A Simple Experimental Test of the Coase Conjecture: Fairness in Dynamic Bargaining

Jack Fanning
Brown University
jack_fanning@brown.edu

Andrew Kloosterman
University of Virginia
ask5b@virginia.edu

February 2, 2018

Abstract

In each round of an infinite horizon bargaining game, a proposer proposes a division of chips, until a responder accepts. The Coase conjecture predicts that incomplete information about responders' preferences for fairness leads to almost immediate agreement on an equal payoff split when discounting between rounds is small. We experimentally test this prediction when chips are equally valuable to both bargainers and when they are worth three times as much to proposers, and compare outcomes to an ultimatum game. Behavior offers strong support for the theory. In particular, when chips are more valuable to proposers, initial offers, initial minimum acceptable offers, and responder payoffs are significantly higher in the infinite horizon game than in the ultimatum game, while proposer payoffs are significantly lower.

Key words: Bargaining, Coase Conjecture, Fairness, Experimental Economics

JEL Codes: C78, C92

The authors would like to thank Charlie Holt, Tim Cason, Emin Karagozoglu, and seminar participants at Brown University, Georgia State University, Purdue University, the 15th Anniversary CESS conference at New York University, the 2017 ESA meetings in Richmond, the 2017 SEA meetings in Tampa, and the 2018 ASSA meetings in Philadelphia.

1 Introduction

The Coase conjecture (Coase (1972)) is central to modern bargaining theory. It predicts that when both agents are patient a seller who faces a buyer with private information about her value will sell her good almost immediately at close to the buyer's lowest possible value.[1,2] The seller's problem is that she must compete against her future self in an infinite horizon. For any price she offers today, high value buyers will have the greatest incentive to accept. If today's offer is rejected, therefore, the seller should likely

[1]So long as she obtains a positive profit from doing so, what is known as the "gap" case in the literature.

[2]In fact, Coase's original idea was that a durable goods monopolist facing downward sloping demand (an equivalent problem) would be forced to sell at (close to) marginal cost.

infer that she faces a low value buyer, creating an incentive to cut her price tomorrow. But in which case high value buyers should only purchase if today's price is already low (this is known as demand withholding).

A rigorous theoretical proof of Coase's idea was ultimately provided by Gul, Sonnenschein, and Wilson (1986). It serves as a foundational result for many recent and more complex bargaining models (e.g. two-sided reputational models in the style of Abreu & Gul (2000)). The core insight is ubiquitous: with patient bargainers, one-sided asymmetric information leads to almost immediate agreement favorable to the informed party (as though negotiation was exclusively with her "toughest" type).

Nonetheless, it is not obvious that bargaining in this setting will proceed this way. Another intuition is that asymmetric information will lead to inefficient delay, with an uninformed seller trying to hold out for a high price. Testing whether it holds in practice is therefore important. Laboratory experiments, providing a controlled environment, would seem ideally suited to that task. Discouragingly for the theory, results from the existing experimental literature have run almost entirely counter to even its qualitative predictions. In particular, the initial prices of an uninformed seller are typically increasing in the discount factor and game's horizon (see our literature review at the end of this section for details). Indeed, Reynolds (2000) concludes by saying: "It appears doubtful that any experimental design would generate results consistent with the Coase conjecture."

One potential confound of such results is that there is an additional source of asymmetric information. An extremely well-established finding in the bargaining literature, exemplified by the simple ultimatum game, is that some people care about fairness. Subjects routinely reject offers which give them a much smaller payoff than their opponent even though rejection leaves them with nothing (for a recent survey see Guth and Kocher (2014)). Preferences for fairness are heterogeneous and private, and therefore represent a second source of asymmetric information. Previous Coase conjecture experiments considered the canonical model of a seller facing a privately informed buyer, but their hypotheses were based on money maximizing subjects, and these need not hold even approximately in the presence of private information about fairness concerns. In particular, a seller who cares about fairness can resist rapidly dropping her price if she obtains sufficient disutility from selling to a high value buyer at an "unfair" low price.[3] Money maximizing sellers may then profitably imitate such behavior.

In this paper we provide a simple experimental test of the Coase conjecture by focusing exclusively on the asymmetric information about naturally occurring fairness preferences, which potentially confounded previous experiments. We consider an infinite horizon bargaining game, where in each round a proposer can propose any division of 100 chips

[3]See Fanning (2014).

between herself and a responder. If the responder accepts this offer the game ends, otherwise the value of chips to both parties are discounted by δ and the game continues into the next round. Each round, we use the "strategy method" to elicit responders acceptance decision by asking them for their minimum acceptable offer.[4]

If there is any positive probability that a responder is a fair type, that is, one who is unwilling to accept any unequal division, then for δ close to 1, the Coase conjecture predicts almost immediate agreement on an almost equal monetary payoff split (we choose $\delta = 0.95$). The translation of the Coasean reasoning is straightforward: if a proposer's offer of less than an equal monetary split is rejected, she should increase her belief that her opponent cares about fairness and so increase her offer in the next round. Anticipating this, even purely money motivated responders should reject offers that are not close to equal monetary payoffs.

The difficulty of statistically testing this prediction is that it is somewhat imprecise. How close to an equal payoff split is close enough? We therefore compare the infinite horizon game outcomes to those of ultimatum games (which are equivalent to $\delta = 0$ in infinite horizon games).

The Coasean theory predicts that initial offers and responder payoffs should be weakly higher in the infinite horizon game, proposer payoffs should be weakly lower, while initial minimum acceptable offers should be strictly higher (due to demand withholding). The weakly qualifier relates to the fact that if fair types are very likely then an equal payoff split offer can be money maximizing in the ultimatum game too. Indeed, the existing literature has documented that in ultimatum games where chips are equally valuable to proposers and responders (i.e. the standard ultimatum game), proposer offers are already close to an equal monetary payoff split.

Theory also predicts, however, that equal payoff split offers are less likely in ultimatum games when chips are worth more to the proposer than the responder (this increases the incentive to make an unequal offer). Moreover, that prediction is supported by previous experimental evidence from Kagel, Kim & Moser (1996). In addition to the standard equal chip value setting, therefore, we consider treatments where chips are worth three times as much to proposers as responders. We view this second comparison as offering the best test of the theory, where we expect the predictions noted above to hold strictly.

When chips are equally valuable to both parties, the weak inequalities of our predictions are indeed weak: initial offers and payoffs in the infinite horizon are statistically indistinguishable from the ultimatum game. Nonetheless, initial minimum acceptable offers

[4] We elicit responders' "strategy" in each round, not for the full infinite horizon game. We did this both in order to save time and to have a clear measure of responder demand witholding. Oosterbeek et al (2004) shows that this method tends to increase offers in ultimatum games, but, any such effect should not affect our comparisons of results across treatments that all use this same method.

are significantly higher in the infinite horizon (in line with theory's strict prediction). When chips are worth more to proposers than responders, however, all the treatment differences are strict: initial offers, initial minimum acceptable offers, responder payoffs and efficiency are significantly higher in the infinite horizon, while proposer payoffs are significantly lower. Moreover, the differences are qualitatively large. In particular, responder payoffs were 52% higher in the infinite horizon and proposer payoffs were 28% lower.

Not only are the treatment comparisons in line with the theory, but additionally our infinite horizon results are very close to its limiting point predictions (as $\delta \to 1$). Regardless of chip values, average initial offers, minimum acceptable offers, proposer payoffs and responder payoffs were less than 9% away from an equal monetary division. Remarkably, more than 92% of final chip allocations were within 10% of an equal division.

We believe these results represent strong evidence in favor of the Coase conjecture. In this simple bargaining environment, responders appear to have correctly understood that rejecting an unfair offer would likely lead to a proposer raising her offer in the next round. Anticipating the rejection of unfair offers, proposers made fair offers from the start of the game.

In addition to providing support for the theory generally, it also specifically suggests that fairness concerns may have large effects on bargaining outcomes more broadly. Even in settings where most people care little for fairness, the Coase conjecture implies that the mere possibility that they care can lead to an almost completely fair outcome (although of course determining what is fair may be more difficult in more complex settings). This may, for instance, help explain the findings of Cullen & Pakzad-Hurson (2017) in Task Rabbit data. When multiple workers privately contracted with an employer to do the same work in the same location, low paid workers frequently renegotiated their wages, and conditional on doing so obtained exactly the wage of the highest paid worker (the highest paid workers and those in different locations did not renegotiate). Co-location appears to have enabled workers to discover their peers' wages. Fairness concerns and the Coase conjecture can then explain why renegotiation fully equalized wages.

The remainder of this section discusses related literature. Section 2 then highlights our theoretical predictions, Section 3 provides details on the experimental design, Section 4 presents the results and Section 5 is a conclusion.

Related literature

Most Coase conjecture experiments have found evidence that contradicts the theory's comparative statics.[5] Guth, Ockenfels & Ritzberger (1995) matched a seller to 10 buyers with uniformly distributed values, for either two or three trading rounds and three different discount factor combinations. Subjects were either inexperienced or had training on the theoretical model. Behavior violated theoretical comparative statics on the discount factor and prices were typically above the static monopoly level (i.e. above the single round commitment price, far above the Coasean prediction). Rapoport, Erev & Zwick (1995) paired a seller to a single buyer with uniformly distributed values in a long horizon setting and three different discount factors. Contrary to theory, initial prices were increasing in the discount factor and above the static monopoly level. Reynolds (2000) compared treatments where a seller faced either one buyer or five buyers with uniformly distributed values, and either one, two or six trading rounds (and a constant discount factor). Again, contrary to theory, initial prices increased in the trading horizon, although they were below the static monopoly benchmark. Cason and Reynolds (2005) paired a seller with a buyer holding either a high or low value, for either one or two trading rounds with four possible continuation probabilities, and imposed a restricted grid of possible offers. Initial offers didn't vary much with the continuation probability and not generally in the direction suggested by theory. The authors argue that quantal response equilibrium captures some important features of their results.

On the other hand, other experiments related to the Coase conjecture offer at least partial support for the theory. Cason and Sharma (2001) considered a certain demand setting where one seller faced two buyers, one of whom always had a high value and the other a low value, and an uncertain demand treatment where with small probability both buyers had the same value. The bargaining horizon was indefinite, with treatments also varying the continuation probability that the game would continue for another round. Theory predicts high initial prices for certain demand regardless of continuation probability, and lower prices for uncertain demand, which are declining in the continuation probability (for Coasean reasons). However, the authors hypothesized that subjects would view both treatments as comparable and uncertain. Initial prices in both demand treatments were indeed closer to theory's predictions for uncertain demand. Those prices declined in the continuation probability for certain demand, in line with theory's prediction for uncertain demand, but they did not decline in the uncertain demand treatment. Guth, Kroger and Normann (2004) considered a two round bargaining game where a seller faces a buyer with uniformly distributed values, and both parties have private discount factors. They found

[5] In addition to the one-sided asymmetric information experiments highlighted here, some experiments consider two-sided asymmetric information (e.g. Embrey, Fréchette, and Lehrer (2015) investigate the reputational bargaining model of Abreu and Gul (2000)).

support for theory's prediction that initial prices are increasing in the seller's discount factor.

Slembeck (1999) investigated a repeated ultimatum game with a fixed partner (and equally valuable chips for both players), that might be expected to share a similar Coasean prediction with our infinite horizon game. Subjects only played the repeated game once. Average offers across all twenty rounds were close to an equal split, although this was also true in a control treatment where subjects were randomly rematched after each round. Rejection rates were significantly higher when partners were fixed, however, in line with our finding of higher minimum acceptable offers in the infinite horizon.

Related to our infinite horizon game with unequal chip values, Roth and Malouf (1979) considered a dynamic unstructured bargaining game where two subjects had to agree on a division of the probability of winning a prize, worth three times as much to one of the players. Both subjects could send freeform chat messages and propose any division at any time within 12 minutes. Agreements clustered around two distinct fairness norms, equal probability (50% probability for each subject) and equal expected payoffs (75% probability for the low prize subject). We did not expect an equal chip division in our setting to have a similar normative appeal to an equal probability of getting some prize (as opposed to nothing),[6] however, differences in the bargaining protocols of the two papers may also explain their different results.

2 Theory and hypotheses

We propose a very simple model, in which agents care about fairness to different extents. To simplify our analysis we consider a model with a continuous action space in which responders see proposers' offers before choosing to accept or reject them. Nonetheless, it is possible to show that the conclusions still hold in a fine, discrete action space game where responders simultaneously choose a minimum acceptable (MA) offer and the offer is accepted if and only if it is greater than or equal to the MA offer (the actual game in our experiment), in which there is a vanishingly small probability that agents tremble over their action choices.[7] Our solution concept is Perfect Bayesian equilibrium, which requires sequential rationality at all information sets and beliefs determined by Bayes' rule wherever possible. Off equilibrium path beliefs are unimportant for the analysis.

