Economic Theory **25**, 553–574 (2005)
DOI: 10.1007/s00199-003-0464-2

Economic Theory

Bounded rationality in laboratory bargaining with asymmetric information*

Timothy N. Cason[1] **and Stanley S. Reynolds**[2]

[1] Department of Economics, Krannert School of Management, Purdue University, 100 S. Grant Street, West Lafayette, IN 47907-2076, USA (e-mail: cason@mgmt.purdue.edu)
[2] Department of Economics, Eller College of Business and Public Administration, University of Arizona, Tucson, AZ 85721, USA

Received: April 30, 2003; revised version: December 10, 2003

Summary. This paper reports an experiment on two-player sequential bargaining with asymmetric information that features some forces present in multi-round monopoly pricing environments. Buyer-seller pairs play a series of bargaining games that last for either one or two rounds of offers. The treatment variable is the probability of continuing into a second round. Equilibrium predictions do a poor job of explaining levels of prices and treatment effects. As an alternative to the conventional equilibrium model, we consider models that allow for bounded rationality of subjects. The quantal response equilibrium model captures some of the important features of the results.

Keywords and Phrases: Laboratory, Durable goods monopoly, Logit equilibrium.

JEL Classification Numbers: C78, D82, D42, C91.

1 Introduction

The literatures on sequential bargaining theory and durable goods monopoly theory focus on how repeated interaction between a single seller and one or more

* This research was funded in part by the National Science Foundation (SBR-9809110). The experiments were run at the Economic Science Laboratory of the University of Arizona and the Krannert Laboratory for Experimental Economic Research at Purdue University, using the *z-Tree* software developed at the Institute for Empirical Research at the University of Zurich (Fischbacher [8]). David Cooper, Rachel Croson, Charles Noussair, an anonymous referee, and conference participants at the Economic Science Association and the Society for the Advancement of Economic Theory meetings provided helpful comments. Timothy O'Neill Dang, Thomas Wilkening and Marikah Mancini provided expert research assistance.
Correspondence to: T.N. Cason

buyers influence market outcomes. Three types of forces shape outcomes in this environment. First, a buyer who has the option of deferring a purchase until later may have an incentive to wait for a better price offer even if their current payoff from making a purchase (consumer surplus) is positive. Second, a seller should anticipate strategic behavior by buyers when choosing prices. If a buyer is waiting for a better offer then the seller may have to set a price below the price they would set in a take-it-or-leave-it situation, in order to induce the buyer to purchase earlier. Third, the possibility of making multiple price offers may permit a seller to practice a form of inter-temporal price discrimination and extract more surplus from buyers than is possible in a static setting.

Researchers have utilized a variety of game theoretic formulations to examine how these forces play out. These formulations include models of sequential bargaining with asymmetric information and dynamic market models with complete information. These models yield equilibrium predictions about initial price offers, changes in prices over time, and buyer purchasing behavior.[1] Under some assumptions, the seller's initial price offer in equilibrium falls as the discount factor rises; as the discount factor approaches unity the equilibrium initial price offer converges to marginal cost (see Stokey [21]; Gul, Sonnenschein, and Wilson [12]. This last result seems to capture the essence of Coase's conjecture (Coase [5)] on durable goods monopoly pricing. However, equilibrium results are quite sensitive to assumptions. If one assumes complete information in a setting with a finite number of buyers and discrete demands, rather than a continuum of buyers or bargaining with a single privately-informed buyer, then the equilibrium may involve perfect price discrimination over time rather than prices close to marginal cost (see Bagnoli, Salant and Swierzbinski [1]).

Recent laboratory experimental studies by Güth, Ockenfels, and Ritzberger [13], Rapoport, Erev, and Zwick [18], Cason and Sharma [4], and Reynolds [19] have examined equilibrium predictions from these models. By and large, equilibrium predictions from game theoretic models of multi-round pricing have fared poorly in these experiments. Prices are often far from predicted levels, and buyers' purchasing behavior fails to conform to some predictions. Perhaps most significant is the typical failure of comparative statics predictions about the effects of changes in information, changes in the discount factor, or changes in the time horizon. Moreover, adjustments to the models to take into account risk preferences or preferences for fairness are not adequate to explain the results in these experiments (e.g., see the discussion in Reynolds [19]). And although rules of thumb and social norms about fairness can provide a partial explanation (Rapoport et al. [18]), no formal and systematic model has been proposed to explain the observed deviations from equilibrium.

The present paper reports on a new experimental design for sequential bargaining with asymmetric information. The design specifies a relatively simple decision-making environment, while maintaining enough richness to capture the forces at work in multi-round pricing environments. The simplicity of the experimental de-

[1] Reynolds [19] describes how the durable goods monopoly model has been extended in a variety of directions in recent papers.

sign serves two purposes. First, it gives equilibrium theory its best shot at successfully predicting subjects' behavior. This is important in light of the predictive failures noted above. Second, the simplicity of the design allows us to compute predictions for some general models of bounded rationality.

We have designed experiments in which buyer-seller pairs play a series of bargaining games that last either one or two rounds. The seller is uninformed about the resale value for the buyer she is matched with. The seller simply chooses a price offer from a set of ten possible prices in the first round and the buyer either accepts or rejects the offer. If the buyer rejects the offer then the game moves into a second round with some known, exogenous probability. If there is a second round, the seller chooses a price offer from this same set and the buyer either accepts or rejects.

The treatment variable in our experiments is the probability of continuing into a second round. We ran experiments with the following continuation probabilities: zero percent, 30 percent, 60 percent, and 90 percent. The game theoretic equilibrium prediction is that the opening price will fall as the continuation probability rises from zero to 30 percent and from 30 to 60 percent, and that the opening price will rise as the continuation probability rises from 60 to 90 percent. Strategic withholding of demand by buyers is predicted (on the equilibrium path) only for the case of a 90 percent continuation probability.

The environment is kept simple by employing a finite set of ten possible prices for the seller, two possible buyer values (high or low), and at most two rounds in a "game." This allows for straightforward computation of perfect Bayesian equilibrium predictions. More importantly, this relatively simple setup facilitates computations for models of bounded rationality in games. We present results for two such models: the Noisy Nash Model (Nash equilibrium play, plus random decision errors) and the agent quantal response equilibrium (AQRE) model proposed by McKelvey and Palfrey [17].

2 Experimental background

Several recent studies report on experiments that were designed to test game theoretic predictions of multi-round monopoly pricing models. The experiments of Güth, Ockenfels and Ritzberger [13] match a single seller with 10 buyers. Each experiment consisted of 5 or 6 experimental games. In each game there was a different combination of maximum number of trading rounds (two or three) and discount factors. In one experiment subjects received prior training in durable goods monopoly pricing games. The results for untrained subjects were grossly inconsistent with theoretical predictions from a model with a continuum of buyers. Prices failed to conform to comparative statics predictions and prices tended to be much higher than predicted. The levels of prices with trained subjects were closer to theoretical predictions, but prices still failed to satisfy comparative statics predictions.

Rapoport, Erev and Zwick [18] report on bargaining experiments with time discounting, one-sided incomplete information, and an infinite (unlimited) time horizon. They find that: (1) price offers tended to decline over time, as predicted by the sequential equilibrium (SE), (2) average initial prices were *higher* the higher the discount factor, contrary to the SE, and (3) for some discount factors the average initial price was above the static monopoly price, contrary to the SE. Rapoport et al. suggest that the SE theory may be failing because some buyers are using suboptimal rules of thumb for making purchases and because of a social norm about fair divisions of surplus.

Cason and Sharma [4] focus on the role of information about buyers' values. Game theoretic analyses of durable goods monopoly pricing predict large differences in outcomes depending on whether agents are informed about buyers' values. If all agents are informed about buyers' values then perfect price discrimination can emerge. This is in sharp contrast to asymmetric information environments, which typically yield equilibrium prices below the static monopoly price. Cason and Sharma report on infinite horizon experiments with two buyers and one seller. In the Certain Demand treatment all agents know that one buyer has a high value and the other buyer has a low value. In the Uncertain Demand treatment, buyers are privately informed about values; all agents know that the most likely outcome is that one buyer has a high value and the other has a low value, but outcomes with two low values or two high values are also possible. Prices in both treatments were much closer to equilibrium predictions of the incomplete information model than to the equilibrium predictions of the complete information model.

