Ultimatum Game Bargaining in a Partially Directed Search Market

Andrew Kloosterman and Stephen Paul

University of Virginia

June 12th, 2018

Abstract

We investigate a partially directed search and bargaining market with a laboratory experiment. First, sellers post intervals of possible surplus splits (i.e. the payoffs that would result from posting possible prices) that direct buyers to approach them. Second, after matching occurs, final surpluses are determined by ultimatum game bargaining. We investigate the interaction between bargaining and competition in the preliminary search stage, with a focus on how preferences for fair bargaining outcomes affect search. The main results confirm that behavior in the ultimatum game is consistent with preferences for fair outcomes, and the main effect on search is to drive up the posted buyer surplus lower bounds above the competitive equilibrium towards more equal surplus splits. Our main treatment variable is the number of buyers in the market, and when the number of buyers is increased, lower bounds and ultimatum offers to buyers decrease. This is consistent with fairness perceptions being influenced by competition.

Key words: Experimental Economics, Directed Search, Ultimatum Game

JEL Codes: C78, C92

The authors would like to thank James Cox and participants of the 2015 and 2016 ESA North American Meetings and seminars at the University of Virginia, Georgia State University, and the Atlanta Federal Reserve Bank.

1 Introduction

Many markets operate in two stages. First, buyers search for a seller to purchase from. Second, they bargain with the seller over the price. In the goods market, this is common for expensive durable goods such as homes, cars, etc. It is also common on online marketplaces such as Craigslist. Many labor markets operate this way too. A well-established fact about bargaining is that the participants' concerns for fairness may impact the outcome.[1] We investigate whether fairness still matters when bargaining occurs after search, and once we show that it does, what effects it has on the search stage.

Our search model is directed search as introduced by Burdett, Shi, and Wright (2001) (hereafter BSW)[2] generalized to allow sellers to post an interval of possible buyer surpluses.[3] We call this partially directed search, because sellers post only partial information (an interval of possible surpluses) rather than the precise information posted in BSW (an exact surplus).[4] Buyers observe the posted possible surpluses and then decide which seller to approach. Sellers have 1 unit of the good (which is of no value to them) and buyers have unit demand. This supply and demand structure captures the examples listed above, but our main reason to choose it is that it implies theoretical tradeoffs for both sides of the market. Sellers want to keep as much surplus for themselves as possible, but know this may deter buyers from approaching. Buyers want to approach sellers who they think will offer the most surplus, but know that such sellers may be more congested with buyers. This induces competition on both sides of the market which is the focus of our analysis.[5] Specifically, the effects for search we alluded to at the end of the first paragraph are the effects on this competition, and we investigate a number of questions related

[1] Fairness in bargaining refers to agents that consider their own payoff as well as the division of surplus where, usually, more equal divisions are preferred to more unequal divisions (inequality aversion). Formal models include Fehr and Schmidt (1999) and Bolton and Ockenfels (2000).

[2] Other similar early models of directed search include Peters (1991), Montgomery (1991), and Moen (1997).

[3] In most market games, sellers post prices. In the ultimatum game (see below) or the labor market, sellers usually post buyer surpluses. Of course, this is just framing as one can easily be calculated from the other. Hereafter, we borrow the terminology from the ultimatum game.

[4] We generalize directed search by allowing sellers to post partial information. Another literature has generalized directed search by allowing sellers to post the mechanism they will use to sell their object (see Albrecht, Gautier, and Vroman (2014) for a recent example).

[5] Alternatively, in random search models, there is only competition among sellers, because buyers do not actively search (see Mortensen (2005)).

to them. For example, do sellers post large intervals with the expectation that the final surplus split will be a fair split, or do they post narrow (or possibly even degenerate) intervals that guarantee a fair final split? And if they post large intervals, is the final split fair relative to the intervals or to the total value of the good? The answers to these questions, and others, are of course tied to buyer search behavior, because sellers compete to attract buyers.

The bargaining protocol we implement is ultimatum game bargaining with sellers proposing to buyers. While the actual bargaining structure in many markets is less rigid (in the sense that buyers may be able to make counteroffers), the environment is already quite complex and we think the ultimatum game is the simplest way to identify the basic effects of bargaining on search competition. Additionally, the ultimatum game has perhaps the starkest divergence of behavior and theory due the theoretical prediction being extremely unequal.[6] Therefore, there is a good chance to see behavior in the search stage that also diverges from theory and in ways that are interpretable through fairness.[7] Finally, this protocol is directly comparable to standard directed search where there can only be one offer and so only one (meaningful) round of bargaining.

We analyze results of a laboratory experiment on the following partially directed search and bargaining game. First, sellers post an interval range of possible buyer surpluses. Then, buyers view the ranges and select one seller to approach. A buyer is matched to the seller they approach if they are the only buyer who approaches them, or one of the approaching buyers is randomly selected to match if multiple buyers approach the seller. Second, each matched seller makes an offer from the range they have posted, and this offer is accepted or rejected by the buyer. This final stage is the ultimatum game constrained to the interval of possible divisions posted by the seller.[8] It is important to note that the equilibrium outcomes of partially directed search

[6]The proposer in an ultimatum game should offer nothing to the responder (or perhaps a penny to make the best response strict) and the responder should accept because receiving something is better than nothing. Experiments have shown that offers are usually intermediate, averaging around 30-40% of the endowment, and in fact when offers are small, below about 20% of the endowment, responders routinely reject them (see, for example, Camerer and Thaler (1995)).

[7]More generally, the effects of bargaining outcomes on search go back to pioneering search theory. Diamond (1982) and Pissarides (1985) look at the effects for the number of vacancies posted in labor markets while Mortensen (1982) looks at the effects for how hard agents search when search is costly. Investigating these other kinds of links in the presence of concerns for fairness would be interesting future work.

[8]Constraints such as this are occasionally referred to in the literature as a mini-ultimatum game. For an example, see Falk et al. (2003).

are identical to those of standard directed search. This is because the seller will offer the lower bound of their posted interval, so both sides theoretically view the lower bound as the final offer. We can therefore use the equilibrium predictions of directed search as a baseline prediction, and then compare differences to this prediction to examine fairness effects.

In all treatments, there are two sellers and our main treatment variable is the number of buyers in the market. In one treatment, there are three buyers, and the theoretically predicted surplus for matched sellers is 72.7 and matched buyers is 27.3, which fairly closely coincides with the usual offers in the ultimatum game. In the second treatment, there are five buyers and the theoretically predicted surplus for matched sellers is 91.23 and matched buyers is 8.77, which is far more unequal. We also consider standard directed search where sellers must post an exact offer, although buyers still have the option to reject the offer after they match with the seller. These baseline treatments allow us to understand the effects of offering partial information as opposed to exact information.

We find that the ultimatum game behavior has a substantial effect on how both the buyers and the sellers behave. We confirm that offers are well above the predicted equilibrium, particularly in the treatment with five buyers. This is true in the baseline treatments too, which suggests that just having the power to reject the final offer increases buyer surplus. We also find that, on average, sellers post fairly large intervals. The lower bounds are in fact very close to offers, suggesting that theory predicts well in this subgame and that fair offers are set relative to total value and not the endogenously selected intervals. The upper bounds are quite large, which we interpret as slight evidence of tricking buyers into thinking larger upper bounds signal larger offers although it also may be due to larger upper bounds just not driving away buyers.

We also see some treatment effects which imply that the competition is not completely eroded by concerns for fairness. Offers are lower and buyers are willing to accept lower offers when there are more buyers. We interpret this as the competitive context affecting the norm for what a fair ultimatum offer is. In particular, more competition among buyers drives down the perceptions of what constitutes a fair offer.

Finally, we also investigate efficiency and find that the markets are reasonably (constrained) efficient, because buyers approach strategies turn out to be close to random (for complicated, but interesting, reasons we detail below), which is the constrained efficient strategy, and they rarely end up rejecting offers.

Our experiment contributes to the growing literature on directed search in the laboratory. Cason and Noussair (2007) showed (weak) convergence to the predicted equilibrium in the standard BSW model. Subsequent experiments confirmed these results and generalized them in different directions (Anbarci and Feltovich (2013), Paul (2015), and Kloosterman (2016)). Anbarci and Feltovich (2013) suggest that fairness may be driving their results, which we expand upon in our investigation here. Several very recent papers have also made search partially directed, although in a different way, by allowing only a subset of buyers to observe posted prices. Anbarci and Feltovich (2017) show that exogenously decreasing the number of buyers who can observe prices increases seller surpluses even when theory predicts the opposite, which they attribute to concerns for fairness. Helland, Moen, and Preugschat (2017) find results more in line with equilibrium. However, both these papers are quite different from ours because they study exogenous informational changes while information is endogenous in our case, and our main questions investigate this endogenous choice.