Each agent is a fair type with probability λ and is otherwise normal (money maximizing). A fair type's utility function when she receives $\$x_i$ and her bargaining partner receives

[6]Equal expected monetary payoffs may have appeared very unfair to risk averse high prize subjects, given their large chance of getting nothing.

[7]Without such trembles our use of the "strategy method" in each round prevents sequential rationality from having any bite.

\$$x_j$, is $u_i^F(x_i, x_j) = x_i - x_j \mathbb{1}_{[x_j > x_i]}$. This ensures that a fair type prefers disagreement to any agreement that gives her opponent a higher monetary payoff. A normal type's utility function is simply $u_i^N(x_i, x_j) = x_i$. It does not affect the analysis if proposers are assumed to always have normal preferences. An agreement is a division of chips for the proposer and the responder, (c_P, c_R) where $c_P = 100 - c_R \in [0, 100]$. If agents reach an agreement (c_P, c_R) in round $t \in \mathbb{N}$ (time is discrete), then monetary payoffs are $x_i = \delta^{t-1} e_i c_i$, where e_i is agent i's exchange rate of chips to money in round 1 and $\delta \in [0, 1)$ is the common discount factor. No agreement is captured by $t = \infty$.

The main game for our experiment has an infinite horizon. At the start of each round t (until the game ends), the proposer makes an offer, $c_{R,t} \in [0, 100]$, which the responder observes and then decides to accept or reject. If the offer is accepted, then the game ends and payoffs are received as previously stated. If the offer is rejected, the game moves to the next round. We also consider a baseline ultimatum game which is the same except that it always ends after round 1, and both players receive 0 if the offer is rejected. The ultimatum game is theoretically equivalent to the infinite horizon game with $\delta = 0$.

Let $\bar{c}_R = \frac{100 e_P}{e_R + e_P}$ be the offer of chips which implies equal payoffs for both players, and let $\bar{m} = \bar{c}_R e_R = (100 - \bar{c}_R) e_P$ be the implied monetary payoff (so, for example, $\bar{c}_R = 50$ if $e_P = e_R$ and $\bar{c}_R = 75$ if $e_P = 3e_R$). The next proposition establishes the Coase conjecture for the infinite horizon game. It says that there is a (generically) unique equilibrium. Moreover, for any positive probability of fair types $\lambda \in (0, 1]$, if players are sufficiently patient, i.e. $\delta \approx 1$, then there is almost immediate agreement on an equal payoff split. This is an approximately efficient outcome.

Proposition 1. *In the infinite horizon game there is a generically unique equilibrium path of play. For any $\varepsilon > 0$ and any $\lambda \in (0, 1]$ there is some $\bar{\delta}_{\varepsilon,\lambda} < 1$ such that if $\delta \geq \bar{\delta}_{\varepsilon,\lambda}$, in any equilibrium $c_{R,1} \geq \bar{c}_R - \varepsilon$, while for both players (and both types), payoffs are within ε of $\bar{m}$. If $\lambda = 0$ then $c_{R,1} = 0$, responder payoffs are 0 and proposer payoffs are $100 e_P$ (for both types).*

The proposition's proof is standard and is relegated to the Online Appendix. It involves first showing that for any $\lambda > 0$ a proposer must offer $\bar{c}_R$ within a finite number of rounds (such an offer must be accepted). This holds because the proposer can always guarantee a positive continuation payoff by offering $\bar{c}_R$, and so if she makes less generous offers, she must expect normal types to sometimes accept, but in which case the updated probability of facing a fair responder must increase and reach one. This induced finite horizon allows us to characterize a (generically) unique equilibrium by backward induction. The equilibrium features offers which slowly increase up to $\bar{c}_R$ (making normal types indifferent to waiting). The final step is to show that there is an upper bound, T, on the number of rounds until the proposer offers $\bar{c}_R$, that is uniform for all $\delta \in [0, 1)$. Given the option of waiting to accept, responder payoffs are at least $\delta^{T-1}\bar{m}$ (which is within ϵ of $\bar{m}$ for large δ).

The cutoff δ needed for the Coasean prediction to hold depends on the precise details of the model. Nonetheless, the general Coasean prediction of approximately immediate agreement, and approximately equal payoffs, for sufficiently high δ is seemingly robust to more general forms of fairness preferences (e.g. see Lopomo & Ok (2001)). In particular, this is the case when there is a continuum of types who differ continuously in their concern for fairness with only a positive density refusing any offer less than $\bar{c}_R$. The reasoning is identical: a proposer can guarantee a positive continuation payoff by offering $\bar{c}_R$ and so she must make this offer in a (uniformly) bounded number of rounds.

The Coase conjecture provides predictions for a number of outcomes of the infinite horizon bargaining game which we summarize here as Hypothesis APP (Approximate Point Prediction). We phrase the predictions for responders in terms of MA offers, the actual choice they make in our experimental game.

Hypothesis APP. *In the infinite horizon game, initial offers and initial MA offers and final agreed offers will be close to $\bar{c}_R$, payoffs of all agents will be close to $\bar{m}$, and any efficiency loss through bargaining will be small.*

As its name suggests, Hypothesis APP is somewhat imprecise. Exactly how close to the limit predictions the data need to be for the hypothesis to be satisfied is a matter of judgment. To provide a cleaner test, therefore, we compare theory's predictions in the infinite horizon game and the ultimatum game. In the ultimatum game, there is no guarantee of approximately equal initial offers, equal payoffs or an efficient outcome. Normal responders will accept any positive offer (a MA offer of 0), and so if $\lambda > \bar{\lambda} = \frac{e_P}{e_R+e_P}$, the proposer will offer $c_{R,1} = \bar{c}_R$, which all responders accept . For smaller λ, however, she will offer $c_{R,1} = 0$, which only normal types accept. To see this, notice that $\lambda > \bar{\lambda}$ if and only if $\bar{m} > (1-\lambda)e_P 100$.

This implies that initial offers, responder payoffs and efficiency should be weakly higher in the infinite horizon while proposer payoffs should be weakly lower, although the inequality need not be strict if fair types are very likely ($\lambda > \bar{\lambda}$).[8] Initial MA offers should always be strictly higher, however, due to demand withholding. While fair types set MA offers equal to $\bar{c}_R$ in both games, normal types set MA offers above 0 in the infinite horizon game and equal to 0 in the ultimatum game.

The possibility that offers are close to an equal monetary split even in the ultimatum game is not just a theoretical concern, this is typical in the literature when each agent faces the same chip to money exchange rate. Notice, however, that the cutoff $\bar{\lambda}$ is increasing

[8] In the proof of Proposition 1 we show that when $\lambda > \bar{\lambda}$, there is agreement in round 1 in the infinite horizon game too (not just almost immediate agreement).

in the ratio $\frac{e_P}{e_R}$. This is because a higher ratio requires a proposer to sacrifice more chips to persuade fair types to accept, making this strategy relatively less attractive compared to offering zero chips (payoff $e_P(100 - \bar{c}_R) = e_P 100 \frac{e_R}{e_R+e_P}$ vs $e_P 100(1-\lambda)$).[9] This is in line with the finding of Kagel, Kim and Moser (1995) that offers are far from an equal monetary split and there is significant disagreement when $e_P = 3e_R$. The Coasean prediction of Hypothesis APP is not affected by exchange rates, however, and strict treatment differences are more likely when $\frac{e_P}{e_R}$ is larger.

Hypothesis TC (Treatment Comparison) summarizes the predictions of the analysis above.

Hypothesis TC. *Initial offers, responder payoffs and efficiency will be weakly higher, proposer payoffs will be weakly lower and initial MA offers will be strictly higher in the infinite horizon game than in the ultimatum game. Strict differences in these variables are more likely when* $\frac{e_P}{e_R}$ *is larger.*[10]

Finally, we reiterate that our model of preferences shouldn't be taken literally. The literature is clear that in the ultimatum game (regardless of exchange rates) proposers do not offer $c_{R,1} = 0$, nor do (any) responders accept such offers. Our qualitative predictions, however, are robust to more complex models.

3 Experimental Design

We ran 12 sessions with a total of 236 subjects from the undergraduate student population at the University of Virginia. They were recruited through the VeconLab and Darden BRAD lab pools of students who had signed up to participate in experiments. No subject participated in more than 1 session. The two treatment variables were the horizon of the game, infinite or one-shot, and the exchange rates for chips, equal or 3 times as valuable to proposers. Using the standard 2×2, design, there were 3 sessions of each of the 4 treatments described in the following table.[11]

[9] This seems to be a robust prediction. Suppose the normal type was in fact modeled as in Fehr and Schmidt (1999) as a type with $u_i^N(x_i, x_j) = x_i - \alpha \max\{x_j - x_i, 0\}$ then she will accept any offer greater than $\underline{c}_R = \frac{100 e_P \alpha}{e_R(1+\alpha)+e_P\alpha}$ in the ultimatum game. A proposer should therefore offer $\underline{c}_R$ if $\lambda \leq \underline{\lambda} = \frac{\bar{c}_R - \underline{c}_R}{100-\underline{c}_R}$ (and $\bar{c}_R$ otherwise) where $\underline{\lambda}$ is again increasing in $\frac{e_P}{e_R}$.

[10] This comparison is robust to risk aversion. Responder risk aversion has no effect on the infinite horizon's deterministic equilibrium price path (responders receive the same *monetary* payoff whether they accept or reject on path offers below $\bar{c}_R$). Responders' MA Offers in the ultimatum game always leave them indifferent between accepting or rejecting (e.g. 0 if normal, $\bar{c}_R$ if fair). Risk aversion increases proposers' incentives to make offers closer to $\bar{c}_R$ in both games (so they are more likely to be accepted), but that merely makes the Hypothesis's inequalities more likely to be weak.

[11] We actually used exchange rates of either $e_P = e_R = \$0.25$ or $e_P = 3e_R = \$0.75$, but to simplify the exposition of the results we rescale these to $e_P = e_R = 1$ or $e_P = 3e_R = 3$.

Treatments

Name	Horizon	Exchange Rates	Subjects
1:1 IH	∞	$e_P = e_R = 1$	60
3:1 IH	∞	$e_P = 3$ and $e_R = 1$	60
1:1 UG	1	$e_P = e_R = 1$	60
3:1 UG	1	$e_P = 3$ and $e_R = 1$	56

In each match, subjects were matched in pairs and engaged in the infinite horizon or ultimatum bargaining game (depending on which treatment was run in their session). As previously noted, each round was implemented with simultaneous moves in which the proposer made an offer and the responder an MA offer, and the offer was accepted if and only if it was larger than the MA offer.[12] In the infinite horizon treatments, discounting was implemented by multiplying each side's chips by 0.95^{t-1} when agreement was reached in round t.[13]

Subjects received feedback about their own outcomes at the end of each round. In particular, a responder saw the proposer's offer, while a proposer only saw if their offer was accepted or rejected (but not the MA offer of the responder).

Each session consisted of 10 matches where the matching procedure was the turnpike design.[14,15] That is, the subjects were given a fixed role, proposer or responder, and then matched with every participant of the other role exactly once and in a way that if say proposer A matches to responder B who later matches to proposer C who then matches to responder D, proposer A will match with responder D before proposer C does.

All terms were neutrally framed. The experiment was programmed and run in z-tree (Fischbacher 2007). Subjects were paid for 1 randomly selected match and average earnings were $25.96.

4 Results

Our results are presented in three subsections. The first tests our two hypotheses, the second examines behavior in infinite horizon bargaining in greater detail, and third describes evidence for subject learning. Except for that third subsection we evaluate the

[12]Offers and MA offers could be made to the nearest one-hundredth of a chip.

[13]Potentially these games could have gone on forever. Fortunately, none did.

[14] In one session of Treatment 3:1 UG, recruitment error led to only 16 participants. In this session, only 8 matches were possible.

[15]After the experiment, we discovered that a Table in the instructions, which detailed the proposer-responder pairs implied by turnpike matching, contained typos from match 7 onwards. In particular, responder 8 appeared to be matched twice, while responder 6 was unmatched. We strongly believe that no subject noticed.