Reynolds [19] reports on experiments with one-sided incomplete information, a finite horizon, and either one buyer (bargaining) or five buyers (market). Bargaining experiments were run with one, two and six trading rounds, while all market experiments were run with six rounds. The demand withholding results were roughly in line with game theoretic predictions: withholding is much lower for one round games and for the final round of multi-round games than for earlier rounds of multi-round games. Reynolds nevertheless finds a significant failure of the equilibrium pricing predictions. The perfect Bayesian equilibrium predicts that opening prices will fall as the trading horizon increases from one to two to six rounds. Instead, opening prices in the bargaining experiments *rise* as the trading horizon increases, holding subject experience constant. Güth, Kröger and Normann [14], however, find better support for price comparative statics as the (private) discount factors change in a two-round bargaining experiment with incomplete information regarding buyer values.

3 Experimental design and predictions

We conducted a total of 9 sessions, employing 131 subjects. Fifteen or 13 subjects participated in each session, as summarized in Table 1. In each session subjects participated in a sequence of bargaining games against different anonymous opponents. Each bargaining game had either one round or two rounds of bargaining between a buyer and a seller. Subjects interacted only through a computer network running an application written using the University of Zurich's *z-Tree* program (Fischbacher [8]). All subjects were inexperienced, in the sense that they had never participated in a previous bargaining session that employed this design.

Our goal was to minimize repeated game incentives but allow subjects to obtain feedback and learn from their earlier experience playing the game, so we used the

Table 1. Summary of experimental sessions

Session name	Continuation probability	Session site	Number of subjects	Number of periods
UA00-1	0 percent	Univ. of Arizona	15	30
UA30-1	30 percent	Univ. of Arizona	15	30
UA30-2	30 percent	Univ. of Arizona	13	26
UA60-1	60 percent	Univ. of Arizona	15	30
UA90-1	90 percent	Univ. of Arizona	13	26
PU00-1	0 percent	Purdue Univ.	15	30
PU30-1	30 percent	Purdue Univ.	15[a]	30
PU60-1	60 percent	Purdue Univ.	15	30
PU90-1	90 percent	Purdue Univ.	15	28[b]

[a] Although this session employed 15 subjects, we discovered after the session that one of the subjects had previously participated in session PU00-1. Data from this experienced subject are excluded from the analysis.
[b] The software crashed after period 28 in this session.

matching scheme employed by Cooper et al. [6]. In this "strangers" design all subjects bargained exactly twice with each other subject in their session – once as the seller and once as the buyer – and subjects never knew the identity or history of their bargaining opponent. They never bargained against the same opponent in consecutive periods. Subjects alternated between being seller and buyer, and one subject sat out each period. Thus each session consisted of 30 periods, and each of the 15 subjects bargained 14 times as a seller and 14 times as a buyer. [The sessions with 13 subjects lasted for 26 periods.] Switching roles might have helped the subjects learn how to backward induct in this game, but it could have also led to some initial confusion. The ordering of the subjects and the pair assignments were determined randomly at the start of the session, when the subjects were randomly assigned to a computer.

The instructions (see http://www.mgmt.purdue.edu/faculty/cason/bounded-inst.pdf) were read aloud while subjects followed along on their own paper copy. The instructions explain the matching protocol described above. The instructions also explain that the sellers' cost is zero and the buyers' value is randomly determined, independently for each buyer. With a 40 percent chance the buyers' value is 54 and with a 60 percent chance the buyers' value is 18. An individual buyer's value is her private information. In each bargaining round the seller chooses a price offer from the set $\{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$, and then the buyer either accepts or rejects the offer. When an offer is accepted, seller profits simply equal the price, and buyer profits equal their value minus the price. If no offer is accepted for a particular period both the buyer and seller earn zero. The program did not allow buyers to accept price offers that result in negative profits.

Subjects were recruited from the undergraduate population at the University of Arizona and Purdue University. Upon arrival to the lab subjects were paid a $5 show-up fee. Including instructions, sessions required approximately 100 minutes to complete. Payoffs were converted using an exchange rate of 12 laboratory dollars = 1 U.S. dollar. Total earnings per subject ranged from $15 to $50, with a mean of about $36.

Perfect Bayesian equilibrium (PBE) predictions for risk neutral agents are summarized as follows. The PBE for the zero percent continuation probability is straightforward. The seller sets the highest possible price (50) and the high value buyer accepts this price offer. The expected payoff for the seller from this price $(50 \times 0.4 = 20)$ exceeds the payoff from setting a price that both buyer types would accept (15). There is no demand withholding in this equilibrium.

The form of the PBE is similar for the 30 and 60 percent continuation probability treatments. The high-value buyer accepts the round one offer, and if the game goes to a second round, the low-value buyer accepts the round two offer of 15. The seller's prices must satisfy the following incentive constraint for the high-value buyer (with value $\bar{v} = 54$):

$$(IC) \qquad \bar{v} - p_1 \geq \delta(\bar{v} - p_2)$$

In this inequality, δ is the continuation probability. The price p_2 is the highest price that a low-value buyer will accept; i.e., $p_2 = 15$. The seller chooses the highest price in the first round that satisfies the *IC*. For $\delta = 0.3$ this price is $p_1 = 40$ and for $\delta = 0.6$ this price is $p_1 = 30$. There is no demand withholding along the equilibrium path. Demand withholding would occur only for prices off the equilibrium path (e.g., if $\delta = 0.3$ and the seller sets $p_1 = 45$ then the PBE strategy for a high-value buyer is to reject this offer even though it provides positive consumer surplus).

Note that as the continuation probability increases from zero to 30 to 60 percent, the PBE initial price offer falls from 50 to 40 to 30. As high-value buyers become more likely to receive an attractive second round offer after a rejected first offer, the seller responds by lowering the initial price offer. The logic for this result is essentially the same as the logic behind the predicted effect of an increase in the discount factor in the infinite horizon experiments run by Rapoport, Erev and Zwick [18].

When the continuation probability gets closer to one in a two round game, the character of equilibrium changes. The seller will not set a price of 15 in the second round with probability one, and therefore a high-value buyer cannot count on receiving a low price offer if she rejects the initial price offer.[2] When $\delta = 0.9$ the PBE is as follows:[3]

1. $p_1 = 50$,
2. high-value buyer uses a mixed strategy for round one decision with $\Pr[accept\,|p_1 = 50] \approx 0.357$,
3. seller uses a mixed strategy in round two, with $p_2 \in \{15, 50\}$ and $\Pr[p_2 = 50] \approx 0.988$.

The PBE price predictions are summarized in Table 2.

2 When $\delta = 0.9$ there is no initial offer that is consistent with a screening equilibrium (an equilibrium in which the high value buyer accepts the initial offer and the low value buyer accepts the second offer). When $\delta = 0.9$ a high value buyer will reject any initial offer above 15, if they expect a second offer of 15. An equilibrium must involve a second offer above 15 with positive probability.

3 The derivation of this equilibrium is similar to the derivation in Fudenberg and Tirole [9]. The main difference is that the price space is discrete in our setting, whereas Fudenberg and Tirole have prices chosen from a continuous interval.