The ultimatum game is of course one of the most-studied games in the literature. It was first studied by Güth, Schmittberger, and Schwarze (1982). See Güth and Kocher (2014) for a recent review of the literature. Competition has been introduced along several dimensions. Roth et al. (1991) puts the ultimatum game in a competitive environment by considering many proposers and one responder while Güth et al. (1997) and Grosskopf (2003) consider one proposer and many responders. Fischbacher, Fong, and Fehr (2009) consider both these types of treatments. Behavior tends towards the subgame perfect equilibrium in these experiments with one-sided competition. Fischbacher, Fong, and Fehr (2009) explain that this is consistent with inequality aversion, because aversion is heterogeneous. So when there are multiple responders, one is likely to accept much less. In fact, even the responders who are more averse will now accept less too as they realize that some responder will take an unequal offer and so the outcome is likely to be unequal even if they chose to reject. This differs from our environment, however, where there is competition on both sides and, importantly, a rejected offer can not be accepted by a different responder. Another group of papers has considered competition to be the proposer, via various effort tasks, and this also drives offers towards the unequal equilibrium (Güth and Tietz (1985) and (1986), Hoffman et al (1996), Gächter and Riedl (2005)). This is interpreted as winning the competition giving the proposer the right to a larger share. We find a similar effect when we explain that the number of buyers in the market seems to set a standard for what constitutes a fair offer, although the norm is created by a different mechanism.

On the theory side, Menzio (2007) introduces a model of partially directed search. There are heterogeneous sellers, and the partially directed search that he introduces involves cheap talk signals that allow high-value sellers to attract more buyers. More recently, the literature has shown that the details of the environment determine whether cheap talk equilibria exist (see Kim and Kircher (2015)) or not (see Mirkin and Pycia (2015)). We do not have heterogeneous agents, and therefore our partially directed search model is theoretically equivalent to standard directed search. We instead investigate the behavioral consequences of allowing sellers to post only partial information. Nevertheless, we think that this environment is the most promising of all for future work on the link between fairness in bargaining and search, because, by nature of the information being cheap talk, there is necessarily a bargaining stage that follows matching.

2 Theory

2.1 Model

There are two sellers, each selling one good, and n buyers seeking to buy one good.[9] The value of each seller's good is 0 to themselves and v to the buyers, which we normalize to 100.

The game proceeds in three stages. First, each seller posts a lower bound L and an upper

[9] This can easily be generalized to m sellers, but we use 2 here because there are 2 sellers in all experimental treatments.

bound H satisfying $0 \leq L \leq H \leq 100$ (note that H and L may be equal), which are then shown to the buyers. All the values between L and H inclusive we call the offer range. Second, the buyers each simultaneously select one seller or the other to approach. A seller with one approaching buyer is matched with that buyer. A seller with more than one approaching buyer is matched with one of the buyers with equal likelihood for each. Third, each matched seller makes an offer x that must be in its posted offer range and the buyer to whom they are matched then observes this offer and decides whether to accept or reject it.

For matches where the offer is accepted, the seller gets $100 - x$ and the buyer gets x. For matches where the offer is rejected, both the seller and the buyer get 0. All unmatched agents also get 0.

2.2 Equilibrium

The model modifies the BSW directed search model by changing two parts. First, buyers view a range of possible offers selected by the sellers instead of the sellers' offers themselves and then observe the offer after they match. Second, the buyers may reject offers.

Nevertheless, the symmetric subgame perfect equilibrium outcome of the game is the same because the third stage of this game is the classic *ultimatum game* within the bounds set in the first stage.[10] Since buyers prefer any x to 0 they will accept every offer. Therefore, each seller will choose the smallest possible offer which is $x = L$. So, at the beginning of the second stage, everyone knows that $x = L$ for each seller and so the first two stages are played just as in the BSW model where L is the offer and H is arbitrary. Adapting the solution from BSW we arrive at the following characterization of the symmetric equilibrium that has been the focus of much of the literature.

Proposition 1. *There is a unique (in terms of final offers) symmetric subgame perfect equilibrium where*

[10] Symmetric means the buyers all choose the same strategy and the sellers all choose the same strategy. Coordination among buyers may allow for asymmetric equilibria, although buyers cannot communicate (and are randomly matched each period) in the experiment so there is no reason to believe any such equilibria are plausible.

1. *Both sellers choose*

$$L^* = 100 * \frac{n(1/2)^n}{2 - (2+n)(1/2)^n}$$

and H^* arbitrarily larger than L^*.

2. *All buyers select each seller with probability 1/2.*

3. *All matched sellers choose $x^* = L^*$ and all matched buyers accept.*

Since we expect that sellers may not always post equilibrium lower bounds, it will be useful in the results section to have a complete characterization of the solution to the buyers' problem for any posted lower bounds. Intuitively, if one seller posts a lower bound that is at least n times larger than the other, then all buyers select that seller, because getting this large offer $1/n$ of the time is better than the small offer for certain. In all other cases, the buyers play a mixed strategy which is found by solving their indifference condition.

Proposition 2. *Suppose seller 1 posts lower bound L_1 and seller 2 posts lower bound L_2. Then, the probability p that each buyer approaches seller 1 is*

1. *1 if $L_1/L_2 \geq n$*

2. *0 if $L_1/L_2 \leq 1/n$*

3. *and otherwise solves the equation*[11]

$$\frac{1-(1-p)^n}{np} L_1 = \frac{1-p^n}{n(1-p)} L_2$$

2.3 Fairness

The equilibrium predictions provide a starting point to which we will compare behavior. In fact, past experiments on the standard directed search game find that behavior is mostly consistent with these predictions (Cason and Noussair 2007, Anbarci and Feltovich 2013, Kloosterman

[11] The left hand side is the probability of matching with seller 1 if all three buyers approach seller 1 with probability p multiplied by the payoff L_1. The right hand side is likewise for seller 2. The equation simplifies to a polynomial of degree $n-1$ and thus has no closed-form solution unless $n \leq 4$.

2016). However, we add ultimatum game bargaining to the model and it is well-known that behavior in this game displays concerns for fairness (inequality aversion in particular). We expect that these concerns will have an impact both in this final bargaining stage and also, through forward-reasoning, in the preceding search stage.

To better understand the impact, in Appendix A we develop a model in which individuals have heterogeneous minimum acceptable offers. This model predicts, quite intuitively, that offers will be larger than the equilibrium offers. Depending on the distribution of minimum acceptable offers, sellers may choose offers equal to lower bounds as in the standard model but they also may end up choosing offers equal to upper bounds (if lower offers are too likely to be rejected). Buyer approach behavior is similar to the standard model except when only one of the sellers posts bounds consistent with an offer above their minimum acceptable offer (in which case they always apply to this seller).

However, we also think that fairness concerns may be more complicated in this environment. In particular, fairness may be reference-dependent in that there is a norm that determines what is fair and what is unfair. We think the reference point that determines the norm may be exogenous, endogenous, or both. The exogenous factor is the relative number of buyers to sellers. When there are more buyers, they will accept lower offers because they understand the sellers have greater bargaining power.

The endogenous factor is a little more subtle. The buyers receive information before they decide which seller to approach, and it therefore is reasonable to expect that fairness may be a function of this information. Specifically, it is possible that a fair offer is determined relative to the posted offer ranges, perhaps it's an intermediate value in the range, in which case both bounds will impact buyers' approach decisions. The direction of the impact is not obvious. On the one hand, if a buyer believes the final offer will be an intermediate value of the offer range, then larger upper bounds signal larger offers and therefore attract more buyers. On the other hand, if buyers think lower bounds fully determine offers, then they may react spitefully to large upper bounds which are only there as a trick and therefore larger upper bounds may attract fewer buyers. Perhaps other factors may matter as well (such as the size of the offer range) although, given our data, we will focus on upper bounds.

2.4 Efficiency

Efficiency is lost here when one (or both) of the sellers fails to sell their good. This can happen in two ways. Either a seller may fail to match or a seller's offer may be rejected. In equilibrium, all efficiency loss comes from search frictions, which occur when all buyers select one seller, because all offers are accepted. Since buyers select each seller with probability 1/2, the probability that a seller fails to match is $(1/2)^n$.

A social planner could potentially attain full efficiency by sending at least one buyer to each of the sellers (assuming $n \geq 2$). However, because the equilibrium focuses on the case of symmetry, the literature has considered a constrained social planner who must choose the same strategy for all buyers. The main result is that the equilibrium is constrained efficient.