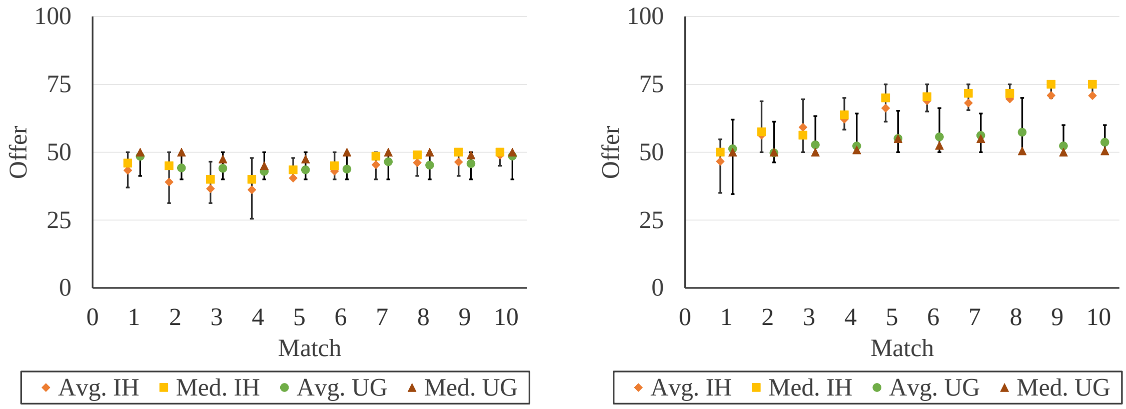

(a) 1:1 treatments (b) 3:1 treatments

Figure 1. Offers by Match

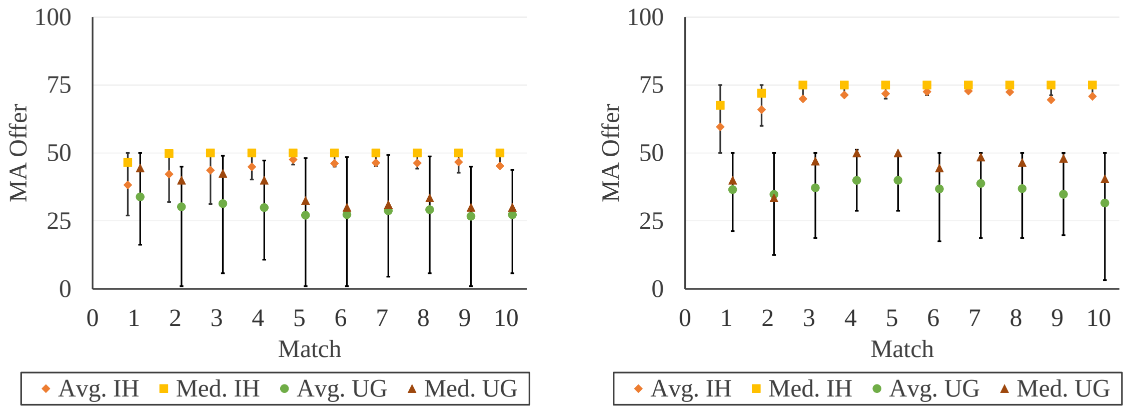

(a) 1:1 treatments (b) 3:1 treatments

Figure 2. MA Offers by Match

last five matches to allow subjects to first learn the game environment.[16]

4.1 Hypothesis tests

Figures 1-5 display means, medians, and the interquartile range for round 1 offers, round 1 MA offers, proposer payoffs, responder payoffs, and efficiency by match and treatment respectively. We illustrate variability with the interquartile range as much of the data is heavily skewed.

Simply eyeballing these figures tells the story of our experiment. They suggest first, that by the last 5 matches there is a clear separation between the treatments with unequal

[16] Guth, Ockenfels and Ritzberger (1995) speculate that the lack of opportunity to learn about their complicated game environment explains why theory performed so poorly in their experiment. In a different infinite horizon setting, the repeated prisoner's dilemma, a survey by Dal Bó and Fréchette 2017 finds that theory performs much better once subjects have had a chance to learn.

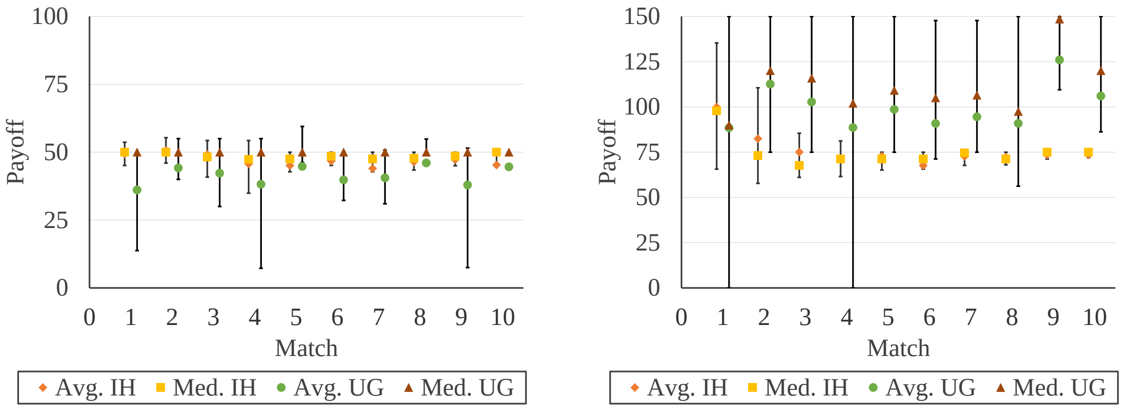

(a) 1:1 treatments (b) 3:1 treatments

Figure 3. Proposer Payoffs by Match

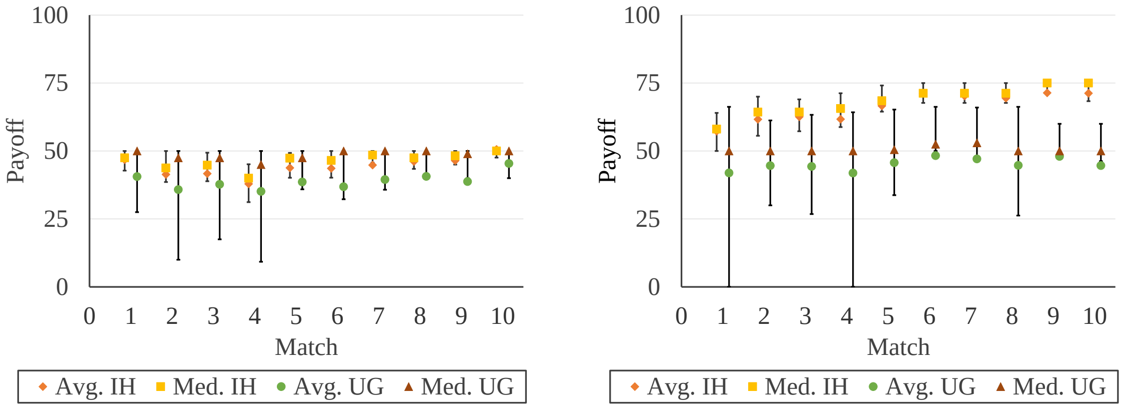

(a) 1:1 treatments (b) 3:1 treatments

Figure 4. Responder Payoffs by Match

exchange rates (3:1 IH and 3:1 UG) in the direction predicted by Hypothesis TC. Second, there is less (if any) separation between treatments with equal exchange rates (1:1 IH and 1:1 UG), consistent with Hypothesis TC. Third, in the infinite horizon games, regardless of the exchange rate, outcomes are close to the limiting Coasean prediction of an immediate equal division (recall that $\bar{c}_R = \bar{m} = 50$ in 1:1 IH and $\bar{c}_R = \bar{m} = 75$ in 3:1 IH), in line with Hypothesis APP. Overall, therefore, the data appear to offer strong support for the Coase conjecture. The rest of this subsection is devoted to more careful examination and statistical confirmation of these impressions.

We assess Hypothesis TC using Wilcoxon-Mann-Whitney tests on individual data match by match (with each subject contributing a single data point).[17] We call these Non-Parametric Individual (NPI) tests. It would not be correct to interpret a significant

[17] For payoffs, efficiency and final accepted offers, independence would be badly violated by averaging across matches 6-10. For example, Proposer A would directly affect the averaged payoffs of Responder B and C whom she is paired with in matches 6 and 7 respectively.

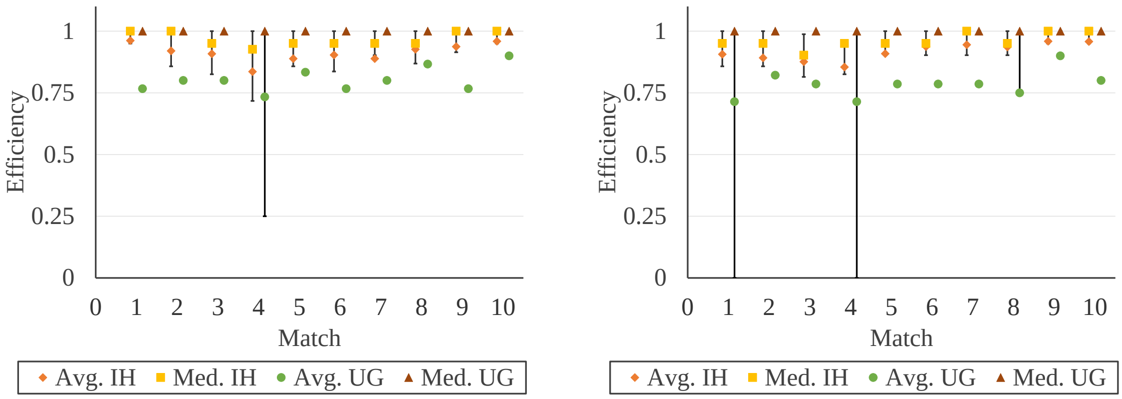

(a) 1:1 treatments (b) 3:1 treatments

Figure 5. Efficiency by Match

difference in any one match as clear evidence of a (non-random) treatment effect because multiple hypotheses are being tested, however, we typically find that either all 5 matches deliver significant differences or few do (this is particularly true for the choice variables offer and MA offer).

The turnpike matching design ensures that each subject should independently maximize utility in each match, without considering the effect of their behavior on future matches. However, to the extent that subjects learn about population behavior from previous matches independence might still be violated. We therefore check the robustness of the NPI tests in two ways. We consider similar Wilcoxon-Mann-Whitney tests using session-average data, which we call Non-Parametric Session (NPS) tests, as well as OLS regressions of individual data on a treatment dummy variable with standard errors clustered at the session level, which we call Parametric Session Cluster (PSC) tests.

Table 1: Number of Matches (from Matches 6-10)
with Significant Treatment Effects

	1:1 Treatments			3:1 Treatments		
Variable	NPI	NPS	PSC	NPI	NPS	PSC
Offer	0	0	0	5	5	5
MA Offer	5	5	5	5	5	5
Prop. Payoff	1	0	0(2)	5	3	3
Resp. Payoff	0	0	3	5	5	5
Efficiency	1	3	2	0	4	1

Table 1 reports the number of matches (out of matches 6-10) where there are significant differences ($p < .05$) between our treatments in the direction predicted by Hypothesis TC for our NPI tests as well as the NPS and PSC robustness tests. For each exchange rate, we ran 6 sessions with 10 subjects in each role per session, which implies 60 observations for each NPI and PSC test and 6 observations for each NPS test. Due to a recruitment issue detailed in footnote 14, however, for 3:1 exchange rates in matches 9 and 10, we only have 50 observations for each NPI and PSC test and 5 observations for each NPS test. We therefore use only a 10% significance threshold for those two NPS tests, the highest one feasible (when all sessions of one treatment are above all those of the other). When comparing proposer payoffs with 1:1 exchange rates we found significant differences in 2 matches for the PSC test that ran counter to theory's prediction (i.e. suggesting that proposers earned more in the infinite horizon). We report this number in Table 1 in parentheses.

We first assess the implications of these results for Hypothesis TC, when exchange rates were equal (1:1 IH vs 1:1 UG).

Confirming the impression of Figure 1(a) that round 1 offers are very similar for both game horizons, we find that there are no statistical differences for any of the last five matches for any of our statistical tests. Of course, theory allows for no strict treatment difference in offers, because if fair types are sufficiently likely, an equal split offer may be optimal even in the ultimatum game (this is why Hypothesis TC allows for weak differences). In fact, we can verify that this was the case. Given the empirical distribution of MA offers in the last 5 matches, the offer which maximizes expected proposer payoffs in the ultimatum game is exactly $\bar{c}_R = 50$.[18]

Even with no difference in offers, Hypothesis TC predicted that round 1 MA offers should be strictly higher in the infinite horizon than in the ultimatum game (due to demand withholding). This is indeed the impression of Figure 2(a). Table 1 confirms that the treatment difference is significant for all of the last 5 matches regardless of the statistical test. Moreover, the magnitude of the difference is large. Across the last 5 matches, the average MA offer is 66% larger in the infinite horizon than in the ultimatum game (46.17 vs 27.83). This suggests that most subjects who didn't care much about fairness, nonetheless seemed to understand the logic of the Coase Conjecture (i.e. that they could withhold demand to achieve close to an equal split in the infinite horizon).