4 Results: comparison with PBE predictions

Table 2 presents the mean and median offer prices in each treatment, pooled across periods. (The rightmost column is discussed later, in Sect. 5.) Figure 1 below indicates that the time trends for these prices are not substantial, and there is no evidence for significant time trends in the second half of the sessions (periods 16–30).[4] In this table and in the subsequent analysis we pool across sites for the zero, 60 and 90 percent continuation probability treatments, since results did not differ significantly at Arizona and Purdue for these treatments. For reasons we are unable to explain, however, prices tended to be lower at Arizona than Purdue for the 30 percent continuation probability treatment, so to err on the side of caution we do not pool those datasets.[5]

In the first round the equilibrium prices are highly sensitive to the continuation probability, falling from 50 to 40 to 30, and then rising back to 50 as the continuation probability rises from zero to 90 percent. Observed opening offer prices, by contrast, do not vary much across the continuation probability treatments. Although (as predicted) mean prices are highest in the zero percent continuation treatment, in this treatment 47 percent of the opening offer prices are less than or equal to 30, while only 31 percent equal the prediction of 50 (and 10 percent are 45). Even more striking are the low opening offer prices for the 90 percent treatment. Fully 74 percent of these prices are less than or equal to 30, and only 4 percent are either 45 or 50. It is tempting to attribute the mean price level to “fairness,” since prices in the range of 25 to 30 provide a roughly equal split of the exchange surplus between the seller and the value=54 buyer. But as we discuss below, the subjects in this experiment appear substantially less sensitive to fairness concerns than are subjects in other related experiments.

The lower panel of Table 2 presents the summary statistics for the second round of price offers. For the 30 and 60 percent continuation treatments the second round mean prices are within 5 laboratory dollars of the prediction of 15. For the 90 percent continuation treatment the second round mean price is greater than the mean prices for the other treatments, but it is far below the PBE predicted price of 50.

Figure 1 displays the time series of mean opening offer prices for the four continuation probability treatments. This figure pools adjacent periods because subjects alternated between being buyers and sellers, so each subject contributes at most one

[4] To evaluate whether significant time trends exist in the data, we regressed prices on period number (and in an alternative specification, on 1/period). The coefficient estimates on these time trend variables are not significantly different from zero in any treatment when dropping the first 15 periods.

[5] In particular, according to nonparametric Wilcoxon Signed-Rank tests conducted for individual pairs of periods (so that each subject contributes one price offer to each test), opening offer prices are lower in UA30-1 than in PU30-1 at the 5 percent significance level in over half of the period pairs. According to this same test, opening offer prices in the two Arizona sessions UA30-1 and UA30-2 are not significantly different in any period pair. For the other session-site comparisons in the other three continuation probability treatments (i.e., UA00-1 vs. PU00-1, UA60-1 vs. PU60-1, and UA90-1 vs. PU90-1), these offer prices are significantly different at the 5 percent level in only 5 of the 43 period pairs. Similar tests for second round prices are only possible in the 60 and 90 percent continuation probability treatments, since few observations exist for second round prices in individual period pairs when the continuation probability is 30 percent. For these tests in the 60 and 90 percent continuation probability treatments, the prices are generally not significantly different across experiment sites.

Table 2. Summary statistics for price offers for each continuation probability treatment

Continuation probability	Obser-vations	Perfect Bayesian equilibrium offer	Actual offers: Mean offer (std. error)	Actual offers: Median offer	AQRE model mean (median)
Panel A: Opening offer prices					
0 percent	420	50	33.1 (0.73)	35	30.3 (30)
30 percent (Arizona)	366	40	22.7 (0.53)	20	30.1 (30)
30 percent (Purdue)	196	40	30.6 (0.70)	30	30.1 (30)
60 percent	420	30	27.7 (0.50)	30	29.1 (30)
90 percent	352	50	28.8 (0.41)	30	28.1 (30)
Panel B: Second round offer prices					
0 percent	–	–	–	–	–
30 percent (Arizona)	36	15	15.7 (1.09)	15	22.0 (15)
30 percent (Purdue)	28	15	20.0 (2.29)	15	22.0 (15)
60 percent	156	15	19.5 (0.90)	15	24.3 (20)
90 percent	245	50*	25.0 (0.89)	15	26.9 (25)

* The second round price of 50 is approximate. The seller sets a price of 50 with probability 98.8 percent and a price of 15 with probability 1.2 percent.

(price offer) observation to each adjacent pair of periods. In the opening periods the mean prices are highest in the 90 percent continuation probability treatment. Mean prices tend to fall in the 90 percent continuation treatment, however, while mean prices rise for the first third of the sessions in the zero percent continuation treatment. Prices also tend to fall in the 30 percent continuation probability sessions conducted at Arizona.

For statistical tests we cannot pool (non-independent) choices made by the same subject, so we conduct conservative nonparametric Wilcoxon tests separately for each pair of periods. As noted above, each subject contributes at most one price observation to each adjacent pair of periods because they alternated between buyer and seller roles. In the following discussion, "period pair 2" refers to periods 1 and 2, "period pair 4" refers to periods 3 and 4, and so on.

Table 2 shows that the PBE prices in round one vary substantially as the continuation probability varies. Consistent with the visual impression provided by Figure 1, however, in many period pairs the prices are not significantly different across treatments. Prices in the Arizona 30 percent continuation probability treatment are often significantly lower than prices in the zero and 90 percent continuation probability treatments, consistent with the PBE.[6] Also consistent with equilibrium, prices are

[6] In particular, nonparametric Wilcoxon tests indicate that prices are significantly lower in the Arizona 30 percent continuation probability treatment than in the zero percent continuation probability treatment

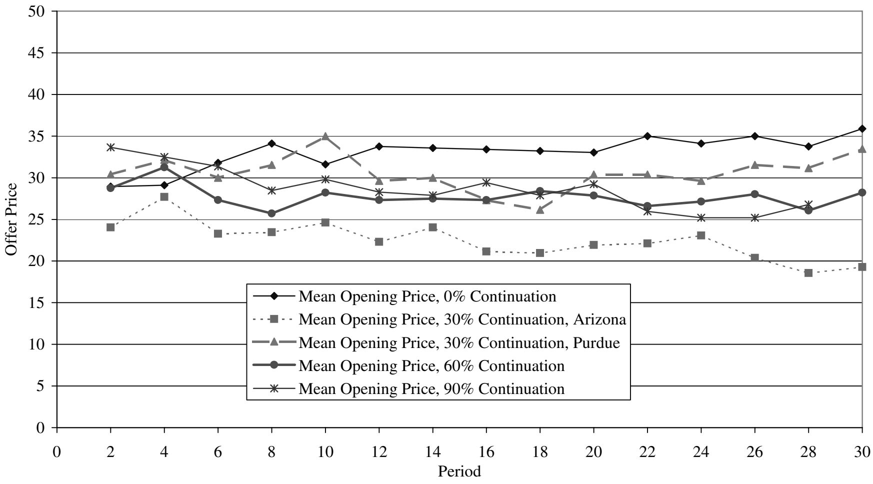

Figure 1. Mean opening prices by period pair for four continuation probability treatments

higher in the zero percent continuation probability treatment than the 60 percent continuation probability treatment in period pairs 8, 12, 14, and 22 through 30. But contrary to the PBE, prices are also significantly lower in the Arizona 30 percent continuation probability treatment than in the 60 percent continuation probability treatment in period pairs 2, 12 and 16 through 30, and they are significantly lower than the Purdue 30 percent continuation probability treatment in period pairs 6, 8, 10 20, 22, 26, 28 and 30. Prices are not significantly different at the five-percent level for 53 of the 58 other pairwise comparisons with a price difference predicted by the PBE (i.e., between zero and Purdue 30, zero and 60, Purdue 30 and 60, Purdue 30 and 90, and 60 and 90 continuation probability treatments).

A series of Wilcoxon signed-rank tests also clearly rejects the PBE null hypothesis for opening round prices in three of the four treatments. Figure 1 indicates that mean prices typically range between 20 and 35, while PBE prices vary between 30 and 50. In every period pair, the data reject at the five-percent level the null PBE hypotheses that (a) median prices equal 50 when the continuation probability is 90 percent; (b) median prices equal 40 when the continuation probability is 30 percent; and (c) median prices equal 50 when the continuation probability is zero percent.[7] The data reject the PBE hypothesis that median prices equal 30 when the continuation probability is 60 percent (at the five-percent significance level) only in period pairs 8 and 28.