Proposition 3. *The constrained efficient outcome is for each buyer to select each seller with probability 1/2. The probability that a given seller matches is* $1 - (1/2)^n$.

From the perspective of fairness, low offers are likely to be rejected which will decrease efficiency. However, to the extent that sellers respond by increasing their offers to avoid this contingency, the efficiency loss from rejection may be mitigated.

3 The Experiment

We conducted 20 sessions at the Veconlab Experimental Economics Laboratory at the University of Virginia with a total of 360 participants in the fall of 2015 and the springs of 2016 and 2018. The participants were drawn from the undergraduate student population at the University of Virginia and were recruited through the Veconlab and the Darden Sona email recruitment systems. No subject participated in more than one session of the study.

There were 4 treatments, consisting of 2 main treatments and 2 baselines. In both main treatments, participants played the game described above. The main treatment variable was that in one of the main treatments there were 3 buyers (Treatment *3B*) while in the other there were 5 (Treatment *5B*) . The motivation for this treatment variable is the following. We start with Treatment *3B* to replicate the most common market size in the previous directed search literature. However, the theoretical offer is fairly intermediate and in fact close to previous offers in the ultimatum game, so we consider Treatment *5B*, where the theoretical offer is much closer to 0. The baseline treatments were standard directed search games, sellers posted x immediately rather than H and L, except that the matched buyers could still reject the offers at the end. There was a baseline treatment for each market size, (Treatment *3B_Base* and Treatment *5B_Base*). The baselines allow us to compare sellers posting partially informative intervals to fully informative offers in an environment where buyers can reject final offers.

For each treatment and baseline session, Table 1 reports the the number of participants and equilibrium lower bound and offer calculated from Proposition 1. Each session consisted of 40 periods of the game. The experiment was programmed and conducted with the z-Tree software (Fischbacher 2007).

Table 1: Treatments

Treatment	# of buyers	# of Sessions/Participants	L^*	x^*
3B	3	5/75	27.3	27.3
3B_Base	3	5/75	-	27.3
5B	5	5/105	8.77	8.77
5B_Base	5	5/105	-	8.77

Each of the 40 periods of the main treatment sessions proceeded in the following way. The participants were randomly divided into three groups of 5 or 7 participants, depending on the treatment, and randomly assigned to be either a seller or a buyer. Sellers were prompted to choose a lower bound, an upper bound, and an offer.[12] After all sellers had chosen these values,

[12]We had sellers choose offers here, even though they weren't revealed to buyers until after the matching occured, in order to have

buyers were shown the lower and upper bounds for the sellers in their groups. The buyers then selected one of the two sellers. After all buyers had made their choices, sellers were matched with a buyer that had selected their posting. If more than one buyer selected one seller, one of the selecting buyers was randomly matched with the seller. The buyers that were matched with a seller then observed that seller's offer. They were able to either accept or reject the offer. Buyers that accepted the offer received a number of points equal to the offer. Sellers whose offers were accepted received a number of points equal to 100 less their offer. At the end of each period, the participants observed the following outcomes for their group: the number of buyers who approached each seller, the upper and lower bounds, the offers, the acceptance decisions (if applicable), and their own points received. All participants were provided with a history screen that showed this information in all previous periods.

The baseline sessions proceeded in a similar way, except that sellers chose only an offer in the first stage and this was immediately shown to buyers. Upon matching, the buyer still chose to either accept or reject their seller's offer.

The sessions lasted about 1-1.5 hours and participants were paid for 4 of the 40 periods that were selected randomly. They received points in the experiment that were converted to dollars at a given rate (rounding up to the nearest dollar at the end). They earned on average $31.46. Instructions for Treatment *3B* are appended.

4 Results

We will go through the results in reverse order of the game, starting with the ultimatum game stage and then proceeding to the search market stage. The experimental results will be compared to the equilibrium predictions made by Propositions 1 and 2. As mentioned above, to the extent that these predictions fail to hold up, we will explain how fairness considerations related to previous findings on the ultimatum game are likely the reason. The results will conclude with a discussion of efficiency, comparing results to the predictions of Proposition 3.

an offer even for sellers that do not match. That is, we use a strategy method for sellers.

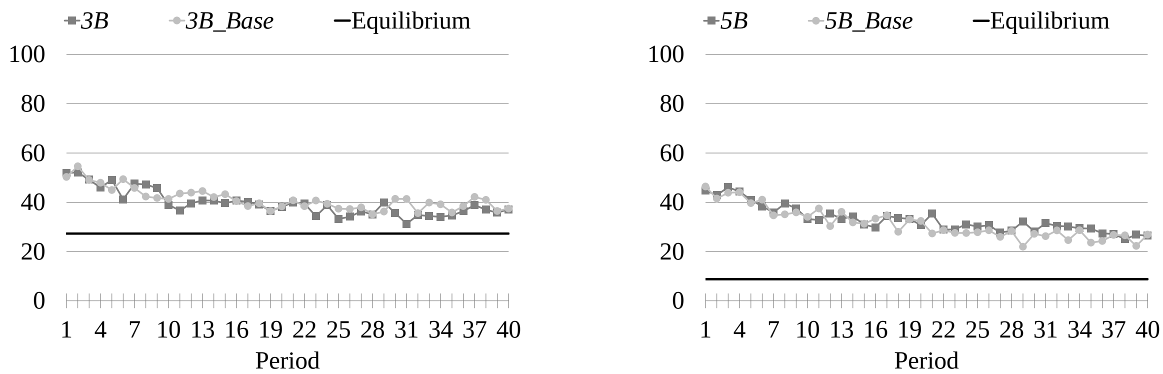

(a) Treatments *3B* and *3B_Base*

(b) Treatments *5B* and *5B_Base*

Figure 1: Average Offers by Period

4.1 Ultimatum Bargaining

We begin the analysis of the results with the final bargaining stage. Figure 1 displays average offers for each of the 40 periods of the experiment in each of the 4 treatments.

As seen in Figure 1, offers start around an even split, but decrease over the course of the experiment. In order to incorporate experience, we will focus on the last half of periods.[13] Table 2 provides the average offers for these periods, as well as average posted bounds for the main treatments which we will analyze later.

Table 2: Sellers' Choices for Periods 21-40

Treatment	Eq. Offer/Lower Bd.	Avg. Offer	Avg. Lower Bd.	Avg. Upper Bd.
3B	27.3	36.07	32.24	74.81
3B_Base	27.3	38.60	-	-
5B	8.77	29.27	24.65	72.11
5B_Base	8.77	26.47	-	-

There are two main results for offers. First, offers are larger than equilibrium offers in every treatment ($36.07 >^{**} 27.3$, $38.60 >^{**} 27.3$, $29.27 >^{**} 8.77$, and $26.47 >^{**} 8.77$, Wilcoxon signed-

[13]Offers seem to have converged in *3B* and *3B_Base* for these periods. They may still be slightly decreasing in *5B* and *5B_Base*, although the rate of decrease is small.

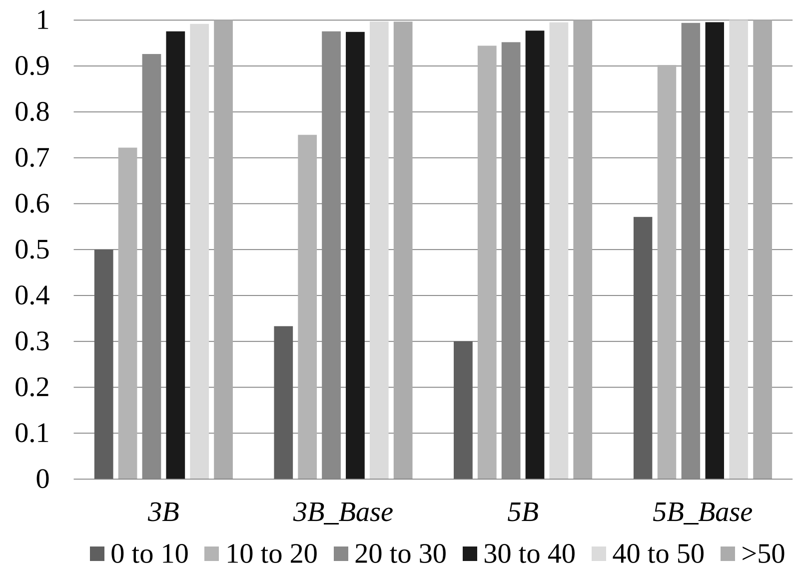

Figure 2: Acceptance Rates by Offer

ranks test).[14] Second, offers are larger in the treatments with 3 buyers than in the treatments with 5 buyers for both the main and baseline treatments (36.07 $>^{**}$ 29.27 and 38.60 $>^{**}$ 26.47, Wilcoxon-Mann-Whitney test).