Given that initial offers are close to an equal split for both game horizons it is unsurprising that Figures 3(a) and 4(a) show that both proposer and responder payoffs are very similar across game horizons. Only one (out of 15) tests in Table 1 indicates that proposers earned significantly more in the ultimatum game, while the PSC robustness test indicates significant results in the opposite direction in 2 (out of 5) matches. For responders, none

[18] Notice that while mean and median MA offers are approximately 25 in 1:1 UG, offering 25 would result in rejection about 50% of the time for an expected proposer payoff of approximately 37.5. Increasing the proposer's offer has two effects on her payoffs: it increases the probability of acceptance and lowers her payoff conditional on acceptance. The large dispersion of MA offers ensures that the net effect is positive up to an offer of 50, which is (almost) always accepted.

of the non-parametric tests are significant while the PSC robustness test gives significant differences in the predicted direction for 3 (out of 5) matches. We interpret these findings as indicating no treatment effect (or a minimal effect), which is nonetheless, still consistent with Hypothesis TC.

Figure 5(a) suggests that mean efficiency is slightly higher in the infinite horizon than in the ultimatum game. Averaging over the last 5 matches, efficiency is 82% in the ultimatum game compared to 92% in the infinite horizon. Nonetheless, Table 1 indicates few significant differences: 1 (out of 5) for the NPI tests and 5 (out of 10) for the robustness tests. This is perhaps not surprising given that the medians and interquartile ranges are tightly clustered around 100% for both treatments. However, it may also reflect the difficulty in statistically comparing values when efficiency is binary in the ultimatum game, and close to continuous in the infinite horizon.

We next consider treatments where chips are worth three times as much to proposers (3:1 IH vs 3:1 UG). Hypothesis TC suggests that strict treatment differences are more likely in this case. We find that not only do these predicted differences exist, but they are large in magnitude.

Figures 1(b) and 2(b) illustrate a clear separation between the treatments for offers and MA offers respectively. Averaging across the last 5 matches, round 1 offers are 26% higher in the 3:1 IH than in 3:1 UG (69.69 vs 55.29), while MA offers are 98% higher (71.64 vs 36.10).[19] The treatment differences are significant for every one of the last 5 matches for both variables under all statistical tests. For equal (1:1) exchange rates, we explained the lack of a treatment effect for offers (and by extension payoffs and efficiency) by the fact that equitable offers ($\bar{c_R} = 50$) were optimal in 1:1 UG given the distribution on MA offers. Similarly, we can explain the presence of a treatment effect with unequal (3:1) exchange rates by the fact that equitable offers ($\bar{c_R} = 75$) were not optimal in 3:1 UG. Rather, offers of 50 delivered the highest expected payoff.

Figures 3(b) and 4(b) illustrate the differences in payoffs. Averaging across the last 5 matches, proposer payoffs are 28% lower (71.95 vs 99.81) while responder payoffs are 52% higher (70.71 vs 46.57). Table 1 shows that for responders these differences are significant for all matches and for all statistical tests. For proposers, the differences are significant for all 5 matches using the NPI test, but are significant for only 3 (out of 5) matches for each of the robustness checks.

[19]The average MA Offer of 36.10 in 3:1 UG may seem quite low, indicating a responder's acceptance of an offer which gave her only 19% of the proposer's monetary payoff. Note, however, that the median MA Offer is higher at 47.5 (giving a responder at least 30% of the proposer's payoff) and that both mean and median are larger than in 1:1 UG (respectively 27.83 and 30). Comparing our results with Kagel, Kim & Moser (1996) is difficult because they did not use a strategy method, but there is some suggestion that our responders were less demanding. While offers in both experiments are broadly comparable, they report a higher rejection rate across all 10 matches (39% vs 22%), although the difference is lower for the last 7 matches (33% vs 22%).

Finally, Figure 5(b) shows that mean efficiency is again higher in the infinite horizon. Averaging across the last 5 matches, it is 95%, compared to 80% in the ultimatum game, which is larger than the difference between the equal exchange rate treatments (92% vs 82%). Nonetheless, our NPI test finds no significant differences between the treatments for any of the last 5 matches, while 5 (out of 10) of the robustness tests are significant. Again, this may reflect the fact that the interquartile ranges remain tightly clustered around 100% for both treatments, as well as the difficulty of statistically comparing a binary and (close to) continuous variable.

In summary, subject behavior is very much consistent with Hypothesis TC. While large strict treatment differences for some variables are only observed when the exchange rate ratio is 3:1 rather than 1:1, this is exactly what theory led us to expect.

We now turn to Hypothesis APP, which concerns infinite horizon outcomes exclusively. While the Hypothesis TC offered cleaner (clearly statistically falsifiable) tests of theory's predictions, Hypothesis APP is clearly also important. If infinite horizon outcomes were significantly closer to immediate agreement on a equal split than ultimatum game outcomes, but still in fact far from that limiting prediction (as $\delta \to 1$), it would seem to offer limited support for the Coase conjecture. We use Wilcoxon signed rank tests to statistically compare individual data match by match to these limiting theoretical benchmarks. Notice, however, that both the presence and absence of a significant difference can be consistent with Hypothesis APP so our main focus is on magnitudes.[20]

We go back to Figures 1 and 2, but now focus only on the infinite horizon treatments. Not only are means and medians of both variables close to $\bar{c_R}$ (an equal monetary split) for both exchange rates, but the interquartile ranges exhibit very little variation by the last few matches. For matches 6-10, the average offers are 45.96 in 1:1 IH and 69.69 in 3:1 IH, while the average MA offers are respectively 46.17 and 71.64. Each of these averages is less than 9% from $\bar{c_R}$. Statistically, offers are significantly less than $\bar{c_R}$ for every one of the last 5 matches for each treatment while MA offers are significantly less than $\bar{c_R}$ in 3 of the last 5 matches for each treatment. Whenever either variable differs from $\bar{c_R}$, they are almost always less than $\bar{c_R}$. Close to, but possibly slightly below, $\bar{c_R}$ is precisely what the Coase Conjecture predicts and so this is consistent with Hypothesis APP.

Hypothesis APP's prediction that payoffs should be close to $\bar{m}$ for both players also appears to be born out by Figures 3 and 4. The average proposer payoff in the last 5 matches is 45.96 and 71.95 for 1:1 IH and 3:1 IH respectively, while the average responder payoff is 46.32 and 70.71. Each average is less than 9% from $\bar{m}$. Moreover, again, the interquartile ranges indicate little variation. Statistically, payoffs are significantly less

[20]We do not additionally report these tests using session average data or use session clustered OLS regressions. Notice that the former could never yield a significant result given only 3 sessions per treatment.

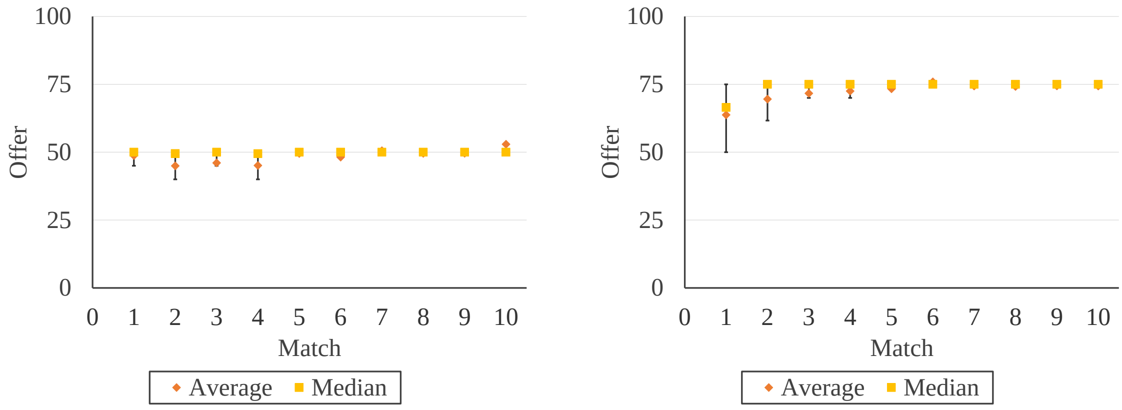

(a) Treatment 1:1 IH

(b) Treatment 3:1 IH

Figure 6. Accepted Offers by Match

than $\bar{m}$ for all of the last 5 matches for proposers in Treatment 1:1 IH and for responders in both treatments, while proposer payoffs in 3:1 IH are significantly different in only 2 matches. The difference between proposer and responder payoff in a match can also be directly compared, again using a (matched pairs) Wilcoxon signed ranks test. This difference is significant (and positive) for only 1 of the last 5 matches of each treatment which is consistent with the subjects achieving an equal monetary split. While the small differences in magnitude of payoffs from $\bar{m}$ is again in line with the Coasean prediction, it should be noted that theory would predict proposer payoffs slightly larger than $\bar{m}$ if not exactly equal (because proposers should always be able to offer $\bar{c}_R$ and have it immediately accepted).

Figure 5 exhibits evidence of minimal inefficiency for either exchange rate. Across the last 5 matches, average efficiency is 92% for 1:1 IH and 95% for 3:1 IH. Indeed, a majority of pairs reached agreement in round 1 and at least 70% did so by round 2 for both treatments. Statistically, efficiency is nonetheless significantly less than 100% for all of the last 5 matches in both treatments.[21] This is again consistent with Hypothesis APP.

We finish by reporting on one final variable, final accepted offers. These are the offers that end each bargaining game (there is no comparable variable for the ultimatum game). Match by match summary statistics are displayed in Figure 6. The average final accepted offer across the last 5 matches is 50.16 for 1:1 IH and 74.67 for 3:1 IH, less than 0.5% from $\bar{c}_R$. Moreover, there is very little variation around this average, highlighted by the degenerate interquartile ranges. Indeed, in 1:1 IH, 67% of accepted offers are exactly 50, and 92% are within 5 of 50. In 3:1 IH, 76% of accepted offers are exactly 75 and 97% are within 5 of 75. This lack of variation is quite remarkable. It is again in line with Hypothesis APP.

[21] Statistical significance is not surprising here as efficiency is bounded above by 100%, and so all deviations are negative.

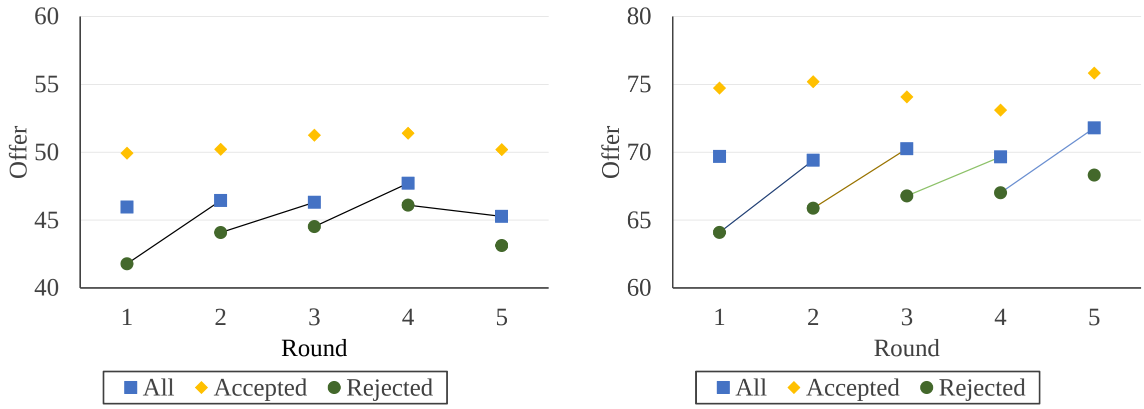

(a) Treatment 1:1 IH

(b) Treatment 3:1 IH

Figure 7. Average Offers by Round

In summary, we view the small differences in magnitude (less than 9%) between observed outcomes and an immediate agreement on an equal monetary split as strong evidence in support of Hypothesis APP and by extension the Coase conjecture. Combined with similar supportive findings for Hypothesis TC, we are left to conclude that the Coase conjecture seems to perform very well in this environment.

4.2 Infinite Horizon Bargaining

In this subsection, we explore subject behavior in the infinite horizon game in greater detail. As previously mentioned, a majority of matches reached agreement in round 1 for both infinite horizon treatments. Those matches which did not reach immediate agreement are of special interest, however, because theory makes additional predictions in this case. Moreover, comparing subject behavior to those predictions can help explain why most agreements were in fact immediate.