The second round offer prices provide better support for the equilibrium predictions, although only the 60 and 90 percent continuation probability treatments have sufficient observations for individual period pair statistical tests. Consistent

in period pairs 6, 8, and 12 through 30; and are lower than in the 90 percent continuation probability treatment in period pairs 2 through 12, 16 through 22, 26 and 28 (all five-percent significance level, one-tailed tests).

[7] The one exception is period pair 5 for the 30 percent continuation probability treatment at Purdue, which is not significantly different from 40.

with the PBE, mean second round offer prices are higher for every period pair in the 90 percent treatment than in the 60 percent treatment, and the differences are statistically significant in 5 of the last 8 period pairs (pairs 14, 16, 20, 24 and 28). Moreover, second round offer prices are not significantly different from the PBE prediction of 15 in any period pair of the 60 percent continuation probability treatment. Prices are, however, significantly less than the PBE prediction of 50 in every period pair of the 90 percent continuation probability treatment.

Finally, consider the demand withholding (price rejection) choices by buyers. As we show in the next section, prices do not vary significantly over time in the second half of the sessions. Table 3 therefore presents the frequency distribution of opening price offers and acceptances for these later periods. Recall that we did not allow buyers to accept price offers that result in negative profits, which is why the acceptance rate for the value=18 buyer is uniformly zero for all prices greater than 15. The value=18 buyers usually accept offers less than or equal to 15 in the zero percent and Arizona 30 percent continuation probability treatments, but in the other treatments they accept offers of 15 less frequently. The value=54 buyers accept most price offers in the zero percent and Arizona 30 percent continuation probability treatments, but in the other treatments they often reject high price offers. These rejections for both types of buyers are contrary to the PBE. Such demand withholding is less common in the second round (not shown); in the second round buyers reject only 22 out of 165 acceptable offers (13 percent).

Although the opening round offer rejections are not consistent with the PBE, they are also not consistent with the social "fairness" norms that have been modeled and calibrated to related games. For example, the zero percent continuation probability results at the top of Table 3 indicate how results from this game are strikingly different from the typical rejection behavior in the well-known ultimatum game. In the ultimatum game buyers frequently reject offers that give them less than half of the exchange surplus. By contrast, our buyers accept offers of less than 17 percent of the surplus over 80 percent of the time (e.g., price offers of 15 to a value=18 buyer, and price offers of 45 to a value=54 buyer). High value buyers accept an offer of 7.4 percent of the surplus (i.e., a price of 50) two-thirds of the time, whereas in the ultimatum game such "unfair" offers are almost always rejected.

Fehr and Schmidt [7] provide a summary of ultimatum game results and model the fairness concerns that lead to equitable outcomes in that and other games. It is straightforward to show that their approach does a poor job describing the buyer choices in this experiment. In Fehr and Schmidt's [7] model preferences include "inequity aversion," so that subjects prefer higher earnings for themselves but also more equal earnings across players. Utility is equal to monetary payoffs less inequity costs that rise as the difference between a subject's own and the other's monetary payoff increases.[8] Fehr and Schmidt also derive parameter distributions for the relative tradeoff of monetary gains and inequity aversion that describes behavior across a variety of games, which we can use to assess the effectiveness of this approach in describing the rejection rate data reported here. Applying their

[8] In particular, for a two-person game player i's utility is $U_i(x) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\}, i \neq j$, where x_k denotes monetary earnings ($k = i, j$), $\alpha_i \geq \beta_j$, and $1 > \beta_i \geq 0$.

Table 3. Opening offer price acceptances for late periods (periods 16–30)

Continuation probability	Buyer value		Prices 5	10	15	20	25	30	35	40	45	50
0 percent	18	Price frequency	0	0	41	2	1	10	8	4	14	36
0 percent	18	Number accepted	0	0	35	0	0	0	0	0	0	0
0 percent	18	Acceptance rate	**–**	**–**	**0.85**	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0 percent	54	Price frequency	0	1	25	1	5	6	8	3	12	33
0 percent	54	Number accepted	0	1	25	1	5	6	7	2	10	22
0 percent	54	Acceptance rate	**–**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**0.88**	**0.67**	**0.83**	**0.67**
30 percent (Arizona)	18	Price frequency	0	9	48	9	6	14	6	3	3	1
30 percent (Arizona)	18	Number accepted	0	9	39	0	0	0	0	0	0	0
30 percent (Arizona)	18	Acceptance rate	**–**	**1.00**	**0.81**	0.00	0.00	0.00	0.00	0.00	0.00	0.00
30 percent (Arizona)	54	Price frequency	0	10	27	3	13	9	4	4	2	0
30 percent (Arizona)	54	Number accepted	0	10	27	3	13	9	3	3	0	0
30 percent (Arizona)	54	Acceptance rate	**–**	**1.00**	**1.00**	**1.00**	**1.00**	**1.00**	**0.75**	**0.75**	0.00	–
30 percent (Purdue)	18	Price frequency	0	0	12	0	4	19	10	6	8	2
30 percent (Purdue)	18	Number accepted	0	0	3	0	0	0	0	0	0	0
30 percent (Purdue)	18	Acceptance rate	**–**	**–**	**0.25**	–	0.00	0.00	0.00	0.00	0.00	0.00
30 percent (Purdue)	54	Price frequency	0	0	4	1	4	17	7	2	2	0
30 percent (Purdue)	54	Number accepted	0	0	4	1	4	17	6	0	1	0
30 percent (Purdue)	54	Acceptance rate	**–**	**–**	**1.00**	**1.00**	**1.00**	**1.00**	**0.86**	**0.00**	0.50	–
60 percent	18	Price frequency	0	5	20	13	17	41	12	14	1	5
60 percent	18	Number accepted	0	3	14	0	0	0	0	0	0	0
60 percent	18	Acceptance rate	**–**	**0.60**	**0.70**	0.00	0.00	0.00	0.00	0.00	0.00	0.00
60 percent	54	Price frequency	0	3	22	2	8	27	6	9	0	5
60 percent	54	Number accepted	0	3	22	2	7	17	3	1	0	0
60 percent	54	Acceptance rate	**–**	**1.00**	**1.00**	**1.00**	**0.88**	**0.63**	0.50	0.11	–	0.00
90 percent	18	Price frequency	1	1	7	8	32	30	6	6	3	0
90 percent	18	Number accepted	1	1	4	0	0	0	0	0	0	0
90 percent	18	Acceptance rate	**1.00**	**1.00**	**0.57**	0.00	0.00	0.00	0.00	0.00	0.00	–
90 percent	54	Price frequency	0	0	9	6	30	10	1	6	1	0
90 percent	54	Number accepted	0	0	9	6	19	5	0	0	0	0
90 percent	54	Acceptance rate	**–**	**–**	**1.00**	**1.00**	**0.63**	**0.50**	**0.00**	**0.00**	**0.00**	–

Notes: The perfect Bayesian equilibrium price offer is shown in the box. This model predicts acceptance rates of 1.00 in the **bold** highlighted regions, and acceptance rates of 0.00 elsewhere, except for $p_1 = 50$ in the 90 percent continuation probability treatment, in which the acceptance rate for the high value buyer is 0.357.

distribution of preferences to our subjects, it is straightforward to show that in the 0 percent continuation treatment, 70 percent of the value=18 buyers should reject a price of 15.[9] Table 3, however, indicates that these buyers only reject 15 percent of these offers. Likewise, in this treatment the Fehr-Schmidt model and parameters indicate a 70 percent rejection rate for the value=54 buyers when prices are 45 or 50; but the observed rejection rate is only 29 percent.

