather than sequential, i.e., they may treat the situation as if all buyers make offers at the same time. To explore this, we ran treatments 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 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 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 does not explain the absence of an effect of price transparency.

Another possible concern is that offer transparency interacts with how one models time frictions. To explore this, in three sessions of *Private*, *Private 6B*, *Public* and *Public 6B*, after the main experiment had been completed, subjects played five additional bargaining rounds in which we used discounting instead of a breakdown probability. 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*. These differences are insignificant, suggesting that behavior is unaffected by whether time frictions take the form of a breakdown probability or discounting.

We have established that the absence of a price transparency effect is a robust phenomenon. This indicates that the theoretical predictions for public offers are too subtle to be picked up by humans, at least in a laboratory environment. To understand why, it is useful to consider an equilibrium concept called analogy-based expectation equilibrium (Jehiel, 2005). This equilibrium concept is motivated by the observation that in complex situations agents often 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 but not, or not to a sufficiently large degree, what the previous offers were. If this is the case, the environment with public offers essentially reduces to the one with private offers.

Analogy-based expectation equilibrium suggests that helping buyers understand the strategic richness of the environment with public offers may bring behavior closer to the theoretical predictions. One way to achieve this is to use the strategy method, which explicitly asks subjects to make decisions conditional on other subjects' behavior. Our final two treatments *Private Strategy* and *Public Strategy* implement the strategy method. In *Private Strategy*, we asked 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?" Sellers also made conditional choices by stating a minimum acceptable offer. In *Public Strategy*, buyers made offers conditional not only on stage t but also on the offer made in stage $t-1$.[19]

Strikingly, the strategy method leads to significant differences between the public and private offer settings:

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.*

[19]We refer to the experimental design section in this article and the instructions in Appendix II for more details on how we implemented the strategy method. Notice that subjects choices are incentivized 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.

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 strong effect on behavior. 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 that exceeds 16. □

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.*[20] We next take a more detailed look at behavior in *Public Strategy*.

Opening offers in treatment *Public Strategy* are around 7.3 (see Table 6). From there, subjects raised their offers to a price of 10 within the next few stages. However, buyers are less likely to increase offers to reach a price of 16 and above than in the other treatments. To confirm this, we ran a probit regression with the dependent variable being 1 if an offer exceeds $c_H = 16$ and 0 otherwise and independent variables treatment (*Private Strategy* and *Public Strategy*), period (a dummy variable for each repetition of the bargaining game), stage (the number of an offer in a given bargaining game), and the offer made in the previous stage. We find that, at the mean, buyers were 11.7 percentage points less likely to make an offer of 16 or above in treatment *Public Strategy* than in treatment *Private Strategy* (p<0.001, bootstrapped standard errors).

The lower inclination of buyers to raise offers in *Public Strategy* cannot be explained through seller behavior. One way to show this is to look at screening effectiveness, i.e., the increase in the probability that the seller is a high type over time. In stages 1 to 7, respectively, the probabilities are: 0.35, 0.39, 0.43, 0.50, 0.65, 0.77, 0.88 in *Public*, 0.35, 0.42, 0.48, 0.53,

[20] Brandts and Charness (2011) present a survey of the literature regarding whether the strategy method leads to different experimental results than the direct-response method. They find that out of twenty-nine studies, sixteen find no difference between the two approaches, four do find differences, and nine comparisons find mixed evidence. They also find that in no case there is a treatment effect with the strategy method if it is not also observed with the direct-response method. Interestingly, our study is such a case, as we show that the strategy method can help players use contingent reasoning. This is in line with the recent literature on "obviously strategy-proof mechanism," stressing that an equilibrium should be expected to be played when equilibrium strategies can be recognized by cognitively limited agents but may not be observed otherwise (Li, 2017).

0.68, 0.81, 0.92 in *Private*, 0.36, 0.41, 0.46, 0.56, 0.67, 0.78, 0.82 in *Public Strategy* and 0.36, 0.46, 0.59, 0.61, 0.71, 0.69, 0.77 in *Private Strategy*. As can be seen, screening of low-type sellers was equally successful in the treatments with public and private offers. High-type sellers behaved as expected, that is, on average they chose to accept offers of around 17 (their reservation cost is 16).