The unique equilibrium of the infinite horizon game, which provides the basis for the Coase conjecture, predicts that if a proposer offers less than $\bar{c}_R$ and this is rejected, then her subsequent offer should be larger (although not above $\bar{c}_R$). In the simple two type model developed in the theory section this rate of increase of offers is $\frac{1-\delta}{\delta} \approx 5\%$ in order to make money maximizing responders indifferent to waiting. With a richer set of "somewhat fair" types, the rate of increase may be smaller, but it should not be larger or else all responders would find it worthwhile to wait for $\bar{c}_R$.

Figure 7 plots the average offers in each round for all offers, accepted offers, and rejected offers for the first 5 rounds.[22] Lines are placed between all offers in round t and rejected offers in round $t-1$ to highlight the change in offer after an offer is rejected.

[22] 92% of matches last less than or equal to 5 rounds.

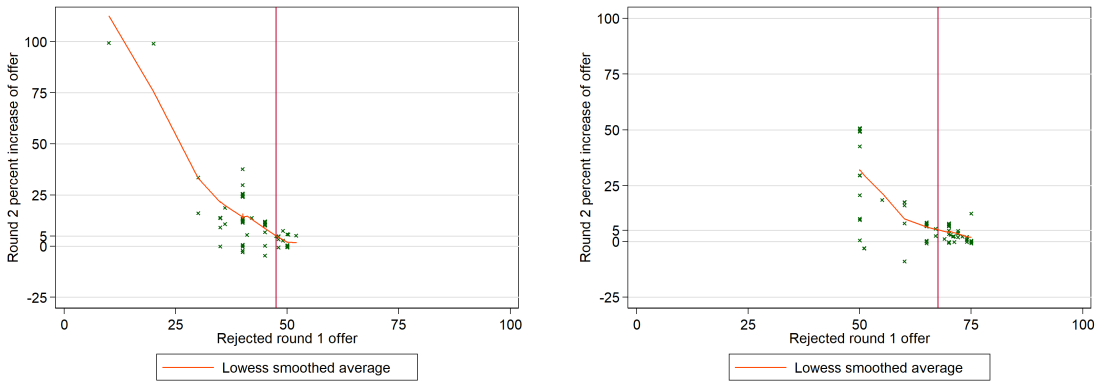

(a) Treatment 1:1 IH

(b) Treatment 3:1 IH

Figure 8. Percentage increase of round 2 offer by rejected round 1 Offer[23]

The figure illustrates several things. First, we confirm the previous result that accepted offers are almost exactly $\bar{c_R}$. This seems to hold regardless of how long the bargaining game lasts. Second, we see that proposers do indeed increase their offers after rejection. This is true for all but the change from round 4 to 5 in 1:1 IH (and that case accounts for only $n = 7$ observations).

The average increase of round 2 offers given a rejected round 1 offer is 11% in 1:1 IH (41.78 to 46.44) and 8% in 3:1 IH (64.09 to 69.41), although the rate is slightly lower between later rounds. This suggests that it was profitable on average for responders to reject the round 1 offers which they in fact rejected. It does not, of course, mean that it was profitable to reject all offers. To illustrate that it is low offers that should be rejected, Figure 8 plots rejected round 1 offers against the percentage increase of offer in round 2 as well as a smoothed lowess line of best fit. The lowess line is monotonically decreasing in a proposer's initial offer for both treatments indicating that small offers are likely to be substantially increased and therefore these offers should be rejected by responders. In particular, responders should reject round 1 offers that will be increased by 5%. In 1:1 IH the lowess line falls below 5% at 47.5, which is exactly $\delta\bar{c}_R$. Note that if proposers' never offer above $\bar{c}_R$ (and they rarely do) then money maximizing responders should always accept offers above $\delta\bar{c}_R$. In 3:1 IH the lowess line falls below 5% at 67.5. The high levels of these cutoffs illustrate the key insight of the Coase conjecture: it is very difficult for an uniformed party (here proposers) to commit not to substantially increase an offer, unless it is already close to $\bar{c}_R$, if it is rejected.

Theory's predictions for MA offers are less clear. In our simple two-type model, a money maximizing responder mixes over her acceptance decision. This mixing could potentially

[23]The 3:1 IH figure excludes one round 1 offer of 60 which fell 100% to 0. We add a uniformly distributed noise variable to the percentage increase variable (over [-1,1]) in order to make each scatter point visible, but still calculate the lowess line using the true percentage increase.

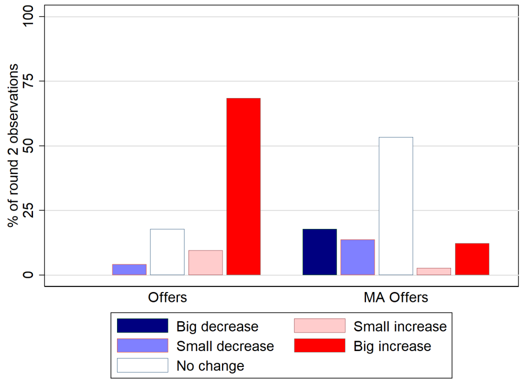

(a) Treatment 1:1 IH

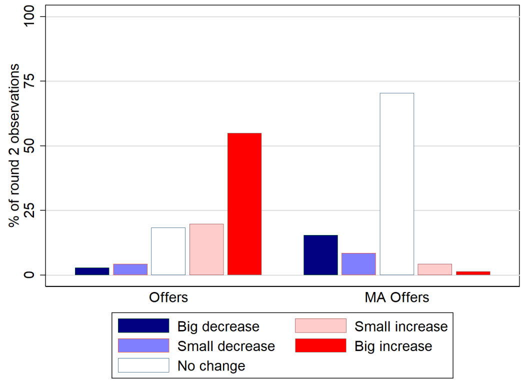

(b) Treatment 3:1 IH

Figure 9. Change in Offers and MA Offers from Round 1 to Round 2

be done at the population level, however, so that each responder's MA offers are constant. In fact, this was typically the case. Figure 9 shows the fraction of round 2 offers and MA offers which represent a big (strictly more than 5%) decrease, a small (between 5% and 0%) decrease, no change, a small increase and a big increase compared to round 1. If their round 1 MA offer resulted in rejection, 53% (in 1:1 IH) and 70% (in 3:1 IH) of MA offers did not change at all. On the other hand, 68% (in 1:1 IH) and 54% (in 3:1 IH) of rejected round 1 offers were increased by more than 5%. In line with predicted equilibrium dynamics, it seems that proposers' increased their offers following a rejection (eventually up to $\bar{c}_R$), while responders waited for that to happen. To statistically assess that impression, we compare the change in round 2 offers averaged at the proposer level (4.11 in 1:1 IH and 4.77 in 3:1 IH) with the negative change (change$\times(-1)$) in round 2 MA offers averaged at the responder level (0.78 in 1:1 IH and 2.41 in 3:1 IH). The null hypothesis of no difference between these changes is rejected at the 5% level (in fact, $p < .01$) in Wilcoxon Mann Whitney tests for both exchange rates.

Figure 8 shows that given proposer behavior, responders should have optimally rejected (round 1) offers which were not close to an equal monetary split. Figure 2 shows that on average responders did in fact do this. Given responder behavior, what was a proposers' optimal strategy? We can attempt an answer to this question by treating a proposer's round 1 offer as a proxy for her (infinite horizon) strategy. Figure 10 plots a proposer's payoff in a match against her round 1 offer as well as a smoothed lowess line. In both treatments the lowess line is initially monotonically increasing and then monotonically decreasing, reaching a maximum at exactly $\bar{c}_R$. Notice that although there is little data on offers above $\bar{c}_R$, payoffs are necessarily decreasing in this region so long as responders always accept those high offers (which they almost always do). Figure 1 established that proposers' average round 1 offers were indeed very close to $\bar{c}_R$. Such approximate optimality from both proposers and responders suggests that behavior was in fact not

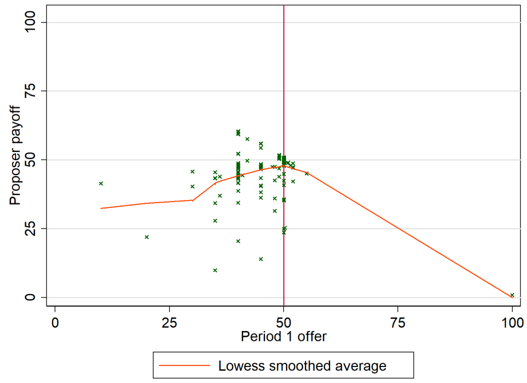

(a) Treatment 1:1 IH

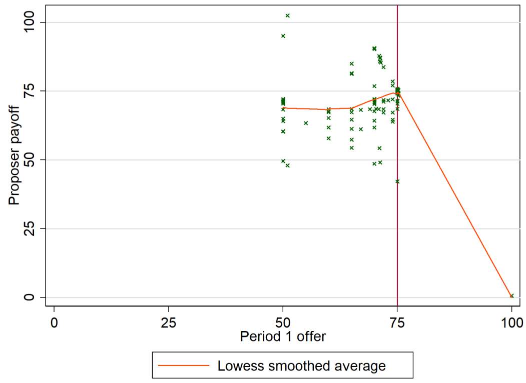

(b) Treatment 3:1 IH

Figure 10. Initial Offers and Proposer Payoffs[24]

very far from (a Coasean) equilibrium.

4.3 Learning

Finally, we investigate how behavior changed over the course of the experiment. Figures 1-6 suggest that learning was more important for 3:1 IH than for other treatments, and that much of this learning occurs in the first few matches.

Comparing the first and last match in 3:1 IH, proposers' round 1 offers increased by 52% (46.62 to 70.85), while responders' MA offers increased by 19% (59.62 to 70.83). A (matched-pairs) Wilcoxon signed ranks test finds that both changes are significantly different from zero. Proposer payoffs decreased 26% (99.98 to 73.60) from the first to the last match, while responder payoffs increased by 24% (57.28 to 71.25), with both changes significant. Average efficiency increased from 91% to 96%, which is marginally significant at the 10% level (for the efficiency change associated with proposers).[25]

Learning in 3:1 IH appears to occur in two stages. First, responder behavior converges to the Coasean prediction, and only then does proposer behavior adapt. Round 1 MA Offers (and final accepted offers) are already close to $\bar{c}_R = 75$ by match 3, with an essentially degenerate interquartile range. Round 1 offers, meanwhile take longer to converge (the interquartile range overlaps with $\bar{c}_R = 75$ for the first time only in match 5).

By contrast, in the other treatments, there was little evidence for learning. In 1:1 IH, the 13% increase of round 1 offers between the first and last match (43.33 to 48.88) is

[24]We add a uniformly distributed noise variable to the proposer's payoff variable (over [-1,1]) in order to make each scatter point visible, but still calculate the lowess line using the true payoff.

[25]For payoffs and efficiency, the assumption of independence may be quite suspect here, because Proposer A and B's change of payoff may both be affected by the behavior of Responder C (with whom they are paired with in matches 1 and 10 respectively).

significant at the 10% level, but all other changes are insignificant. In 1:1 UG, the 24% increase in proposer payoffs (36.07 to 44.60) is significant at the 10% level, but all other changes are insignificant. Finally in 3:1 UG the 11% increase of offers (51.25 to 56.80) is significant at the 10% level, but all other changes are insignificant.

It is perhaps not surprising that learning was most important in one of the infinite horizon treatments, as it probably the most complicated bargaining environment. Within the infinite horizon treatments, the fact that significant learning occurs only for the 3:1 exchange rate, is perhaps because of proposers' stronger temptation to make unequal payoff offers in this case (as in 3:1 UG). Only with experience did proposers seem to realize that such low offers were not worthwhile in the infinite horizon, and merely led to delay.

Table 3: Offer Learning Models

	(1)	(2)	(3)	(4)	(5)	(6)
	ΔOff	ΔOff	ΔOff	ΔOff	ΔOff	ΔOff
$\mathbb{1}_{rej}$	3.801**	16.86*	7.670	6.016**	32.04**	10.13*
	(0.616)	(5.429)	(6.098)	(1.314)	(5.455)	(2.663)
$\mathbb{1}_{acc}$	-2.970*	16.93	17.07	-2.530*	15.48	15.69
	(0.765)	(10.89)	(10.74)	(0.602)	(15.11)	(15.42)
$R1Off\mathbb{1}_{rej}$		-0.351	-0.169		-0.449**	-0.157**
		(0.131)	(0.144)		(0.0788)	(0.0182)
$R1Off\mathbb{1}_{acc}$		-0.423	-0.423		-0.253	-0.253
		(0.230)	(0.231)		(0.219)	(0.220)
# Rounds			-0.140			-0.203
			(0.209)			(0.289)
$AccOff{-}R1Off$			0.269*			0.392*
			(0.0759)			(0.0964)
N	270	270	270	270	270	270
adj. R^2	0.105	0.264	0.282	0.204	0.412	0.437

Matches 2-10

Session-clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To help further understand changes in subject behavior in the infinite horizon we turn to regression analysis. Table 3 displays regressions where the dependent variable is the change in round 1 offer between match t and match $t-1$. We start by including dummies for whether the round 1 offer was rejected or accepted in match $t-1$ (column (1) for Treatment 1:1 IH and (4) for Treatment 3:1 IH).[26] We expected that a previous rejection would cause players to increase their new round 1 offer, while acceptance would cause them to decrease it. Our second specification adds interaction terms between the size of the round 1 offer in match $t-1$ and the reject and accept dummies (column (2) for Treatment 1:1 IH and (5) for Treatment 3:1 IH). We expected that rejected offers would only be increased if they were low to begin with (i.e. below $\bar{c}_R$). Finally, we added variables for the number of rounds of bargaining in match $t-1$ and the difference between the final accepted offer and the round 1 offer in match $t-1$ (column (3) for Treatment 1:1 IH and (6) for Treatment 3:1 IH). We expected that both of these additional variables should have a positive effect on the new round 1 offer because they are associated with proposers learning about tough responders who don't accept low offers.