Why don't we see the pronounced fairness effects that others often observe? We conjecture that there are two main reasons. First, we have uncertainty about buyer type, which could have weakened the perceived fairness norm. For example, a particular price might be fair to one type of buyer but not to the other type (Rapoport et al. [18]; Güth et al. [14]). Second, in this experiment subjects alternated between buyer and seller roles, which could also reduce the importance of fairness concerns – particularly the "unfairness" of highly asymmetric distributions of exchange surplus in different rounds. Subjects essentially were able to "take turns" taking advantage of the privileged (seller) position (Thaler [22]). Future experiments can systematically vary the role alternation and the incomplete information about buyer types in order to test these conjectures.

We close this section with a few remarks about the implications of risk aversion. The PBE predictions shown in Tables 2 and 3 are based on an assumption of risk neutrality. In the next section we estimate a model of boundedly rational, risk averse agents, so it is useful to first check whether risk aversion of subjects alone can account for deviations of observed play from PBE predictions. We show below that the estimated value of the Constant Relative Risk Aversion (CRRA) index is 0.2 for these data. If we recompute PBE predictions using this estimated risk aversion index, then PBE predictions change very little. Equilibrium predictions for 0 and 60 percent continuation probabilities do not change, and the structure of the equilibrium for the 90 percent continuation probability treatment remains the same, with only small changes in equilibrium mixing probabilities.[10] The initial price prediction for the 30 percent continuation probability treatment rises from 40 to 45. Thus, overall a change from risk neutrality to CRRA with a modest level of risk aversion results in fairly small changes in PBE predictions, and what change there is tends to be in a direction away from what we observe.

There are qualitative changes in PBE predictions for higher levels of risk aversion. For example, in the 0 percent continuation treatment, if the CRRA risk aversion index is above 0.24 then the initial equilibrium price offer drops from 50 to 15. There are similarly large changes in PBE predictions for high indices of risk aversion for the other continuation treatments. However, these changes in predictions are generally not in the direction toward what we observe in the experiments.

[9] For this calculation one only needs the distribution of players' disutility from disadvantageous inequality (i.e., the α parameter in the previous footnote), because the buyer's earnings are lower than the seller's earnings. We use Fehr and Schmidt's distribution of $\alpha = \{0, 0.5, 1, 4\}$ in proportions of $\{0.3, 0.3, 0.3, 0.1\}$.

[10] In particular, the probability that a high-value buyer accepts an initial price offer above 15 drops from 0.357 to 0.292. The probability that the seller sets a low price in period 2 ($p_2 = 15$) rises slightly, from 1.3% to 2.1%.

5 Bounded rationality

For this bargaining game the perfect Bayesian equilibrium predictions perform poorly in both quantitative and qualitative dimensions. Are game theoretic predictions essentially useless for predicting behavior in settings such as these bargaining experiments, or can game theoretic predictions be modified in some fruitful way?

In this section we examine the predictions of two models of bounded rationality, as alternatives to standard equilibrium models. The first model is the Noisy Nash Model (NNM), which is described in McKelvey and Palfrey [17]. The NNM posits that each agent plays their equilibrium (PBE, for our analysis) strategy with probability γ and randomizes (uniformly) over all strategies with probability 1-γ. The NNM allows for mistakes and suboptimal play on the part of subjects; the NNM generalizes the equilibrium model in a way that allows for variation in the data around the equilibrium point predictions. Of course, the PBE is a special case of the NNM, with $\gamma = 1$. Since this model puts equal weight on all nonequilibrium strategies, it cannot explain why certain deviations from equilibrium are more common and more plausible than others.

The second model is the quantal response equilibrium (QRE) model, and it is much more intuitively appealing because it can potentially explain the pattern of deviations from equilibrium (McKelvey and Palfrey [16]). Agents in the QRE model do not always select a strategy that maximizes their expected utility, but they choose actions that yield higher expected payoffs with higher probability. In contrast to the NNM, the probabilities of (non-Nash) choices in a QRE are sensitive to the expected payoffs for these choices. The QRE is related to an earlier analysis by Rosenthal [20] that examined the implications of probabilistic choice by boundedly rational players in games. It can capture the notion that strategies which are nearly optimal might be chosen almost as frequently as the optimal strategy. For example, as shown at the top of Table 3, buyers accepted the equilibrium price offer of 50 in 22 out of the 69 times a seller offered it in the later periods, resulting in an average seller payoff of 15.9. In these same later periods buyers accepted the price offer of 15 in 60 out of the 66 times sellers offered it, resulting in an average seller payoff of 13.6. This similarity in average payoffs might explain the similar frequency of these two, very different, price choices.

McKelvey and Palfrey's [16] paper develops the idea of QRE for normal form games with finite strategy sets. A quantal response is a smoothed-out best response, in the sense that a player chooses actions that yield higher expected payoffs with higher probability but does not choose a best response with probability one. The QRE is calculated as a fixed point in probability space. Each agent's expected payoffs determine the agent's choice probabilities. These expected payoffs are based on the choice probabilities of the other agents. This choice framework may be modeled by specifying the payoff associated with a choice as the sum of two terms. One term is the expected utility of a choice, given the choice probabilities of other players. The second term is a random variable that reflects idiosyncratic aspects of payoffs that are not modeled formally. The *logit*-QRE is derived by assuming that these random variables are independent and follow an extreme value distribution. In a *logit*-QRE each agent's choice probabilities follow a multinomial

Table 4. Maximum likelihood estimates for NNM and AQRE models* (data after period 15)

δ : Continuation probability	zero	30 percent Arizona	30 percent Purdue	30 percent Pooled	60 percent	90 percent	Pooled data
Noisy Nash Model							
γ	0.34	0.0	0.04	0.02	0.28	0.0	0.12
$S.E.$	(0.03)	(0.02)	(0.04)	(0.03)	(0.02)	(0.0)	(0.02)
$\mathcal{L}^*$	−514.4	−550.7	−364.7	−915.5	−724.8	−731.7	−3873.4
AQRE Model							
λ	2.8	1.5	2.0	1.7	1.9	2.1	1.9
$S.E.$	(0.35)	(0.07)	(0.15)	(0.08)	(0.12)	(0.11)	(0.06)
$\mathcal{L}^*$	−504.6	−494.9	−326.5	−824.0	−684.2	−651.8	−2671.8
# observations	210	171	105	276	210	157	853

* Standard errors are in parentheses. Standard errors were computed using the bootstrap method. The sample size for the bootstrap was equal to the sample size of the data (see the last row in the table). 150 iterations were run to create a sample of bootstrap estimators.

logit distribution with parameter λ. As λ increases each agent puts less weight on choices that yield sub-optimal expected payoffs. As λ approaches zero, each agent's strategy converges to a mixed strategy with equal choice probabilities for each possible action.

This concept is extended to extensive form games (which allow for asymmetric information) in McKelvey and Palfrey [17]. The equilibrium concept in this setting is termed an agent quantal response equilibrium (AQRE). At each information set a player decides on the probabilities of different actions that he/she can take. A *logit*-AQRE may be derived in which choice probabilities at each information set follow a multinomial logit distribution. The appendix presents the application of McKelvey and Palfrey's *logit*-AQRE model to our experimental environment, which provides the basis for maximum likelihood estimation reported below.

There are features of both buyer behavior and seller behavior that suggest that bounded rationality might be a useful way to interpret the experimental results. Buyers are observed to withhold sometimes in the last round even though the payoff would have been higher if a purchase was made. In addition, buyers do not use the kind fixed "cut-off" rule for making purchases that is prescribed by the PBE (e.g., when the continuation probability is 60 percent, the PBE strategy for a high value buyer is to purchase with probability one if $p_1 \leq 30$ and purchase with probability zero if $p_1 > 30$). Instead, as Table 3 indicates buyers appear to use decision rules in which the probability of purchase rises gradually as consumer surplus rises. A smooth response such as a probit or logit distribution describes buyers' decision rules better than a cut-off rule. Seller behavior also appears to be "noisy." Substantial variation in price choices exists both within and across subjects, even in a relatively simple environment such as the zero continuation probability treatment.