The first result indicates evidence for the usual fairness motivations present in the ultimatum game as opposed to pure equilibrium play. This is particularly true for the treatments with 5 buyers where the equilibrium offer is much closer to 0. Consistent with previous findings, offers are in the range of 25-40% of the total surplus. However, the second result shows that buyer competition also plays a role as the difference in offers between treatments with 3 and 5 buyers is in the same direction as predicted by theory.

Next, we turn to the decision of buyers to accept or reject offers. Figure 2 plots the acceptance rates for offers between 0 and 10, 10 and 20, 20 and 30, etc. (the left endpoint is inclusive). We consider all periods, because few very low offers are made in the last half of periods.

[14]Significance is calculated with non-parametric tests on session-average data. It is noted on the $>$ sign with * if $p < .1$, ** if $p < .05$, and *** if $p < .01$.

First, offers less than 10 are rejected about half the time and offers greater than or equal to 20 are accepted almost all the time. This result confirms the well-known finding that low ultimatum game offers are routinely rejected, although one small difference is that offers close to 20 are often rejected in the ultimatum game but not here.

Second, the number of buyers has an effect here as well. Offers between 10 and 20 are still rejected about 25-30% of the time in the treatments with 3 buyers, but less than 10% of the time in the treatments with 5 buyers. The differences are statistically significant ($.72 <^{***} .94$ and $.75 <^{**} .90$, Wilcoxon-Mann-Whitney test). This suggests that the number of buyers does seem to set a norm for what constitutes a fair offer. It is important to note that the explanation given by Fischbacher, Fong, and Fehr (2009) for 1 seller and multiple buyers is not sufficient to explain this result. They see lower acceptance thresholds as the number of buyers increases too, but their explanation relies on other buyers having the opportunity to accept if a given buyer rejects which is not the case in our experiment.

Finally, there are no significant differences between each treatment and its respective baseline. We believe this is quite surprising, because the offers in the baselines were known before the decision of who to approach was made so approaching a given seller would perhaps be an implicit agreement to accept the offer. This was not the case, and we do not know why.[15] Result 1 summarizes our main findings.

Result 1. *Bargaining behavior displays the same well-documented fairness properties found in previous studies of the ultimatum game.*

1. *Sellers choose offers above the equilibrium offers.*

2. *Buyers reject very low offers.*

However, the number of buyers in the market also plays a role for determining offers and acceptance decisions.

[15] The buyers were forced to approach one seller or the other so it's possible that this result could be driven by low offers from both sellers. However, in 78% of the cases where the offer was rejected, the buyer had one offer of at least 20 so usually they had at least one decent option.

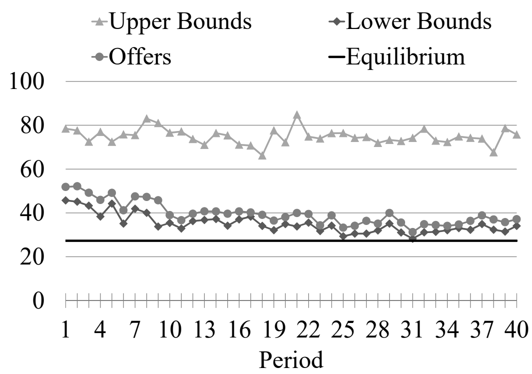

(a) Treatment *3B*

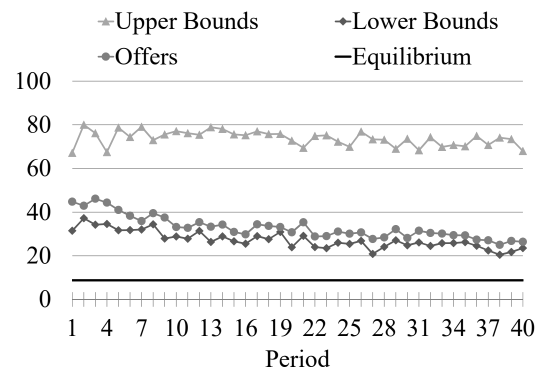

(b) Treatment *5B*

Figure 3: Average Bounds and Offers by Period

1. *Buyers are more likely to accept small offers in treatments with 5 buyers.*

2. *Sellers choose smaller offers in treatments with 5 buyers.*

Previous experiments have found that competition can drive behavior away from the fairness norms and towards the equilibrium (Roth et al. (1991), Güth et al. (1997), Grosskopf (2003), and Fischbacher et al. (2009)). Why have we found that fairness is still important, even in our more competitive treatments? We think the crucial difference is that our competition happens before the ultimatum game is played. It is competition among sellers to get a buyer to play their ultimatum game, but once a buyer has been obtained there is no competition left and so the final offers are still significantly above the equilibrium offer.

4.2 Search Market

Now, we investigate behavior in the search market. We begin with the offer ranges chosen by sellers. Figure 3 illustrates the means of the bounds of the ranges for each of the 40 periods of the experiment (along with the subsequent offers copied from Figure 1). The last two columns of Table 2 provide the average values for the last half of periods.

The results for lower bounds are the same as for offers. First, the lower bounds are larger than predicted by equilibrium although only significantly here for Treatment *5B* ($32.04 > 27.3$ with $p = .14$ and $25.07 >^{**} 8.77$, Wilcoxon signed-ranks test). Second, there are larger lower bounds in Treatment *3B* than in *5B* ($32.04 >^{**} 25.07$, Wilcoxon-Mann-Whitney test). This treatment effect is likely an artifact of both the theoretical prediction that less competition among buyers in *3B* leads to larger offers and, as seen in the previous section, that buyers will reject more low offers in *3B*.

Upper bounds are much larger than lower bounds for both treatments ($74.71 >^{**} 32.04$ and $69.07 >^{**} 25.07$, Wilcoxon signed-ranks test). While theory does not provide a prediction for the upper bounds, the persistence of large upper bounds throughout the experiment provides evidence that buyers continue to select sellers that post large upper bounds. That is, there does not appear to be evidence of spiteful behavior by buyers despite the fact that offers never come close to the upper bounds. In the next section, we will argue that buyers learn that the upper bounds are meaningless and simply ignore them. They probably stay persistently large, because they start large in the initial periods of the experiment.

Perhaps the most interesting result here is that offers and lower bounds track each other very closely. Offers are only slightly larger than lower bounds.[16] This finding is consistent with the theory that lower bounds are essentially offers. From the perspective of fairness, it seems that it is the total value of 100 that determines a fair offer, and not reference-dependent on the posted range.

There is one caveat to this finding. To see this, consider Figure 4 which plots lower bounds against their corresponding offer for the last half of periods. Larger bubbles indicate more data, so the figure indicates that most offers are exactly equal to lower bounds. Except when lower bounds are 10 or less. There are some sellers that do choose very small lower bounds, but then offer more than the lower bound. Perhaps these sellers believe that the lower bound does set a standard of fairness. The near 0 lower bounds allow them to simultaneously offer significantly more than their lower bound and yet still make a small offer. This strategy is effective as long

[16] They are significantly different, but given that the lower bound is a lower bound for the offer, this would only not be the case if every offer was exactly equal to the lower bound.

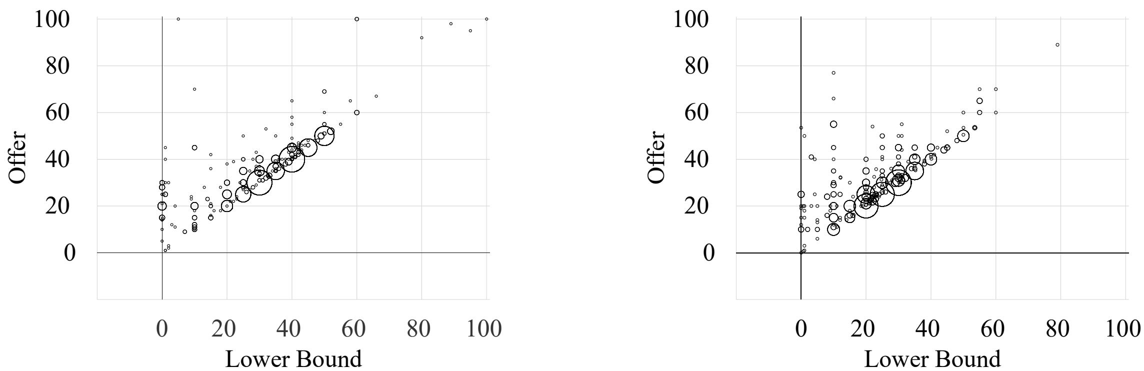

(a) Treatment *3B*

(b) Treatment *5B*

Figure 4: Average Offers by Lower Bound: Periods 21-40

as buyers expect the offers to be larger, or else they would never approach these sellers. We will momentarily show that, at least in Treatment *3B*, they do make this inference.