Table 4 displays regressions where the dependent variable is the change in round 1 MA offer between match t and match $t-1$. The first two specifications (columns (1) and (2) for Treatment 1:1 IH and (5) and (6) for Treatment 3:1 IH) are similar to their counterparts for the change of offers, just replacing the round 1 offer in match $t-1$ with the round 1 MA offer in match $t-1$, where applicable. In line with our previous predictions, we expected that rejection should cause MA offers to decrease, while acceptance would cause them to increase. The corollary to our prediction for offers would be that larger MA offers in match $t-1$ are associated with larger decreases in that MA offer after rejection. However, while this might seem reasonable for money maximizing types, a larger MA Offer in match $t-1$ may also indicate a greater concern for fairness and thus less willingness to reduce the MA offer. An inherent asymmetry between the roles is that responders got more detailed feedback: they observed offers, while proposers didn't observe MA offers. Our third specification therefore adds the proposer's round 1 offer in match $t-1$ and the % change in that offer from round 1 to round 2, which is recorded as zero if the offer was accepted (column (3) for Treatment 1:1 IH and (7) for Treatment 3:1 IH). We expected that both of these variables would have a positive effect, because they should be associated with responders learning that proposers (are willing to) make larger offers. Our final regression adds the number of rounds of bargaining in match $t-1$ and the difference between the MA offer in the final round of bargaining and in round 1 offer in match $t-1$ (column (4) for Treatment 1:1 IH and (8) for Treatment 3:1 IH). These are comparable to the final outcome variables in the change of offer regressions. In line with our predictions for those regressions, we expect that both variables should have negative effects.

[26] We suppress the constant and run with both dummies to ease the exposition, but of course, could also omit one dummy and run the regressions with a constant to get the same results.

Table 4: MA Offer Learning Models

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$	$\Delta MA\ Off$
$\mathbb{1}_{rej}$	-1.144	18.61***	19.79**	16.92*	-0.0424	35.94	40.25*	41.20*
	(0.763)	(1.773)	(2.459)	(5.636)	(0.320)	(13.99)	(11.98)	(11.39)
$\mathbb{1}_{acc}$	2.941**	7.347	9.392	10.09	3.269**	18.37	27.34*	29.13*
	(0.401)	(4.009)	(5.294)	(4.184)	(0.527)	(13.26)	(8.822)	(7.426)
$R1\ MA\ Off \mathbb{1}_{rej}$		-0.372***	-0.354***	-0.270*		-0.483	-0.393	-0.364
		(0.0252)	(0.0308)	(0.0656)		(0.194)	(0.181)	(0.185)
$R1\ MA\ Off \mathbb{1}_{acc}$		-0.125	-0.110	-0.105		-0.244	-0.188	-0.179
		(0.0860)	(0.0836)	(0.0893)		(0.195)	(0.202)	(0.208)
$R1 Off$			-0.0545	-0.0754			-0.174	-0.205
			(0.0786)	(0.0635)			(0.0724)	(0.0982)
% Change $R1$ to $R2$			-0.00244	-0.00468			-0.0606	-0.0737
			(0.00455)	(0.00275)			(0.0455)	(0.0575)
# Rounds				0.0917				-0.212
				(0.169)				(0.154)
$Acc\ MA\ Off – R1\ MA\ Off$				0.167				0.0390
				(0.133)				(0.0931)
N	270	270	269	269	270	270	270	270
adj. R^2	0.064	0.179	0.177	0.194	0.042	0.229	0.250	0.249

Matches 2-10

Session-clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

What do the regression results for the change of offers (Table 3) imply? First, as expected, offers increase significantly after rejection, by about 4 in Treatment 1:1 IH and about 6 in Treatment 3:1 IH. On the other hand, offers decrease by about 3 in both treatments after acceptance, although this is only marginally significant. The size of the previous round 1 offer has the expected negative effect when that offer was rejected in both treatments although this is significant only in Treatment 3:1 IH. The size of the previous round 1 offer also has a negative effect after acceptance, but this isn't significant. Counter to expectations, the number of rounds of bargaining has a negative estimate, although it is not significant. The differences between the final accepted offer and the round 1 has the expected positive effect, although it is only marginally significant for both treatments.

Turning to Table 4, we see that in line with our predictions, MA offers increased significantly after acceptance in both treatments (by about 3). While the coefficients on dummies for rejection also have the expected (negative) sign, they are not significant. Our second specification shows that larger MA offers decrease more after a rejection (as expected for money maximizing responders), although this is significant only in 1:1 IH. All other variables are not significant, including perhaps most notably the size of the previous round 1 offer and the percentage increase in round 2 offer (the coefficients aren't even in the predicted directions). Few significant results is perhaps partially due to less variation in MA offers across matches.

An interesting story which emerges from these learning regressions is that subjects behavior after rejection and acceptance is not symmetric. Proposers significantly increase their offer after rejection, but only marginally significantly decrease it (and this decrease is less in magnitude) after acceptance. Responders significantly increase their MA offer after acceptance, but don't significantly decrease it after rejection. This suggests a ratchet effect whereby each acceptance causes MA offers to increase (leaving offers unchanged) and each rejection causes offers to increase (with MA offers unchanged) until subjects reach the limiting Coasean outcome of an immediate agreement on a equal split.

5 Conclusion

It is often difficult to obtain enough experimental control to convincingly test important theoretical results in lab. Subjects' heterogenous, unobserved preferences for fairness represent an important potential confound for game theoretic predictions. In this paper, rather than ignoring such naturally occurring preferences, we embraced them, and sought to directly use this source of private information to test the Coase conjecture. For an infinite horizon bargaining game in which proposers make all offers, the theory predicts that patient players will agree almost immediately on close to an equal monetary split.

This prediction seems to perform very well. Initial offers, initial minimum acceptable offers, proposer payoffs, responder payoffs and efficiency are all less than 9% away from the immediate equal split benchmark (theory's limit prediction as players become infinitely patient).

One problem with our theoretical prediction is that it lacked precision (how close is close?). We, therefore, also compared the infinite horizon game outcomes to those in an ultimatum game. Our finding that it is impossible to statistically distinguish outcomes (apart from minimum acceptable offers) for the two different game horizons when chips are equally valuable to both players is consistent with the theory (because immediately offering an equal split may always be profitable when fair types are very prevalent). According to theory, however, strict treatment differences are more likely when chips are worth more to proposers than responders, which is exactly what we find. Not only are initial offers, initial minimum acceptable offers and responder payoffs significantly higher in the infinite horizon, and proposer payoffs significantly lower, but the magnitude of the treatment differences is large. We view our findings as strong evidence in support of the basic message of the Coase conjecture, that with one sided asymmetric information the uniformed party must "give in" to the informed party almost immediately (as if she was bargaining with the informed party's toughest type).

It might be argued that while the Coase conjecture performs well in our narrowly defined setting, this is at the expense of sacrificing its reach. As a result, we may not be able to make a clear prediction about what will happen in Coase's original setting of a seller facing a buyer who is privately informed about her value (because the real world is messy and has additional sources of private information, including those about fairness preferences). Two responses are in order. First, bargaining theorists have long acknowledged that clear predictions are rare when there are multiple sources of private information. Some progress has been made, however, in two-sided private information models (e.g. Abreu & Gul (2000)), and those models are still built on Coasean foundations (with one-sided private information). In this light, our experiment is important because it shows that those foundations appear secure. Second, our result highlights how fairness preferences in particular, may play an outsized role in real world bargaining situations. Even in settings where people are unlikely to care much about fairness (e.g. CEOs of multinational corporations), the mere possibility that they do combined with the Coase conjecture may be enough to profoundly affect outcomes.

References

[1] Abreu, Dilip and Gul, Faruk. Bargaining and Reputation. *Econometrica* 2000, 68(1), pp. 85-117.

[2] Cason, Timothy and Sharma, Tridib. Durable Goods Monopoly, Coasean Dynamics, and Uncertainty: Theory and Experiments. *Journal of Political Economy* 2001, 109(6) pp. 1311-1354.

[3] Cason, Timothy and Reynolds, Stanley. Bounded Rationality in Laboratory Bargaining with Asymmetric Information. *Economic Theory* 2005, 25, pp. 553-574.

[4] Coase, Ronald. Durability and Monopoly. *The Journal of Law and Economics* 1972, 15(1), pp. 143-149.

[5] Cullen, Zoe and Pakzad-Hurson Bobak. Equilibrium Effects of Pay Transparency. *Working Paper* 2017.

[6] Dal Bó, Pedro and Fréchette, Guillaume. On the Determinants of Cooperation in Infinitely Repeated Games: A Survey. *Journal of Economic Literature* 2017, forthcoming.

[7] Embrey, Matthew, Fréchette, Guillaume, and Lehrer, Stephen. Bargaining and Reputation: Experimental Evidence on Bargaining in the Presence of Irrational Types *The Review of Economic Studies* 2015, 82(2), pp. 608-631.

[8] Fanning, Jack. Bargaining, Fairness, and Commitment. *Working Paper* 2014.

[9] Fehr, Ernst and Schmidt, Klaus. A Theory of Fairness, Competition and Cooperation. *Quarterly Journal of Economics* 1999, 114, pp. 817-868.

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

[11] Gul, Faruk, Sonnenschein, Hugo, and Wilson, Robert. Foundations of Dynamic Monopoly and the Coase Conjecture. *Journal of Economic Theory* 1986, 39(1), pp. 155-190.

[12] Guth, 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.

[13] Guth, Werner, Kroger, Sabine, and Norman, Hans-Theo. Durable-Goods Monopoly with Privately Known Impatience: A Theoretical and Experimental Study *Economic Inquiry* 2004, 42(3), pp. 413-424.

[14] Guth, Werner, Ockenfels, Peter, and Ritzberger, Klaus. On Durable Goods Monopolies: an Experimental Study of Intrapersonal Price Competition and Price Discrimination over Time. *Journal of Economic Psychology* 1995, 16(2), pp. 247-274.

[15] Kagel, John, Kim, Chung, and Moser, Donald. Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs. *Games and Economic Behavior* 1996, 13(1), pp. 100-110.

[16] Lopomo, Giuseppe and Ok, Efe. Bargaining, Interdependence, and the Rationality of Fair Division. *RAND Journal of Economics* 2001, 32(2), pp. 263-283.

[17] Oosterbeek, Hessel, Sloof, Randolph, and Van de Kuilen, Gijs. Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis. *Experimental Economics* 2004, 7, pp. 171-188.

[18] Reynolds, Stanley. Durable-Goods Monopoly: Laboratory Market and Bargaining Experiments. RAND Journal of Economics 2000, 31(2), pp. 375-394.

[19] Rapoport, Amnon, Erev, Ido and Zwick, Rami. An Experimental Study of Buyer-Seller Negotiation with One-sided Incomplete Information and Time Discounting. *Management Science* 1995, 41(3), pp. 377-394.

[20] Roth, Alvin and Malouf, Michael. Game-Theoretic Models and the Role of Information in Bargaining. *Psychological Review* 1979, 86, pp. 574-594.

[21] Slembeck, Tilman. Reputations and Fairness in Bargaining - Experimental Evidence from a Repeated Ultimatum Game With Fixed Opponents. *Discussion Paper* 1999.