These observations imply that the mechanism leading to low rates of trade with high-type sellers in *Public Strategy* is not the one proposed by the theoretical model in a strict sense. In particular, we do not find that observable offers lead to less effective screening than unobservable offers and we do not observe offers close to 0. Nonetheless, the two main predictions of the public offers model hold up: price sequences are flatter than with private offers and trade with high-type sellers is less common. The strategy method makes explicit the possibility of 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 a high offer, too." Indeed, buyers typically chose offers above 16 conditional on a previous offer of 16 and above. Taken together, this suggests that buyers in *Public Strategy* were reluctant to raise offers, because they realized that the next buyer will react by making a high offer as well; i.e., offering more than 16 is too risky, because in case of a low-type seller it results in a loss while in case of a high-type seller there will be fierce competition from other buyers.

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 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 observation that price transparency does not affect efficiency is not in line with the theoretical predictions. We provide robustness checks for this finding and provide an explanation: remarkably, when we use 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.

Diamond (1971) observed that a small search or time cost between periods can suffice to give the offering party (the buyers in our case) substantial monopoly power. Intuitively, at each possible price level a buyer can offer slightly less such that the seller still prefers accepting this offer to incurring additional search costs. This unravels until the price is at 0. A related force is present in our environment when offers are public. Thus, the nature of competition in our setting is relatively weak. In this context, following theoretical work by Fuchs et al. (2016), an interesting extension of our experiment would be to allow for intra-period competition, that is, to consider the case of multiple buyers in every period.[21]

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. In a model without fairness preferences, there exists an offer such that the low-type seller receives a utility of 0 when accepting. The same is true for models with outcome-based fairness preferences. However, if the perception of fairness is based on intensions, a low-type sellers may reject offers below a price that leads to a 50-50 split of the surplus, considering such an offer as unfair, but at the same time receive a strictly positive utility when accepting a 50-50 split. In other words, there is no offer that, when accepted, results in a utility of 0 for the buyer. Without such an offer, the equilibrium construction in our treatments with observable offers breaks down and trade with high-type sellers occurs with positive probability. On the other hand, the fact that sellers extract almost 100% of the gains from trade suggests that fairness motives do not play a dominant role in our experiment.

[21] Fuchs et al. (2016) show that with intra-period competition, trade occurs gradually over time even if offers are public (i.e., there is no point in time after which no more offers are ever accepted, a situation that can occur in our setting). Hence, Diamond paradox effects are eliminated. Nevertheless, under natural assumptions, markets with private offers are still predicted to Pareto-dominate transparent markets.

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 existing theories to better explain the interaction between risk aversion and competition. They also document the need for modeling how agents perceive trading environments, that is, which strategic opportunities humans focus on and which parts they neglect. The literature following Jehiel (2005) is an important step in this direction. The continued interplay between bargaining theory and experiments is a promising avenue to improve our understanding of markets, in particular in the presence of asymmetric information.

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.

Barseghyan, Levon, Jeffrey Prince, and Joshua C Teitelbaum, "Are risk preferences stable across contexts? Evidence from insurance data," *American Economic Review*, 2011, *101* (2), 591–631.

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).

_ **and** _ , "Data and code for: Competiton and Price Transparency in the Market for Lemons: Experimental Evidence," *American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/EV1*, 2020.

Bolton, Gary E. and Emin Karagözoğlu, "On the influence of hard leverage in a soft leverage bargaining game: The importance of credible claims," *Games and Economic Behavior*, 2016, *99*, 164–179.

Brandts, Jordi and Gary Charness, "The strategy versus the direct-response method: a first survey of experimental comparisons," *Experimental Economics*, 2011, *14* (3), 375–398.

Camerer, Colin F, Gideon Nave, and Alec Smith, "Dynamic unstructured bargaining with private information: Theory, experiment, and outcome prediction via machine learning," *Management Science*, 2018, *65* (4), 1867–1890.

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.

Cox, James and Glenn W Harrison, *Risk aversion in experiments*, Vol. 12 2008.

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.

Fuchs, William, Aniko Öry, and Andrzej Skrzypacz, "Transparency and distressed sales under asymmetric information," *Theoretical Economics*, 2016, *11* (3), 1103–1144.

Fudenberg, Drew and Jean Tirole, "Game theory, 1991," *Cambridge, Massachusetts*, 1991, *393* (12), 80.

Gächter, Simon and Arno Riedl, "Moral property rights in bargaining with infeasible claims," *Management Science*, 2005, *51* (2), 249–263.

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.

Karagözoğlu, Emin and Arno Riedl, "Performance information, production uncertainty, and subjective entitlements in bargaining," *Management Science*, 2014, *61* (11), 2611–2626.

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.

Li, Shengwu, "Obviously strategy-proof mechanisms," *American Economic Review*, 2017, *107* (11), 3257–87.

Miettinen, Topi, "The partially cursed and the analogy-based expectation equilibrium," *Economics Letters*, 2009, *105* (2), 162–164.

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.