Now, we consider buyer behavior. The theoretical prediction for buyers is to select each seller half of the time. Of course, this prediction is based on the prediction that sellers will all post the same lower bound (or offer in the baselines), which is usually not true. However, Proposition 2 gives a solution for the buyers' problem for any pair of lower bounds (offers) to which actual approach behavior can be compared.

Figure 5(a) illustrates the approach rate to the seller that posts the larger lower bound in bins corresponding to predicted approach rates between .5 and .55, .55 and .6, .6 and .65, etc. for the main treatments and last half of periods.[17] Figure 5(b) creates the same plot for the baseline treatments, but based on offers rather than lower bounds.

Given how the figure is drawn, the 45 degree line represents the theoretical prediction for the approach probabilities. The figure indicates that approaches are reasonably close to predicted for Treatments *3B_Base* and *5B_Base*, but not as close for the main treatments. For Treatment

[17] So, for example, the diamond plotted at (.53,.51) in Figure 5(a) means that if we consider Treatment *3B* and all markets where the posted lower bounds are such that Proposition 2 predicts the approach probability to the seller that posted the larger of the two lower bounds is in the interval (.5,.55] then the average of the predicted applications is .53 whereas the average of the actual applications is .51. The right endpoint is inclusive except for the last bin. Predicted rates of 1 correspond to many different cases, and for the illustrative purposes here, are entirely excluded.

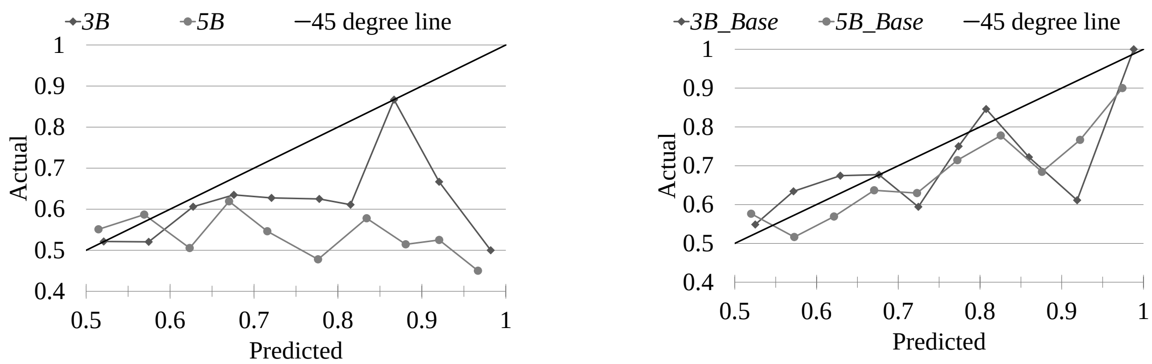

(a) Main Treatments

(b) Baselines

Figure 5: Predicted and Actual Approach Rates: Periods 21-40

3B, rates are close to 50% in the first two bins, 60% in the next 5 bins, and 70% for the last 3 bins so there is a positive relationship although it's not exactly linear nor is it as strong as predicted.[18] For Treatment *5B*, there seems to be almost no relationship between the predicted and actual approach rates at all.

To better understand what affects the buyers' approach decisions, we estimate a probit model of buyer approach to the seller that posts the larger lower bound (choosing seller 1 if the lower bounds are equal). We think that both the current observed lower and upper bounds as well as the outcomes of previous periods may influence choices. The variables of the first type we consider are the difference in lower bounds ($LBdiff$),[19] the difference in upper bounds ($UBdiff$), a dummy variable equal to 1 when the smaller lower bound is less than or equal to 10 ($LB0to10min$), and a dummy variable equal to 1 when both lower bounds are less than or equal to 10 ($LB0to10both$). We predict the sign of the coefficient on $LBdiff$ will be positive due to Proposition 2, the coefficient on $UBdiff$ could be positive if buyers think large upper bounds signal large offers or negative if buyers know this is not true and react spitefully, and the coefficients on $LB0to10min$ and $LB0to10both$ will be negative due to buyers recognizing that

[18]There is much variation for the last three bins, but this is due to a small number of observations ($n = 15$, $n = 21$, and $n = 6$ respectively). If we average across all $n = 42$ observations in these three bins, the average approach rate is 71%.

[19]We also could use the ratio of lower bounds to be more consistent with theory, but this causes problems when lower bounds are close to or equal to zero.

these small lower bounds do not represent the final offer.

The outcomes of previous periods may influence choices too which we investigate with a number of variables from the previous period in which the subject played in the role of buyer. The variables of this type that we consider are a dummy variable equal to 1 if the subject previously approached the seller who posted the larger lower bound and they matched ($LastChoosemaxMatch$), a dummy variable equal to 1 if the subject previously approached the seller who posted the larger lower bound and they did not match ($LastChoosemaxNoMatch$), a dummy variable equal to 1 if the subject previously approached the seller who posted the smaller lower bound and they matched ($LastChooseminMatch$), and the number of buyers who previously approached the seller who posted the larger lower bound ($LastNumAppsMax$). We predict the sign of the coefficient on $LastChoosemaxMatch$ will be positive because approaching the seller who posted the larger lower bound was a successful strategy, the coefficient on $LastChoosemaxNoMatch$ could be positive if buyers stick to choosing the seller offering the larger lower bound regardless of the outcome or negative if they switch when they get the bad outcome, the coefficient on $LastChoosemin$ will be negative because approaching the seller who posted the smaller lower bound was an unsuccessful strategy,[20] and the coefficient on $LastNumAppsMax$ will be negative because the more congested the seller was the less likely matching will occur.

Table 3 provides estimates for this model for each of the two treatments for the last half of periods.[21] The notable takeaway from the table is that buyer behavior is different in Treatment *3B* than in Treatment *5B*. In the first case, the current period seems to determine behavior, while in the second, it is the outcomes of the previous period in which the subject was a buyer.

[20] These predictions are relative to the omitted dummy for the case where the buyer approached the seller who posted the smaller lower bound and did not match. Formally, the predictions would come from a learning model where buyers use the same approach strategy (strategy here being to approach the seller posting the smaller or the larger lower bound) when it was successful and sometimes switch strategies when it was not.

[21] All buyers have previously been buyers by period 21 so we do not need to exclude any data. In the appendix, we look at learning effects by considering the first half of periods where we must exclude data for each subject from the first period in which they were a buyer.

Table 3: Buyer Approach Model, Periods 21-40

	3B	*5B*
$LBdiff$	.0179***	-.0017
	(.0054)	(.0038)
$UBdiff$	.0023	.0012
	(.0014)	(.0020)
$LB0to10min$	-.2905	-.0352
	(.1992)	(.0351)
$LB0to10both$	-.3642**	-.0185
	(.1763)	(.1057)
$LastChoosemaxMatch$	.1994	.3316***
	(.1807)	(.1359)
$LastChoosemaxNoMatch$	.2436	.2687*
	(.1593)	(.1596)
$LastChooseminMatch$	-.2435	-.1864**
	(.1841)	(.0753)
$LastNumAppsMax$	-.1094	-.0594***
	(.1104)	(.0161)
$constant$	.1676	.1973***
	(.1042)	(.0349)

Dependent Variable: 1 if approach to seller posting larger lower bound

Probit, s.e. (in parentheses) clustered at session level

* if $p < .1$, ** if $p < .05$, and *** if $p < .01$

More specifically, for Treatment *3B*, the coefficients for $LBdiff$, $LB0to10min$, and $LB0to10both$ are in the direction predicted and two of the three are significant (the other has a $p = .15$ so is close to significant). That is, as seen in Figure 5(a), buyers are responsive to the difference in lower bounds. Furthermore, they understand that very small lower bounds are not indicative of final offers. The coefficient on $UBdiff$ is positive although not significant ($p = .101$ so it is very close to marginally significant) indicating that buyers seem to ignore the upper bounds (or possibly have a slight inclination to apply more to the sellers posting them).[22] The coefficients on the variables corresponding to the previous period are all consistent with the qualitative predictions although none are significant.