6 Online Appendix (not for publication)

Proof of Proposition 1

We first claim that both proposers and responders must accept $\bar{c}_R$ immediately if offered. To see this, let $\bar{c}$ be the supremum of any offer made in any equilibrium after any history and suppose that $\tilde{c} > \bar{c}_R$. In that case both responder types certainly accept any offer strictly above $\max\{\bar{c}_R, \delta\tilde{c}\} < \tilde{c}$. But in which case, the proposer would never find it optimal to make offers arbitrarily close to $\tilde{c}$ as strictly lower offers would be accepted, a contradiction which implies $\tilde{c} \leq \bar{c}_R$. Moreover, both types must accept $\bar{c}_R$ if offered because their best possible continuation payoff after rejection is $\delta\bar{c}_R < \bar{c}_R$. Identical logic implies that the proposer always offers 0 when $\lambda = 0$, and this is accepted.

Next define a series of equilibrium offers $q(n) = \delta^n \bar{c}_R$, which will apply when there are n rounds of bargaining remaining. Notice that this makes a normal responder indifferent between accepting immediately or waiting for an additional round.

The equilibrium probability that the responder is fair with n rounds of bargaining remaining is $\eta(n)$, and the proposer's value function divided by her exchange rate e_P is $W(n)$. We shall refer to this as simple her value function (rescaling by e_P in this way does not affect the analysis). A simple application of Bayes' rule implies that if beliefs are given by $\eta(n)$ today and $\eta(n-1)$ tomorrow and only normal types accept, then a fraction $\frac{\eta(n-1)-\eta(n)}{\eta(n-1)}$ of responders do in fact accept today. The variables $\eta(n)$ and $W(n)$ are defined recursively from $\eta(n-1)$ and $W(n-1)$ starting with $\eta(0) = 1$ and $W(0) = 100 - \bar{c}_R$ using the equality:[27]

$$\begin{aligned}
&\frac{\eta(n-1)-\eta(n)}{\eta(n-1)}(100 - q(n)) + \frac{\eta(n)}{\eta(n-1)}\delta W(n-1) \\
=&\frac{\eta(n-1)-\eta(n)}{\eta(n-1)}(100 - q(n-1)) + \frac{\eta(n)}{\eta(n-1)}W(n-1)
\end{aligned} \tag{1}$$

with $W(n)$ defined as the value at which equality is obtained. This means that $\eta(n)$ represents the belief at which the proposer is just indifferent between offering $q(n)$ today followed by $q(n-1)$ tomorrow, or immediately offering $q(n-1)$. Although the lower offer may be desirable other things equal, it involves some delay. Meanwhile, $W(n)$ simply represents good accounting for the proposer's value function. Using the equality

[27] The second line of this equality is equal to $W(0)$ if $n = 1$ and $\frac{\eta(n-2)-\eta(n)}{\eta(n-2)}(100 - q(n-1)) + \frac{\eta(n)}{\eta(n-2)}\delta W(n-2)$ if $n = 2$.

$q(n-1) - q(n) = (1-\delta)q(n-1)$ and rearranging the above equation gives:

$$\frac{\eta(n-1) - \eta(n)}{\eta(n-1)} = \frac{W(n-1)}{q(n-1) + W(n-1)} \tag{2}$$

Plugging this in to evaluate $W(n)$ gives:

$$W(n) = \frac{100W(n-1)}{q(n-1) + W(n-1)} \tag{3}$$

We are now ready to characterize the equilibrium.

Lemma 1. *There is a generically unique equilibrium path. If $\lambda \in (\eta(N), \eta(N-1))$ then bargaining must finish by round N. On the equilibrium path offers are:*[28]

$$c_{R,t} = q(N-t) \text{ for } t \in \{1, ..., N\}$$

Beliefs that responder is a fair type are:

$$\mu_t = \eta(N-t) \text{ for } \text{ for } t \in \{2, ..., N\}$$

Notice that equation 2 implies $\eta(1) = \frac{\bar{c}_R}{100} = \frac{e_P}{e_R+e_P}$, so that if $\lambda > \bar{\lambda} = \frac{e_P}{e_R+e_P}$ then we must have $c_{R,1} = \bar{c}_R$, as claimed in the main text.

Given the characterization of the lemma the proof of the Proposition is almost immediate. Notice that $W(n-1) \geq 100 - \bar{c}_R$ and $q(n-1) \leq \bar{c}_R$ imply that the probability of acceptance in each round is bounded away from zero:

$$\frac{\eta(n-1) - \eta(n)}{\eta(n-1)} = \frac{W(n-1)}{q(n-1) + W(n-1)} \geq \frac{100 - \bar{c}_R}{100} > 0$$

Hence, given any $\lambda > 0$, Bayesian updating then implies for $t \geq 1$ that $\mu_{t+1} \geq \mu_t \frac{100}{\bar{c}_R} \geq \lambda \left(\frac{100}{\bar{c}_R}\right)^t \geq 1$. Clearly, therefore, bargaining must end before round $T = \left\lceil \frac{ln\left(\frac{1}{\lambda}\right)}{ln\left(\frac{100}{\bar{c}_R}\right)} \right\rceil + 1$ for any δ, or else $\mu_{T+1} > 1$. Given this, we must have that $c_{R,1} \geq q(T-1) = \delta^{T-1}\bar{c}_R$, and responder payoffs for both types are certainly greater than $\delta^{T-1}\bar{c}_R e_R$ (given their option of waiting to accept the offer $\bar{c}_R$), while proposer payoffs are certainly greater than $\bar{m}$ (given the option to offer $\bar{c}_R$ immediately). Choosing $\bar{\delta}_\lambda < 1$ appropriately, therefore, gives the result (notice in particular that if the responder expected payoff is greater than $\bar{m} - \epsilon$ then the proposer's payoff must be smaller than $\bar{m} + \frac{e_P \epsilon}{e_R}$).

Proof of Lemma 1

[28]For $\lambda = \eta(N)$ there are two possible price and belief paths one corresponding to the equilibrium where $\lambda \in (\eta(N), \eta(N-1))$, the other corresponding to the equilibrium where $\lambda \in (\eta(N+1), \eta(N))$.

First, we claim that bargaining must end if the proposer ever offers at least $\bar{c}_R$. Let $\bar{q}$ be the supremum of equilibrium offers and initially suppose that $\bar{q} > \bar{c}_R$. In which case, there exists some $\varepsilon \in (0, \bar{q} - \bar{c}_R)$ such that $\bar{q} - \varepsilon > \delta\bar{q}$ so that all responders should accept, but in which case the proposer should never make an offer above $\bar{q} - \varepsilon$, a contradiction. Hence, $\bar{q} \leq \bar{c}_R$, and if the proposer every offers $\bar{c}_R$ in equilibrium, it must certainly be accepted (as $\bar{c}_R > \delta\bar{c}_R$).

A history h_t is sequence of past offers. Let $\mu_{t+j}(h_t)$ and $c_{R,t+j}(h_t)$ describe beliefs and offers in round $t + j$ consistent with some continuation equilibrium after h_t (such that offers have not yet hit $\bar{c}_R$, ending bargaining). The proposition then follows from the statements below, which are proved by induction on n.

1. If $\mu_t(h_t) > \eta(n)$ then bargaining last at most n more rounds including the current one, $c_{R,t+j}(h_t) \geq q(n - 1 - j)$ for $j \in \{0, ..., n - 1\}$ and $\mu_{t+j}(h_t) \geq \eta(n - 1 - j)$ for $j \in \{1, ..., n - 1\}$.
2. If $\mu_t(h_t) \in (\eta(n), \eta(n - 1))$ then the proposer's continuation value is given by:

$$V_{n-1}(\mu_t(h_t)) = \frac{\eta(n-1) - \mu_t}{\eta(n-1)}(100 - q(n-1)) + \frac{\mu_t}{\eta(n-1)}W(n-1)$$

3. If $\mu_t(h_t) \leq \eta(n)$ and $c_{R,t} \in (q(n), q(n - 1))$ then $\mu_{t+1}(h_t, c_{R,t}) = \eta(n - 1)$.
4. If $\mu_t(h_t) \leq \eta(n)$ and $c_{R,t} < q(n)$, then $\mu_{t+1}(h_t) \leq \eta(n)$.
5. If $\mu_t(h_t) = \eta(n)$ then the proposer's continuation value is $W(n)$.
6. If $\mu_t(h_t) < \eta(n)$ then bargaining lasts at least $n + 1$ more rounds including the current one with $c_{R,t+j}(h_t) \leq q(n - j)$ and $\mu_{t+j}(h_t) \leq \eta(n - j)$ for $j \in \{0, ..., n\}$.

Consider the above claims for $n = 0$, and let $\eta(-1) = 1$ and $W(-1) = q(-1) = 100 - \bar{c}_R$. 1) and 2) and 3) are then true immediately because their antecedents cannot be activated. 4) is true because $\mu_{t+1}(h_t, c_{R,t}) \leq 1 = \eta(0)$. For 5), recall that any offer of at least $\bar{c}_R$ must be accepted. If $\mu_t(h_t) = 1$ then the best possible continuation value the proposer can obtain is from proposing $\bar{c}_R$ immediately, because the fair responder will not accept less. This gives a continuation value of $W(0) = 100 - \bar{c}_R$. For 6) given that the game has not ended it is immediate that if bargaining lasts at least 1 round, and given that both types will immediately accept an offer of $\bar{c}_R = q(0)$ we must have $c_{R,t}(h_t) \leq q(0)$.

Given that the claims are true for arbitrary $(n - 1)$ we proceed to show that they must also be true for n, with $n > 0$. The true statements for (n-1) referred to by **1),...,6)** while the statements to be proven for n are instead referred to by **1'),...,6')**.

1') we need only consider $\mu_t(h_t) \in (\eta(n), \eta(n-1)]$ because for $\mu_t(h_t) > \eta(n-1)$ this is true by claim **1**. We first make a subclaim that there can be at most a finite number of rounds j such that $\mu_{t+j}(h_t) \leq \eta(n-1)$. This is ultimately because the proposer always has the option of offering (fractionally more than) $\bar{c}_R$, which will be accepted and thus ensures a continuation value of *at least* $100 - \bar{c}_R$.

Suppose there is an equilibrium offer path which lasts an infinite number of rounds without agreement and gives the proposer $100 - \bar{c}_R$ but never has $\mu_{t+j}(h_t) > \eta(n-1)$ for any $j \in \mathbb{N}$. Let $a_{t+s}(h_t)$ be the equilibrium implied probability of acceptance in round $t+s$ (conditional on h_t). Given that the proposer can obtain $100 - \bar{c}_R > 0$ for any $\mu_t(h_t) \in (\eta(n-1), \eta(n)]$ we must have:

$$\lim_{m \to \infty} \sum_{j=0}^{m} \delta^j a_{t+j}(h_t) 100 \geq 100 - \bar{c}_R$$

However, for N such that $\delta^N < \frac{100-\bar{c}_R}{200}$ we must have:

$$\sum_{j=0}^{N} a_{t+j}(h_t) > \frac{100 - \bar{c}_R}{200}$$

And so in N rounds at least fraction $\frac{100-\bar{c}_R}{200}$ of the responders accept (all normal types). The same argument can be repeated k times implying that:

$$\mu_{t+kN}(h_t) \geq \mu_t \left(\frac{200}{100 + \bar{c}_R} \right)^k$$

The right hand side of this equation is greater than 1 for large k, giving a contradiction, and so eventually we must have $\mu_{t+j}(h_t) > \eta(n-1)$.

Let t' be the supremum of times such that $\mu_{t'}(h_t) \leq \eta(n-1)$. For $n = 1$ bargaining must end in round t' with $c_{R,t'} = \bar{c}_R$. For $n > 1$, we must have $\mu_{t'+1}(h_t) > \eta(n-1)$. This immediately implies that $c_{R,t'+1}(h_t) \geq q(n-2)$ by claim **1)**. If $c_{R,t'} < q(n-1)$ the normal type responder would then strictly prefer to wait an extra round to obtain the lower price (because $c_{R,t'} < \delta q(n-2)$), which would imply $\mu_{t'}(h_t) = \mu_{t'+1}(h_t)$ a contradiction. This means that $c_{R,t'}(h_t) \geq q(n-1)$.

If $\mu_{t'}(h_t) < \eta(n-1)$ we must have $c_{R,t'+j}(h_t) \leq q(n-j)$ for $j \leq n$ by claim **6)** and so in conjunction with the argument of the previous paragraph $c_{t'+j} = q(n-j)$. This in turn implies the proposer's continuation payoff in round t' is given by $V_{n-1}(\mu_{t'}(h_t))$ where this is defined in equation 3. If $\mu_{t'}(h_t) = \eta(n-1)$ on the other hand, then the proposer's continuation value is $W(n-1) = V_{n-1}(\mu_{t'}(h_t))$ by claim **5)**.