Table 4 reports maximum likelihood estimates for the NNM and AQRE models of bounded rationality. A single parameter is estimated for each model: γ for the NNM and λ for the AQRE. Below each parameter estimate we report the standard error and the log likelihood ($\mathcal{L}^*$) for the model. Following the standard practice for laboratory analyses using the QRE model, for these estimates we only include data from the later part of the sessions (e.g., see Goeree et al. [11] or Capra et al. [3]). In particular, we exclude data from periods 1–15. During the early periods of an experiment subjects are likely to be learning about how their rivals play the game and adjusting their own choices in response to what they learn. For some treatments there are trends in opening prices during the early periods. But as noted at the start of Section 4, there do not appear to be any trends in prices after period 15 in any of the experiments. Parameter estimates are reported for each treatment individually and also for pooled data. The hypothesis that the site for the experiment (Arizona or Purdue) yields different results was rejected for all values of δ except $\delta = 30$ percent. We therefore report parameter estimates for 30 percent Arizona and Purdue datasets separately.

The estimated value of γ for the NNM is well below 1/2 for all treatments, and is essentially zero for the 30 percent and 90 percent continuation probability treatments. For the pooled data the estimated γ is 0.12. This means that under the hypothesis that the NNM is the correct model, only 12 percent of choices correspond to Nash (perfect Bayesian) equilibrium choices with the remaining 88 percent of choices best characterized as random.

The frequency of these apparently random choices does appear to be related to their expected payoffs, however. The lower half of Table 4 indicates that the AQRE model performs better than the NNM in the sense that the log likelihood is higher for the AQRE for every treatment. The estimated values of the noise parameter λ appear reasonable compared to other estimates for experimental games in the literature. For example, in most of the games they study McKelvey and Palfrey [17] estimate λ coefficients that range between one and two.[11] The estimated λ for the zero percent continuation treatment is somewhat higher than in the other treatments. This implies less noise in subjects' best response behavior for the one-round bargaining game, which is not surprising since this game is less complex than the two-round games.

In order to improve the predictive power of the AQRE we allow for the possibility of subject risk aversion. Prices in the zero probability continuation treatment were much lower than the risk neutral equilibrium prediction of 50; indeed, 47 percent of the prices in this treatment were 15 or less. Risk aversion on the part of seller subjects could account for such low prices.[12] In our computation of the AQRE we posit a constant relative risk averse utility function for each subject of

[11] Comparisons of λ across experiments are of limited value, however, because λ is sensitive to the scale of payoffs, and the scale often varies across studies. For example, in a similar specification but with a power functional form of $\mu = 1/\lambda$, Capra et al. [2] and [3] estimate $\mu = 8.3$ and $\mu = 6.7$, respectively, in the Traveler's Dilemma and a Bertrand price competition game. The implied λ of $1/8.3 = 0.12$ and $1/6.7 = 0.15$ are an order of magnitude lower than our estimates, but this is due partly to the differences in payoff scales. Capra et al. express payoffs in cents while we express payoffs in dollars, and the per-period earnings differ across these experiments anyway, even in real terms.

[12] Risk aversion was also advanced as an explanation of prices for the one-shot pricing experiments reported by Reynolds [19]. An alternative strategy is to introduce other-regarding preference parameters, as in Goeree and Holt [10]. As noted in the previous section, however, fairness considerations appear less important in our data compared to other studies. Moreover, Goeree and Holt need to introduce three new parameters to extend the QRE model for Fehr-Schmidt [7] inequality aversion utility, so this strategy is less parsimonious than the one pursued here.

Table 5. Maximum likelihood estimates for AQRE model with risk aversion* (data after period 15)

δ : Continuation probability	zero	30 percent Arizona	30 percent Purdue	30 percent Pooled	60 percent	90 percent	Pooled data
AQRE Model							
λ	4.4	1.0	2.0	1.4	1.9	2.0	1.7
$S.E.$	(0.38)	(0.10)	(0.15)	(0.08)	(0.18)	(0.11)	(0.08)
α	0.25	0.6	0.0	0.35	0.0	0.1	0.2
$S.E.$	(0.02)	(0.06)	(0.0)	(0.06)	(0.08)	(0.06)	(0.04)
$\mathcal{L}^*$	−459.4	−468.6	−326.5	−814.4	−684.2	−651.0	−2660.8
# observations	210	171	105	276	210	157	853

* Standard errors are in parentheses. Standard errors were computed using the bootstrap method. The sample size for the bootstrap was equal to the sample size of the data (see the last row in the table). 150 iterations were run to create a sample of bootstrap estimators.

the form, $u(c) = c^{1-\alpha}/(1-\alpha)$, where c is the dollar payoff for the (one round or two round) game and α is the index of relative risk aversion.

Table 5 reports estimates for the AQRE model with risk aversion. The estimated values of the risk aversion index α vary from a risk neutral level of zero for the 60 percent continuation treatment to moderately risk averse ($\hat{\alpha} = 0.25$) for the zero percent continuation treatment, to highly risk averse ($\hat{\alpha} = 0.6$) for the 30 percent continuation treatment conducted at Arizona. Using the same functional form for the utility function, Goeree et al. [11] estimate $\hat{\alpha} = 0.52$ in QRE estimates for independent private value first-price auction experiments. The variation in the risk aversion index is an indication that this model and utility function are too simple and are mispecified, but the variation arises largely from one or two anomalous treatments. A likelihood ratio test for our estimates strongly rejects the null hypothesis of risk neutrality in the zero and the Arizona 30 percent continuation treatments ($\chi^2_{1\text{ d.f.}} = 90.4$ and 52.6, respectively), but the estimates are consistent with risk neutrality in the other datasets.

Parameter estimates for the pooled data shown on the right of Table 5 indicate a relatively small amount of risk aversion and a moderate value for the noise parameter. Although this pooling "averages out" the significant differences across treatments discussed above, we use these pooled estimates to summarize the implications of this AQRE model with risk aversion. The rightmost column of Table 2 lists mean and median opening offer prices for the AQRE based on ML parameter estimates for the pooled data in Table 5. These predicted mean and median prices are hardly sensitive to changes in the continuation probability. In fact, median opening offer prices for the AQRE are equal to 30 for each continuation probability. This is consistent with our result (Sect. 4) that these prices are typically not significantly different across continuation probability treatments.[13]

[13] Median opening offer prices are also 30 for all treatments according to the fitted Noisy Nash Model with the pooled estimate of $\hat{\gamma} = 0.12$. But the NNM does not fit the bimodal distribution of offers as well as the AQRE model does, and the NNM provides an inferior fit for the second round offer prices. In particular, the fitted second round NNM mean and median prices range between 25 and 30, and are always further than the AQRE model from the observed second round offers.

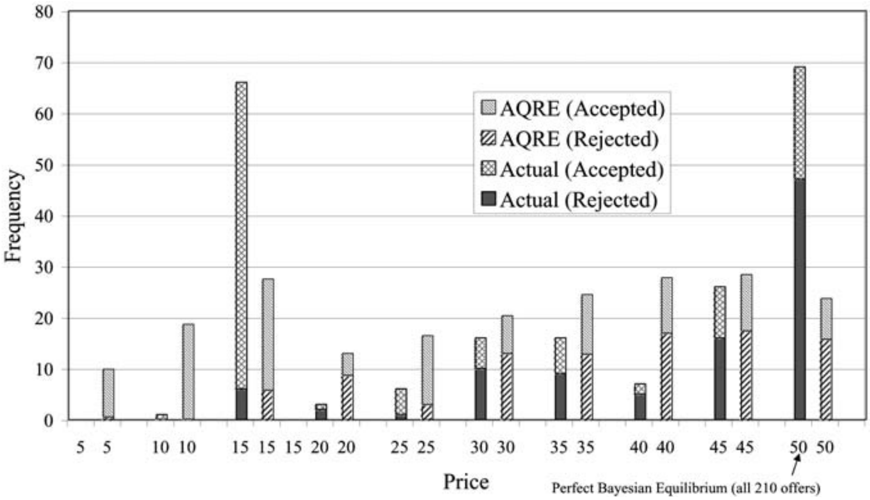

Figure 2. Actual opening prices and AQRE price distribution using pooled ML estimates: 0% continuation, data after period 15