On the other hand, for Treatment *5B*, no estimates for the variables related to the current period are significant. However, the variables related to the previous period in which the subject was a buyer are all in the predicted directions and significant. The positive value on $LastChoosemaxNoMatch$ suggests that buyers stick with their strategy of approaching the seller offering the larger lower bound even when it is not successful.

The natural question to ask is why does the current period matter in Treatment *3B* and the previous period in Treatment *5B*? We believe this may be due to the greater difficulty of matching as a buyer in *5B*. Specifically, buyers may be more focused on determining the seller that is less congested in this treatment, rather than determining which offer will be better, for which the data from the previous periods is quite useful.

Result 2. *Lower bounds closely track offers and so we get the same results for lower bounds as we did for offers.*

1. *For fairness: lower bounds are above the equilibrium predictions.*

2. *For competition: lower bounds are smaller when there are 5 buyers.*

Also, as predicted by theory, the upper bounds are essentially meaningless. However, while buyer approach behavior is focused on the current lower bounds in Treatment 3B it is focused on previous outcomes in Treatment 5B.

[22] This helps explain why upper bounds are so large throughout the experiment. They start large, and then sellers have no reason to decrease them so they just maintain the status quo. Also, in the appendix, we look at the regression results for the first half of periods and find a highly significant positive estimate for $UBdiff$ there so perhaps they were initially tricked and only later learned to ignore them.

4.3 Efficiency

The final section of the results addresses efficiency. Recall that the level of efficiency in this setting is simply the proportion of accepted offers. Inefficiency can arise in two ways, either when no buyer approaches a seller or when an offer is rejected. In equilibrium, all offers are accepted so inefficiency is created only in the first way. The equilibrium is also constrained efficient, but this relies on sellers posting the same lower bounds (or offers in the baselines). To the extent that these values are not equal, we can still calculate expected efficiency if buyers behave as Proposition 2 predicts which we call the interim expected match rate. Making matters more complex, fairness concerns may lead to rejection, the other source of inefficiency. Table 4 describes the interim expected match rate, the match rate, which is just the proportion of sellers that get at least one buyer, and the match and accept rate, which is the proportion of sellers that get at least one buyer who subsequently accepts the offer (i.e. the efficiency rate). We again focus on the last half of periods.

Table 4: Efficiency for Periods 21-40

Treatment	Equilibrium	Int. Exp. Match Rate	Match Rate	Match and Accept Rate
3B	.875	.761	.850	.820
3B_Base	.875	.827	.837	.825
5B	.969	.873	.953	.915
5B_Base	.969	.903	.957	.925

The experimental efficiency is always less than the equilibrium efficiency (which is also the constrained efficient outcome) and the differences are significant (.820 $<^{*}$ .875, .825 $<^{**}$ .875, .915 $<^{**}$ .969, and .925 $<^{**}$ .969, Wilcoxon signed-ranks test). Furthermore, as the equilibrium predicts, there is better efficiency when there are more buyers (.820 $<^{***}$ .915 and .825 $<^{***}$ .925, Wilcoxon-Mann-Whitney test).

The other striking result in the table is the low interim expected match rates in Treatments *3B* and *5B*. The actual match rates are much larger by comparison (.850 $>^{**}$ .761 and .953 $>^{**}$ .873, Wilcoxon signed-ranks test). Buyers induce more efficiency than predicted given sellers' choices because they do not approach the seller offering the larger lower bound quite as more often as Proposition 2 predicts (see Figure 5). This is particularly true for the sellers who post lower bounds equal to or near 0, which buyers understand will not be the final offers these sellers make.[23]

Finally, if we break it down by whether efficiency loss is due to matching inefficiency or rejected offers, we see that it's slightly more due to rejected offers in all treatments except Treatment *3B_Base*. However, they are similar and both types of loss are fairly small so we are hesitant to make a definitive statement as to which problem is greater.

5 Conclusions

In this paper we report the results of an experiment that studies the effects of embedding the ultimatum game in a competitive search model. In theory, the competitive equilibrium of the BSW model of directed search should predict the outcome of our experiment. We examined whether these predictions continue to hold with the introduction of ultimatum game bargaining.

We found that seller offers were significantly above the predicted equilibrium, particularly in our *5B* treatment, and that these offers were the lower bounds of the wider intervals that they initially posted. Buyers, at least in the treatment where there are just 3 of them, understand that the lower bound is the offer they will receive and approach sellers accordingly.

The literature on the ultimatum game shows that mean offers are generally around 30% of the endowment. We observe similar means regardless of the theoretical outcome. This leads us to believe that the behavior we observe is consistent with the behavior exhibited in prior experiments. The same fairness norm or fear of rejection that leads to divergence from the ultimatum game equilibrium leads to similar divergence from our equilibrium. However, the competitive element of our game does have an impact on this norm too. We find that the ratio of buyers to sellers does have an impact on determining what constitutes a fair split. Buyers are

[23] This result was also found in Kloosterman (2016), but is even more pronounced here due to the near 0 lower bounds.

willing to accept lower offers when they are oversupplied, and sellers make lower offers in this situation too.

Further research could examine either side of the market in greater detail by employing automated sellers or buyers that make a variety of predetermined choices or restricting the partial information that sellers could post further via fixed lower bounds, upper bounds, or interval widths. Another avenue would be to investigate whether introducing a competitive element in the ultimatum game itself, for instance, by rematching rejected sellers with a buyer that was not matched, would diminish the effect of the ultimatum game norms. It would also be interesting to investigate whether ultimatum game norms might work in the other direction by decreasing offers relative to the theoretical equilibrium when there are more sellers than buyers. Other similar environments; such as the random search, costly search, or cheap talk information mentioned in the introduction create further opportunities. Our paper shows that the interaction between the structure of competitive environments and the ultimatum game norms merits continued investigation.

We have provided evidence that bargaining norms distort equilibrium outcomes, even when embedded in a larger competitive environment. Previous experiments have shown that introducing competition into the ultimatum game weakens the influence of the fairness norms and pushes the result toward the game-theoretical equilibrium. We have shown that these norms can reassert themselves when subjects make a take-it-or-leave-it decision at the conclusion of a competitive process. This shows that the norms observed in the literature on the ultimatum game have implications beyond the narrow construct of the ultimatum game itself. Behavioral norms evidenced in our result might influence a wide range of bargaining outcomes, especially those characterized by search frictions and where few substitutes are available.

References

[1] Albrecht, James, Gautier, Pieter, and Vroman, Susan. Efficient Entry in Competing Auctions. *American Economic Review* 2014, 104(10), pp. 3288-3296.

[2] Anbarci, Nejat, and Feltovich, Nick. Directed Search, Coordination Failure, and Seller Profits: An Experimental Comparison of Posted Pricing with Single and Multiple Prices. *International Economic Review* 2013, 54(3), pp. 873-884.

[3] Anbarci, Nejat, and Feltovich, Nick. Pricing in Competitive Search Markets: the Roles of Price Information and Fairness Perceptions. *Management Science* 2017, 64(3), pp. 1101-1120.

[4] Bolton, Gary and Ockenfels, Axel. ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review* 2000, 90(1), pp. 166-193.

[5] Burdett, Kenneth, Shi, Shouyong, and Wright, Randall. Pricing and Matching with Frictions. *The Journal of Political Economy* 2001, 109(5), pp. 1060-1085.

[6] Camerer, Colin and Thaler, Richard. Anomalies: Ultimatums, Dictators and Manners. *Journal of Economic Perspectives* 1995, 9(2), pp. 209-219.

[7] Cason, Timothy and Noussair, Charles. A Market with Frictions in the Matching Process: An Experimental Study. *International Economic Review* 2007, 48(2), pp. 665-691.

[8] Diamond, Peter. Wage Determination and Efficiency in Search Equilibrium. *The Review of Economic Studies* 1982, 49(2), pp. 217-227.

[9] Falk, Armin, Fehr, Ernst, and Fischbacher, Urs. On the Nature of Fair Behavior. *Economic Inquiry* 2003, 41(1), pp. 20-26.

[10] Fehr, Ernst, and Schmidt, Klaus. A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics* 1999, 114(3): pp. 817-868.

[11] Fischbacher, Urs, Fong, Christina, and Fehr, Ernst. Fairness, Errors and the Power of Competition. *Journal of Economic Behavior & Organization* 2009, 72, pp. 527-545.

[12] Fischbacher, Urs. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 2007, 10(2), pp. 171-178.

[13] Gächter, Simon, and Riedl, Arno. Moral Property Rights in Bargaining with Infeasible Claims. *Management Science* 2005, 51(2), pp. 249-263.

[14] Grosskopf, Brit. Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition. *Experimental Economics* 2003, 6(2), pp. 141-158.

[15] Güth, Werner, and Kocher, Martin. More Than Thirty Years of Ultimatum Bargaining Experiments: Motives, Variations, and a Survey of the Recent Literature. *Journal of Economic Behavior and Organization* 2014, 108, pp. 396-409.

[16] Güth, Werner, Marchand, Nadege, and Rulliere, Jean-Louis. On the Reliability of Reciprocal Fairness: An Experimental Study. *Mimeo* 1997.

[17] Güth, Werner, Schmittberger, Rolf, and Schwarze, Bernd. An Experimental Analysis of Ultimatum Bargaining. *Journal of Economic Behavior and Organization* 1982, 3(4), pp. 367-388.

[18] Güth, Werner, and Tietz, Reinhard. Strategic Power versus Distributive Justice: An Experimental Analysis of Ultimatum Bargaining. in Brandsttter, H., Kirchler, E. *Economic Psychology Proceedings of the 10th IAREP Annual Colloquium* Linz, 1985, pp. 129-137.

[19] Güth, Werner, and Tietz, Reinhard. Auctioning Ultimatum Bargaining Positions: How to Act if Rational Decisions are Unacceptable? in Scholz, R.W. *Current Issues in West German Decision Research* Frankfurt, 1986, P. Lang Publisher, pp. 173-185.

[20] Helland, Leif, Moen, Espen, and Preugschat, Edgar. Information and Coordination Frictions in Experimental Posted Offer Markets. *Journal of Economic Theory* 2017, 167, pp. 53-74.

[21] Hoffman, Elizabeth, McCabe, Kevin, and Smith, Vernon. On Expectations and the Monetary Stakes in Ultimatum Games. *International Journal of Game Theory* 1996, 25(3), pp. 289-301.

[22] Kim, Kyungmin and Kircher, Philipp. Efficient Competition through Cheap Talk: The Case of Competing Auctions. *Econometrica* 2015, 83(5), pp. 1849-1875.

[23] Kloosterman, Andrew. Directed Search with Heterogeneous Firms: An Experimental Study. *Experimental Economics* 2016, 19(1), pp. 51-66.

[24] Menzio, Guido. A Theory of Partially Directed Search. *Journal of Political Economy* 2007, 115(5), pp. 748-769.

[25] Mirkin, Kenneth and Pycia, Marek. Directed Search and the Futility of Cheap Talk. *Mimeo* 2015.

[26] Moen, Espen R. Competitive Search Equilibrium. *The Journal of Political Economy* 1997, 105(2), pp. 385-411.

[27] Montgomery, James D. Equilibrium Wage Dispersion and Interindustry Wage Differentials. *The Quarterly Journal of Economics* 1991, 106(1), pp. 163-179.

[28] Mortensen, Dale. Wage Dispersion: Why are Similar Workers Paid Differently. MIT press, 2005.

[29] Mortensen, Dale. The Matching Process as a Noncooperative Bargaining Game. *The Economics of Information and Uncertainty* ed. John McCall, University of Chicago Press, 1982, pp. 233-258.

[30] Paul, Stephen. An Experimental Study of Partially Directed Search. *Mimeo* 2015.

[31] Peters, Michael. Ex Ante Price Offers in Matching Games Non-Steady States. *Econometrica* September 1991, 59(5), pp. 1425-1454.

[32] Pissarides, Christopher. Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages. *American Economic Review*, 1985, 75(4), pp. 676-690.

[33] Roth, Alvin, Prasnikar, Vesna, Okuno-Fujiwara, Masahiro, and Zamir, Shmuel. Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *The American Economic Review* 1991, 81(5), pp. 1068-1095.

Appendix A

In this section, we develop a theory with agents that care about fairness. Suppose each buyer i has a privately known minimum acceptable offer m_i and accepts an offer x iff $x \geq m_i$. Let F denote the commonly known population distribution of minimum acceptable offers. We will solve the game backwards.

In the ultimatum game, seller j chooses an offer

$$x_j \in argmax_{x_j \in [L_j, H_j]}(100 - x_j)\tilde{F}(x_j)$$

where $\tilde{F}$ is the distribution of m_i conditional on i approaching j. Unlike the pure self-interested case, it is not immediately clear that $x_j = L_j$ as the seller may want to choose a larger offer to garner more acceptance.

Moving backwards, the previous stage is the buyer approach decision. In equilibrium, the buyers know what x_1 and x_2 will be. There are three cases to consider. First, if $m_i > x_1, x_2$, then the buyer will reject no matter what so they are indifferent. We assume they approach each seller with probability 1/2. Second, if $x_2 \geq m_i > x_1$ or $x_1 \geq m_i > x_2$, then the buyer would reject the smaller offer and so has a dominant strategy to approach the seller that will make the larger offer. Third, if $x_1, x_2 \geq m_i$, then the buyer's decision is similar to that in Proposition 2, with a little adjustment. Namely, they approach seller 1 if $x_1 \geq nx_2$ and seller 2 if $x_1 \leq 1/nx_2$. Otherwise, suppose they approach seller 1 with probability q. To find q, let p be the total probability a buyer approaches seller 1:

$$p = qF(min\{x_1, x_2\}) + (F(x_1) - F(x_2))\mathbb{1}_{x_1 \geq x_2} + 1/2(1 - F(max\{x_1, x_2\}))$$

When p solves the equation in Proposition 2, then these buyers will be indifferent between the sellers and therefore may be able to mix. To check if they can, we solve for p using the equation in Proposition 2 and then for q in the above equation. If $q \in [0, 1]$, then they mix and we are done. However, it's possible that q would have to be less than 0 or greater than 1 to solve the above equation. If it's less than 0, then we must set $q = 0$. Note that then p will be larger than the p required for indifference which means seller 1 will be overcongested and so approaching seller 2 will be optimal. Likewise, if it's greater than 1, then we must set $q = 1$. Again, note that then p will be smaller than the p required for indifference which means seller 2 will be overcongested and so approaching seller 1 will be optimal.

Now, we need to determine what offers the sellers wish to make. In equilibrium, each seller anticipates the offer of the other seller and buyer behavior. Let $l_1(x_1, x_2)$ be the probability that a given buyer either approaches seller 2 or has $m_i > x_1$. Then, $(1 - l_1(x_1, x_2)^n)$ is the probability that seller 1 gets at least one buyer who will accept their offer. So seller 1 solves the problem

$$max_{x_1}(100 - x_1)(1 - l_1(x_1, x_2)^n)$$

Likewise, let $l_2(x_1, x_2)$ be the probability that a given buyer either approaches seller 1 or has $m_i > x_2$ so seller 2 solves

$$max_{x_2}(100 - x_2)(1 - l_2(x_1, x_2)^n)$$

Equilibrium offers (x_1^*, x_2^*) are any offers that satisfy both sellers' problems. Finally, each seller j needs to choose L_j^* and H_j^* so that x_j^* solves $x_j \in argmax_{x_j \in [L_j^*, H_j^*]}(100 - x_j)\tilde{F}(x_j)$. This is always possible by choosing $L_j^* = H_j^* = x_j^*$ although wider ranges are often possible too. Note that if they deviate from choosing bounds that constrain them to choose x_j^* then buyers will anticipate a different x_j will be chosen which will lead to non-optimal approach behavior and final offer from the seller's perspective.

<u>Predictions</u>

To better understand what this model predicts, and its limitations, we consider a distribution F that consists of a proportion s of buyers with $m_i = 0$ (i.e., self-interested buyers) with the remaining $1-s$ uniformly distributed between 0 and 50:

$$F(m_i) = \frac{1-s}{50} m_i + s$$

This model has two main differences from the standard model. First, there is never a symmetric equilibrium where $x_1^* = x_2^*$. In fact, for each fixed offer x_{-j}, the symmetric offer is a local minimum of the seller's objective function. There are two local maxima, one greater than x_{-j} and one less than x_{-j}, one of which is the global maximum. So there may be an asymmetric equilibrium where $x_1^* < x_2^*$ (and a second asymmetric equilibrium where the sellers' offers are interchanged) or there may be no pure strategy equilibrium at all for the sellers (in which case the mixed-strategy equilibrium is much harder to solve for).[24] We don't have a closed form-solution for the seller's problem so it is a little difficult to understand the intuition here.