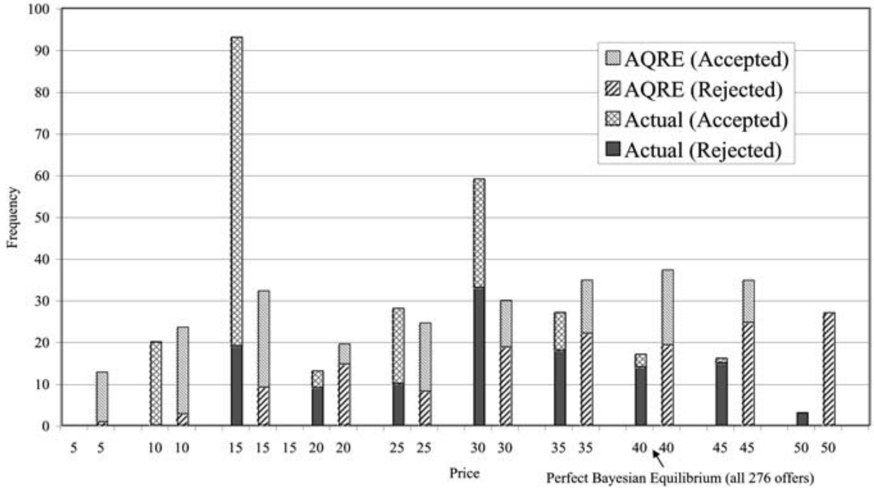

Figure 3. Actual opening prices amd AQRE price distribution using pooled ML estimates: 30% continuation, data after period 15

Figures 2 and 3 illustrate that the AQRE model with risk aversion is also consistent with the wide range of offer prices we observe. These figures show that the frequency distribution of opening offer prices typically varies between 10 and 50, and that the AQRE distribution based on the parameter estimates from the pooled data also varies widely. The AQRE has a significant fraction of offers at 15 (the highest price that both buyer types can profitably accept), but the mode of 15 for the observed offer prices is significantly stronger in all but the 90 percent continuation treatment (not shown).[14] The AQRE also implies more offers of 45 and 50 than observed, except in the one-round game (Fig. 2), and it also implies a smoother and more uniform distribution of prices than we observe. Overall, however, these AQRE distributions are consistent with many of the qualitative properties of the results – particularly when compared to the perfect Bayesian equilibrium.

6 Conclusion

This experiment explores the predictive power of a relatively standard but rather sophisticated game theoretic equilibrium concept – perfect Bayesian equilibrium – in multi-round pricing games. The results reveal significant failures of this theory in terms of levels of prices, comparative statics predictions, and buyer purchasing behavior. These prediction failures occur in spite of an experimental design that specifies a relatively simple decision-making environment.

We find that the agent quantal response equilibrium model of noisy decision-making in games captures important features of the results. In particular, this model correctly predicts two features of the experimental data: (1) variations in the continuation probability have relatively little impact on opening prices and (2) a wide (and in some cases bimodal) distribution of opening prices. Our implementation of this model also includes the possibility of subject risk aversion, and subjects appear to behave as if risk averse in some treatments. Other-regarding preferences have become a popular explanation of deviations from the standard equilibrium in settings like this two-person bargaining environment (e.g., Fehr and Schmidt [7]). But an implication of our results is that an appeal to fairness considerations is not the only way to explain these deviations, and that bounded rationality and risk aversion alone are consistent with a reasonably large proportion of the deviations from equilibrium observed in this game.

Appendix
Agent quantal response equilibrium

This section shows how the *logit*-AQRE may be computed for one and two round bargaining games. The following notation is used (numerical values used for the experiments are also indicated):

v is buyer value; $v \in \{\underline{v}, \bar{v}\}$ ($\underline{v} = 18,\ \bar{v} = 54$)
T is maximum number of trading rounds ($T \in \{1, 2\}$)
δ is continuation probability ($\delta \in \{0.3, 0.6, 0.9\}$ for $T = 2$)
p_t is price in round t ($p_t \in P \equiv \{\, 5, 10, ..., 45, 50 \,\}$)
θ is probability that $v = \bar{v}$ ($\theta = 0.4$)
$u(\cdot)$ is utility function, with dollar payoff as argument

[14] The fit displayed in these figures is improved substantially if we use the treatment-specific rather than the pooled parameter estimates. As Haile, Hortaçsu and Kosenock [15] have recently emphasized, however, it is important to leave the estimated value of the QRE parameter unchanged across treatments to make comparative statics exercises informative.

Buyer decision in final round T

This decision depends on the buyer's value, v, and the seller's price, p_T. If $v < p_T$ then the buyer is not permitted to purchase (according to rules of the experiments). If $v \geq p_T$ then the buyer's expected utility is as follows:

$$E\pi_T^b(1; v, p_T) = u(v - p_T), \quad \text{if buyer accepts offer} \tag{1a}$$
$$E\pi_T^b(0; v, p_T) = u(0) = 0, \quad \text{if buyer refuses offer} \tag{1b}$$

The first argument in the buyer's payoff function represents the decision; a one indicates acceptance of the price offer, a zero indicates rejection.

In a perfect Bayesian equilibrium (PBE) the buyer makes a purchase with probability one as long as $v > p_T$; there is no withholding in the final trading round in equilibrium. In an AQRE payoffs are adjusted by adding a zero-mean random term to the payoff associated with each decision. If these random terms are independent draws from an extreme value distribution with parameter $\lambda > 0$ then we have a *logit*-AQRE model. The buyer's probability of purchase in the final round in a *logit*-AQRE is given by,

$$q_T^b(v, p_T) \equiv \frac{\exp(\lambda E\pi_T^b(1; v, p_T))}{\exp(\lambda E\pi_T^b(0; v, p_T)) + \exp(\lambda E\pi_T^b(1; v, p_T))} = \frac{\exp(\lambda u(v - p_T))}{1 + \exp(\lambda u(v - p_T))} \tag{2}$$

if $v > p_T$. If $v < p_T$ then the probability of purchase is, $q_T^b(v, p_T) = 0$. Note that the purchase probability depends on the buyer's value and the seller's price, but is independent of any other market activity that might have preceded the final round.

Seller decision in single round model $(T = 1)$

The expected utility for a seller who sets a price $p \in P$ is,

$$E\pi_1^s(p) = (1 - \theta) q_1^b(\underline{v}, p) u(p) + \theta q_1^b(\bar{v}, p) u(p). \tag{3}$$

Note that this expected utility depends on the buyer's purchase probabilities. In a *logit*-AQRE the seller chooses price p with probability,

$$q_1^s(p) \equiv \frac{\exp(\lambda E\pi_1^s(p))}{\sum_{p' \in P} \exp(\lambda E\pi_1^s(p'))} \quad . \tag{4}$$

Equations (2) and (4) define choice probabilities for the *logit*-AQRE of the single round model. Given a value for the parameter λ and a specification of the utility function $u(\ )$, the purchase probabilities in (2) may be calculated, and then the seller's expected utility and choice probabilities in (3) and (4) may be calculated. It is not necessary to solve a fixed point problem for the single round model.

The two-round model

The choice probabilities for the buyer in the final round were derived in Eq. (2) above. The choice probabilities for a buyer in round two may be computed independently of other choice probabilities.

Let $q_2^s(p_2; p_1)$ be the probability that a seller sets $p_2 \in P$ after the buyer rejects $p_1 \in P$. The expected utility for a buyer with $v > p_1$ in round one is[15]:

$$E\pi_1^b(1; v, p_1) = u(v - p_1), \quad \text{if buyer accepts offer} \tag{5a}$$

$$E\pi_1^b(0; v, p_1) = \sum_{p_2 \in P} \delta q_2^s(p_2; p_1) q_2^b(v, p_2) u(v - p_2), \quad \text{if buyer refuses offer} \tag{5b}$$

The expected utility for a buyer who rejects an offer in round one depends on the continuation probability, the probabilities of various round two prices, acceptance probabilities for the buyer in round two, and the utility for the buyer for an accepted offer.