Second, sellers may need to choose upper bounds to constrain themselves rather than lower bounds. This result is a little easier to understand. In this case, $\tilde{F}$ is also linear so the optimal offer is a negative quadratic function. This means that if the optimal offer for seller j is less than x_j^*, then we must have $L_j^* = x_j^*$ but H_j^* is arbitrary. But if the optimal offer for seller j is greater than x_j^*, then we must have $H_j^* = x_j^*$ but then L_j^* is arbitrary. As the number of self-interested types increases, there are fewer fair types so we are more likely to be in the standard case.

Now, we calibrate our model. Fehr and Schmidt (1999) postulate approximately 30% of the population are self-interested and the rest have minimum acceptable offers (roughly) uniformly distributed between 0 and 1/2 the surplus (we discuss Fehr and Schmidt (1999) further below). So we started with $s = .3$. For this value, L^* is constraining for the $n = 3$ treatment but H^* is constraining for the $n = 5$ treatment. So we also increased the number of selfish types to $s = .51$, the minimum number required for L^* to constrain in the $n = 5$ case. But for this parametrization, there is no pure strategy equilibrium for $n = 3$. The following table provides the results (as the offers are asymmetric we divide them into the low and high offer).

[24] Sellers' best responses functions are not continuous. There are two local maxima so when the global maxima switches between them, the best response function jumps.

Uniform F Predictions

	$s = .3$				$s = .51$			
# of buyers	x^*_{low}	x^*_{high}	L^*	H^*	x^*_{low}	x^*_{high}	L^*	H^*
3	37.67	41.11	x^*	arb.	-	-	-	-
5	26.42	28.95	arb.	x^*	23.92	22.55	x^*	arb.

The main takeaway from the table is that offers are larger than the equilibrium offers with only self-interested buyers.[25] Furthermore, the offers are larger when there are fewer buyers due to buyer competition still mattering (and the asymmetry does not matter too much).

Finally, we compare our model to the classic Fehr and Schmidt (1999) model of inequality aversion. They postulate the following utility function:

$$u_i(y_i, y_{-i}) = y_i - \alpha_i \frac{1}{N-1} \sum_{j \neq i} \max\{y_j - y_i, 0\} - \beta_i \frac{1}{N-1} \sum_{j \neq i} \max\{y_i - y_j, 0\}$$

where y_i is i's monetary payoff, $\alpha_i \geq \beta_i$ and $1 > \beta_i \geq 0$ are parameters describing i's inequality aversion, and N is the number of players.

We found that plugging this utility function into our game turned out to be intractable, but even if it was tractable, we do not think it makes much sense here, because this model assumes that agents care about everyone in their market. We believe that at the moment in time when a buyer decides whether to accept or reject an offer, they are likely thinking about inequality aversion just between themselves and the seller making the other offer but not about the rest of the agents in the market whose payoffs they can no longer control. One specific oddity about the Fehr and Schmidt model here is that it increases the likelihood of a buyer accepting an offer (relative to the standard ultimatum game) because the other offer in the market is likely to be accepted and so rejecting would still result in an unequal division although just in relation to the agents not involved in this bargaining subgame.[26]

[25] Even though there is no equilibrium for $n = 3$ and $s = .51$ we found that the offers that minimize the distance in best responses are 36.02 and 38.48 and that the difference from the true best response here is only .18 so there is an approximate equilibrium.

[26] This contrasts with Fischbacher et al (2009) where another agent gets the good instead of me if I reject so it makes sense that one would be averse to this.

Appendix B

Here, we look at the buyer approach model for the first half of periods.

Table 5: Buyer Approach Model, Periods 1-20

	3B	*5B*
$LBdiff$	.0070**	.0015
	(.0033)	(.0020)
$UBdiff$	.0041***	-.0004
	(.0011)	(.0010)
$LB0to10min$	.1475	-.0291
	(.1193)	(.1056)
$LB0to10both$	-.4891***	.1734
	(.0406)	(.7482)
$LastChoosemaxMatch$	.7545***	.2577**
	(.0983)	(.1311)
$LastChoosemaxNoMatch$	.4537***	-.0372
	(.0724)	(.1041)
$LastChooseminMatch$	.2438***	-.1686***
	(.0816)	(.0559)
$LastNumAppsMax$	-.2609***	-.0158
	(.0258)	(.0226)
$constant$	-.0013	.1302
	(.0826)	(.1106)

Dependent Variable: 1 if approach to seller posting larger lower bound

Probit, s.e. (in parentheses) clustered at session level

* if $p < .1$, ** if $p < .05$, and *** if $p < .01$

First buyer period omitted for each buyer due to previous outcome variables.

There are two main findings here. First, in Treatment *3B*, the coefficient on $UBdiff$ is significant in the first half of periods. This is consistent with buyers believing large upper bounds are good at first, but then learning that the upper bounds are meaningless. Second, in Treatment *3B* as well, the previous outcome variables are all significant. This suggests that initially, buyers in this treatment were more focused on learning from what they had already experienced.

Appendix C

Instructions for Treatment *3B*

Welcome. This is an experiment in decision making. Various research foundations and institutions have provided funding for this experiment and you will have the opportunity to make a considerable amount of money which will be paid to you at the end. We will go through the instructions together. Make sure you pay close attention because the decisions you make will influence the amount of money you will take home with you today. Please ask questions if any instructions are unclear as we read through them.

Experimental Task

There are 40 periods in this experiment and in each of these periods you will engage in the following task. First, you will be assigned to a group of five participants consisting of two participants in a role we will call A and three participants in a role we will call B. For identifying purposes, one of the A participants will be called A1 and the other A2. There is a total of up to 200 points that may be split among the A and B participants in the group according to following procedure:

- First, the two A participants simultaneously each choose a LOW VALUE, a HIGH VALUE, and an OFFER which are numbers between 0 and 100. The LOW VALUE must be less than or equal to the HIGH VALUE and the OFFER must be greater than or equal to the LOW VALUE and less than or equal to the HIGH VALUE. Note that A participants may choose their LOW VALUE equal to their HIGH VALUE in which case their OFFER must also be equal to them.

- Second, the three B participants are shown the OFFER RANGES for A1 and A2. These are the ranges from the LOW VALUE to the HIGH VALUE that the A participants have respectively chosen. Then, the B participants simultaneously choose to RESPOND to either A1 or A2. Note that the B participants know the OFFERs of A1 and A2 must be in the OFFER RANGES of A1 and A2 respectively, but they do not observe the OFFER before choosing which A participant they will RESPOND to.

- Third, MATCHES are formed as follows. If only one B participant RESPONDS to a given A participant then these two participants will be MATCHED together. If two or three B participants RESPOND to a given A participant then one of the B participants will be chosen randomly, with equal chance for each, to MATCH with that A participant. Note that an A participant is not MATCHED if all three B participants RESPOND to the other A participant and a B participant is not MATCHED if they are not the one selected randomly when more than one B participant RESPONDS to the same A participant as them.

- Fourth: MATCHED B participants observe the OFFER from the A participant they are MATCHED to. They then decide whether to ACCEPT or REJECT this OFFER. If they ACCEPT, the B participant gets OFFER points and the A participant gets 100-OFFER points. If they REJECT, then both the A and B participant get 0 points. Also, any participants that are not MATCHED get 0 points.

A Few More Procedures

At the beginning of each period, you will be randomly assigned to a new group of 5 and randomly assigned the role of A1, A2, or B. This assignment will be shown to you on the top of your screen along with which of the 40 periods you are currently in.

At the end of a period, you will be told your outcomes for the period. This depends on whether you were an A or B participant in that period. If you were an A participant, the information you receive is your LOW and HIGH VALUES, the number of Bs who RESPONDED to you, whether you were MATCHED or not, your OFFER, whether your OFFER was ACCEPTED or REJECTED (this will just be noted NA for not applicable if you did not MATCH), and the number of points you earned. If you were a B participant, the information you receive is the LOW and HIGH VALUES of the A participant you RESPONDED to, which A participant you RESPONDED to, whether you MATCHED or not, the OFFER from the A participant you MATCHED with (NA if you didnt MATCH), whether you ACCEPTED or REJECTED this OFFER (NA if you didnt MATCH), and the number of points you earned.

You will also get more information on the outcomes for your entire group of five. This information is the LOW VALUES and HIGH VALUES from both A participants, the number of B participants that RESPONDED to each of the A participants, the OFFERS from both A participants, and which OFFERS were ACCEPTED and which were REJECTED. All this information will remain on the right hand side of your screen for the rest of the experiment. One box will contain your outcomes (note that this will mean slightly different things depending on whether you were an A or B participant) and a second box below will contain your groups outcomes.

4 of the 40 periods have been randomly selected for payment. The sum of the points you earn in these 4 periods will be converted into dollars at the rate of 6 points per 1 dollar. You will also get 6 dollars for participating.