The expected utilities for the buyer determine the probability of purchase for the buyer in round one:

$$q_1^b(v, p_1) \equiv \frac{\exp(\lambda E\pi_1^b(1; v, p_1))}{\exp(\lambda E\pi_1^b(0; v, p_1)) + \exp(\lambda E\pi_1^b(1; v, p_1))} \tag{6}$$

if $v \geq p_1$. If $v < p_1$ then the probability of purchase is, $q_1^b(v, p_1) = 0$. There are 13 non-zero probabilities defined by Eq. (6); 3 for a low-value buyer and 10 for a high-value buyer.

Now consider the seller's price decision in round two, following a rejected offer in the first round. A key consideration is the way in which a rejected offer influences the seller's beliefs about the buyer's value. Using Bayes' Rule, the probability assessments of a buyer's value after p_1 is rejected are as follows:

$$\Pr[\underline{v}|\, p_1 \ rejected] = \frac{(1-\theta)(1 - q_1^b(\underline{v}, p_1))}{(1-\theta)(1 - q_1^b(\underline{v}, p_1)) + \theta(1 - q_1^b(\bar{v}, p_1))}$$

$$\Pr[\bar{v}|\, p_1 \ rejected] = \frac{\theta(1 - q_1^b(\bar{v}, p_1))}{(1-\theta)(1 - q_1^b(\underline{v}, p_1)) + \theta(1 - q_1^b(\bar{v}, p_1))}$$

The seller's expected utility in round two following a rejected offer is,

$$E\pi_2^s(p_2; p_1) = \Pr[\underline{v}|\, p_1 \ rejected]\, q_2^b(\underline{v}, p_2) u(p_2) + \Pr[\bar{v}\, |p_1 \ rejected]\, q_2^b(\bar{v}, p_2) u(p_2) \quad . \tag{7}$$

[15] If the buyer's value is less than the round one price, then the buyer does not make a decision. The experiment does not permit the buyer to purchase when value is less than price.

Note that the expected utility associated with price p_2 depends upon the particular round one price that was rejected. The round two price probabilities for the seller are given by,

$$q_2^s(p_2;p_1) \equiv \frac{\exp(\lambda E\pi_2^s(p_2;p_1))}{\sum\limits_{p_2' \in P} \exp(\lambda E\pi_2^s(p_2';p_1))} \quad . \tag{8}$$

There are 100 of these probabilities since there are 100 possible combinations of p_2 and p_1.

Buyer purchase probabilities in round one ($q_1^b(\cdot)$ defined in Eq. (6)) and seller price probabilities in round two ($q_2^s(\cdot)$ defined in Eq. (8)) are interdependent. A solution of the system defined by Eqs. (6) and (8) is a fixed point of a vector-valued function with 113 variables.

Once the probabilities in (6) and (8) are found, the seller's expected utilities and choice probabilities for round one may be derived. Expected seller utility as a function of round one price is:

$$\begin{aligned} E\pi_1^s(p_1) = (1-\theta)\Bigg[& q_1^b(\underline{v},p_1)u(p_1) \\ & +\delta(1-q_1^b(\underline{v},p_1)) \sum_{p_2 \in P} q_2^s(p_2;p_1) q_2^b(\underline{v},p_2) u(p_2) \Bigg] \\ & +\theta[q_1^b(\bar{v},p_1)u(p_1) + \delta(1-q_1^b(\bar{v},p_1)) \sum_{p_2 \in P} q_2^s(p_2;p_1) q_2^b(\bar{v},p_2) u(p_2)] \end{aligned} \tag{9}$$

Choice probabilities for round one prices are,

$$q_1^s(p_1) \equiv \frac{\exp(\lambda E\pi_1^s(p_1))}{\sum\limits_{p_1' \in P} \exp(\lambda E\pi_1^s(p_1'))} \quad . \tag{10}$$

Summary of computational steps for two-round logit-AQRE

a) Choose a positive λ parameter and a utility function, $u(\,)$.
b) Calculate the probabilities for buyer purchase in round two, using Eq. (2).
c) Solve the fixed point problem for round one buyer purchase probabilities and round two seller price probabilities, using Eqs. (6) and (8).
d) Calculate seller price probabilities for round one, using Eq. (10).

References

1. Bagnoli, M., Salant, S., Swierzbinski, J.: Durable-goods monopoly with discrete demand. Journal of Political Economy **97**, 1459–1478 (1989)
2. Capra, C.M., Goeree, J., Gomez, R., Holt, C.: Anomalous behavior in a traveler's dilemma? American Economic Review **89**, 678–690 (1999)

3. Capra, C.M., Goeree, J., Gomez, R., Holt, C.: Learning and noisy equilibrium behavior in an experimental study of imperfect price competition. International Economic Review **43**, 613–636 (2002)
4. Cason, T., Sharma, T.: Durable goods, coasian dynamics and uncertainty: theory and experiments. Journal of Political Economy **109**, 1311–1354 (2001)
5. Coase, R.H.: Durability and monopoly. Journal of Law and Economics **15**, 143–149 (1972)
6. Cooper, R., DeJong, D., Forsythe, R., Ross, T.: Selection criteria in coordination games: Some experimental results. American Economic Review **80**(1), 218–233 (March 1990)
7. Fehr, E., Schmidt, K.M.: A theory of fairness, competition, and cooperation. Quarterly Journal of Economics **114**, 817–868 (1999)
8. Fischbacher, U.: z-Tree – Zurich toolbox for readymade economic experiments. Working Paper No. 21, Institute for Empirical Research in Economics, University of Zurich (1999)
9. Fudenberg, D., Tirole, J.: Sequential bargaining with incomplete information. Review of Economic Studies **50**, 221–247 (1983)
10. Goeree, J., Holt, C.: Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. European Economic Review **44**, 1079–1089 (2000)
11. Goeree, J., Holt, C., Palfrey, T.: Quantal response equilibrium and overbidding in private-value auctions. Journal of Economic Theory **104**, 247–272 (2002)
12. Gul, F., Sonnenschein, H., Wilson, R.: Foundations of dynamic monopoly and the coase conjecture. Journal of Economic Theory **39**, 155–190 (1986)
13. Güth, W., Ockenfels, P., Ritzberger, K.: On durable goods monopolies: An experimental study of intrapersonal price competition and price discrimination over time. Journal of Economic Psychology **16**, 247–274 (1995)
14. Güth, W., Kröger, S., Normann, H.-T.: Durable-goods monopoly with privately known impatience: A theoretical and experimental study. Economic Inquiry (forthcoming) (2004)
15. Haile, P., Hortaçsu, A., Kosenok, G.: On the empirical content of quantal response equilibrium, Cowles Foundation Discussion Paper No. 1432, Yale University (2003)
16. McKelvey, R., Palfrey, T.: Quantal response equilibria in normal form games. Games and Economic Behavior **7**, 6–38 (1995)
17. McKelvey, R., Palfrey, T.: Quantal response equilibria for extensive form games. Experimental Economics **1**, 9–42 (1998)
18. Rapoport, A., Erev, I., Zwick, R.: An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. Management Science **41**, 377–394 (1995)
19. Reynolds, S.: Durable-goods monopoly: Laboratory market and bargaining experiments. RAND Journal of Economics **31**, 375–394 (2000)
20. Rosenthal, R.: A bounded-rationality approach to the study of noncooperative games. International Journal of Game Theory **18**, 273–292 (1989)
21. Stokey, N.: Rational expectations and durable goods pricing. Bell Journal of Economics **12**, 112–128 (1981)
22. Thaler, R.: Anomalies: the ultimatum game. Journal of Economic Perspectives **2**, 195–206 (1988)