Now suppose $t' \neq t$ and consider time $t'-1$ where $\mu_{t'-1}(h_t) \in (\eta(n), \eta(n-1)]$. For the proposer's offer to be accepted with positive probability it must be that $c_{R,t'-1}(h_t) \geq q(n)$

or else the responder would wait for the offer of $c_{R,t'}(h_t) = q(n-1)$. But in which case this strategy cannot obtain the proposer the continuation value of $V_{n-1}(\mu_{t-1})$ which we know she could obtain by offering (fractionally more) than $q(n-1)$. To see this notice that her continuation value in the supposed equilibrium can be written as follows:

$$
\begin{aligned}
&\frac{\mu_{t'}(h_t)-\mu_{t'-1}(h_t)}{\mu_{t'}(h_t)}(100-c_{R,t'-1}(h_t))+\frac{\mu_{t'-1}(h_t)}{\mu_{t'}(h_t)}\delta V_{n-1}(\mu_{t'}(h_t))\\
=&\frac{\mu_{t'}(h_t)-\mu_{t'-1}(h_t)}{\mu_{t'}(h_t)}(100-c_{R,t'-1}(h_t))+\frac{\mu_{t'-1}(h_t)}{\mu_{t'}(h_t)}\frac{\eta(n-1)-\mu_{t'}(h_t)}{\eta(n-1)}\delta(100-q(n-1))+\frac{\mu_{t'-1}(h_t)}{\eta(n-1)}\delta W(n-1)\\
\leq&\frac{\eta(n-1)-\mu_{t'-1}(h_t)}{\eta(n-1)}(100-q(n))+\frac{\mu_{t'-1}(h_t)}{\eta(n-1)}\delta W(n-1)\\
<&V_{n-1}(\mu_{t'-1}(h_t))
\end{aligned}
$$

Where the first inequality follows because of the $c_{R,t'-1}(h_t) \geq q(n)$, and the second inequality follows from the fact that $\mu_{t'-1}(h_t) > \eta(n)$, and from remembering that $\eta(n)$ is defined as the lowest belief at which the third line would be equal to the fourth (see equation 1). And so finally we have a contradiction, proving that $t' = t$. We have established that if $\mu_t(h_t) \in (\eta(n), \eta(n-1)]$ then $c_{R,t}(h_t) \geq q(n-1)$ and the proposer's continuation payoff must be at least $V_{n-1}(\mu_{t'}(h_t))$. Furthermore, $\mu_{t+1}(h_t) \geq \eta(n-2)$ for $n > 1$ or else the proposer would not obtain $V_{n-1}(\mu_{t'}(h_t))$. The claims for $c_{R,t+j}(h_t) \geq q(n-j)$ and $\mu_{t+j}(h_t) \geq \eta(n-1-j)$ for $j \geq 1$ then follow from claim **1)**.

2') This is simply accounting given claim **1')** and **6)**.

3') By assumption $\mu_t(h_t) \geq \eta(n)$ and $c_{R,t} \in (q(n), q(n-1))$. Suppose then that equilibrium acceptance of this offer implies $\mu_{t+1}(h_t, c_{R,t}) < \eta(n-1)$, then by claim **6)** $c_{R,t+j+1}(h_t, c_{R,t}) \leq q(n-j)$ for $j > 0$, but in which case all high types would optimally accept $c_{R,t}$ and so $\mu_{t+1}(h_t) = 1$. But that in turn implies $c_{R,t+1}(h_t) = \bar{c}_R$. This presents a contradiction for $n > 1$ and immediately implies the claim **3')** for $n = 1$. In either case, however, it must be true that $\mu_{t+1}(h_t) \geq \eta(n-1)$.

If on the other hand $\mu_{t+1}(h_t, c_{R,t}) > \eta(n-1)$ for $n > 1$ then by claim **1)** we have $c_{R,t+1}(h_t, c_{R,t}) \geq q(n-2)$, but in which case it is not optimal to accept $c_{R,t}$ in round t, and so $\mu_t(h_t, c_{R,t}) = \mu_{t+1}(h_t) < \eta(n-1)$, a contradiction. And so, $\mu_{t+1}(h_t, c_{R,t}) = \eta(n-1)$.

4') By assumption $\mu_t(h_t) \leq \eta(n)$ and $c_{R,t} < q(n)$. Suppose that $\mu_{t+1}(h_t, c_{R,t}) > \eta(n)$. By claim **1')** this implies that $c_{R,t+1}(h_t, c_{R,t}) \geq q(n-1)$, and so the normal type of responder would not optimally accept, ensuring $\mu_t(h_t) = \mu_{t+1}(h_t, c_{R,t}) \leq \eta(n)$, a contradiction.

5') If $\mu_t(h_t) = \eta(n)$, then the proposer can obtain at least the value $V_{n-1}(\eta(n)) = W(n)$ as given by equation 3, by offering (fractionally more than) $c_{R,t} = q(n)$ or $c_{R,t} = q(n-1)$ (by claims **3**, **3'** and **5**). Offering any particular $c_{R,t} \in [0, q(n)) \cup (q(n), q(n-1))$ cannot obtain this value. In particular, notice that offering $c_{R,t} < q(n)$ implies $\mu_{t+1}(h_t, c_{R,t}) = \mu_t(h_t)$ by claim **4')**, and so it simply delays proceedings, and means the proposer cannot obtain a payoff as great as $W(n)$. This means the proposer's equilibrium continuation value must be exactly $W(n)$.

6') Suppose $\mu_t(h_t) < \eta(n)$, then by Claims **3') and 5** the proposer can guarantee a value $V_n(\mu_t(h_t))$ by offering (fractionally more than) $q(n)$. Notice that $V_n(\mu_t(h_t)) > V_{n-j}(\mu_t(h_t))$ for $j > 0$ (for $j = 1$ this follows given how $\eta(n)$ is defined), which implies that offering strictly more than $q(n)$ cannot obtain such a utility. This then must imply that $c_{R,t} \leq q(n)$.

If the time t offer is not accepted with positive probability then $\mu_{t+1}(h_t) = \mu_t(h_t) < \eta(n-1)$ and so claim **6')** is true immediately due to claim **6)**. If on the other hand this time t offer is accepted with positive probability but $c_{R,t+j} > q(n-j)$ for some $j > 0$, then it is not optimal for the normal responder to accept the time t offer, (she prefers instead to wait until time $t + j$), contradicting a positive probability of acceptance. Finally if $\mu_{t+j}(h_t) > \eta(n-j)$ for some $j > 0$ then this would imply $c_{t+j}(h_t) > q(n-j)$ by claim **1)**, but we have just claimed that this leads to a contradiction.

This proves statements **1')** through **6')**. These being true for all n implies that any equilibrium must have the properties described in the Lemma. To prove that such an equilibrium does exist, all that remains is to specify is off path strategies and beliefs for the proposer, the responder's strategy being fully outlined above.

Suppose that $\mu_t \in (\eta(n), \eta(n-1)]$. If the proposer makes an offer, which is strictly less than $q(n)$, then as specified above no responder accepts. Given the responder continuation strategies laid out above then the proposer is content to follow the price path $(q(n-1), q(n-2), ..., q(0))$ from the next round onward making the responders rejection optimal.

If the proposer instead offers $c_{R,t} \in (q(n-k), q(n-k-1)]$ with $k \geq 0$ then as specified above, a responder will accept so that $\mu_{t+1} = \eta(n-k-1)$. Given responder continuation strategies the proposer is then indifferent between charging the offer path $((q(n-k-1), q(n-k-2), ..., q(0))$ or the path $((q(n-k-2), q(n-k-3), ..., q(0))$. By mixing and following the first price path with probability $\frac{c_{R,t} - \delta q(n-k-2)}{\delta(q(n-k-1) - q(n-k-2))}$ and the second with the complimentary probability, we can ensures that responders' behavior is optimal. If an offer above $\bar{c}_R$ is rejected, we can specify that $\mu_t(h_t) = 1$ forever afterwards and the proposer optimally offers $\bar{c}_R$, which is always accepted by both types.

Strategies then represent optimal choices for both the responder and proposer by construction at every possible history given the other's strategy, and beliefs are determined by Bayes' rule wherever possible, hence this is an equilibrium.

Experimental Instructions for 3:1 IH

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. Make sure you pay close attention to the instructions because the choices you make will influence the amount of money you will take home with you today. Please ask questions if any instructions are unclear.

The Experiment

The experiment will consist of 10 identical MATCHES that each proceed as follows. First, you will be matched with one other participant from the room, and one of you will be in the role of PROPOSER while the other will be in the role of RESPONDER. There are 100 CHIPS that will be split between you and the other participant in the following way.

In ROUND 1, the PROPOSER makes an OFFER between 0 and 100 CHIPS (offering partial CHIPS up to 2 decimal places is allowed) which indicates how many of the 100 CHIPS the RESPONDER will receive if the OFFER is ACCEPTED. Simultaneously, the RESPONDER chooses a LOWEST ACCEPTABLE OFFER between 0 and 100 (again, up to 2 decimal places is allowed) which indicates the fewest number of CHIPS they are willing to ACCEPT. If the OFFER is greater than or equal to the LOWEST ACCEPTABLE OFFER, then the OFFER is considered ACCEPTED in which case the RESPONDER receives the number of CHIPS OFFERED while the PROPOSER receives the remaining CHIPS and the MATCH ends.

Alternatively, if the OFFER is less than the LOWEST ACCEPTABLE OFFER, then the OFFER is considered REJECTED and we proceed to ROUND 2. In ROUND 2, the PROPOSER chooses a new OFFER and the RESPONDER a new LOWEST ACCEPTABLE OFFER. Just as before, if the OFFER is greater than or equal to the LOWEST ACCEPTABLE OFFER, then the OFFER is considered ACCEPTED in which case the RESPONDER receives the number of CHIPS OFFERED while the PROPOSER receives the remaining CHIPS and the MATCH ends. Likewise, if the OFFER is less than the LOWEST ACCEPTABLE OFFER, the OFFER is considered REJECTED and we proceed to ROUND 3. ROUND 3, and all subsequent ROUNDS, follow an identical procedure as that of the first and second ROUNDS. The MATCH continues until an OFFER is ACCEPTED.

Although each ROUND of the MATCH proceeds identically, the value of the CHIPS depends on which ROUND the OFFER is ACCEPTED in. Specifically, the dollar value of a CHIP falls by 5% each ROUND for both PROPOSERS and RESPONDERS. If the OFFER is ACCEPTED in ROUND 1, the CHIPS are worth 75 cents for PROPOSERS and 25 cents for RESPONDERS. However, if the OFFER is ACCEPTED in ROUND 2, the CHIPS are worth 5% less. That is, only 71.25 cents each for PROPOSERS ($0.95 \times 75 = 71.25$) and 23.75 cents each for RESPONDERS ($0.95 \times 25 = 23.75$). Likewise, in ROUND 3, the CHIPS are worth 5% less again. That is, 67.69 cents for PROPOSERS ($0.95 \times 0.95 \times 75 = 67.69$) and 22.56 cents for RESPONDERS ($0.95 \times 0.95 \times 25 = 22.56$). In general, the value of CHIPS if the OFFER is ACCEPTED in ROUND n is found by multiplying the original value of a CHIP by 0.95 (n-1) times. That is, $75 \times 0.95^{n-1}$ for PROPOSERS and $25 \times 0.95^{n-1}$ cents for RESPONDERS.

At the end of each ROUND, we will show each pair the PROPOSERs OFFER and reveal whether the OFFER was ACCEPTED or REJECTED. When the match ends, we will also show the earnings of both participants.

Although you will participate in 10 MATCHES, we will only pay you for your earnings in one MATCH which will be randomly determined by the computer at the end of the experiment. In addition to any earnings from CHIPS you will earn 12 dollars for participating. Since it is impossible to earn partial cents, the computer will run a lottery to determine if we round up or down to the nearest cent according to the fraction of a cent which you earned. For example, if you are a RESPONDER who earns 22.67 CHIPS in ROUND 1 then you should earn \$17.6675 ($12 + 0.25 \times 22.67 = 17.6675$) but since we cant give you 0.75 cents you will get \$17.67 with probability 0.75 and \$17.66 with probability 0.25.

Details of matching procedure

At the beginning of the experiment you will be assigned the role of PROPOSER or RESPONDER and you will remain in that role for all 10 MATCHES. Since there are 20 participants this means there will be 10 participants assigned to each role. There are 10 MATCHES and 10 participants in the opposite role to your own. You will be matched with each of those 10 participants in exactly one MATCH. The exact procedure for how you will be matched is described as follows. Lets name the PROPOSERS P1, P2, , and P10 and the RESPONDERS R1, R2, ., and R10 (these names are only for the purpose of explaining the procedure and you wont actually be given one of these names in the experiment). The following table shows who will be matched with who in each MATCH.

Match 1	Match 2	Match 3	Match 4	Match 5	Match 6	Match 7	Match 8	Match 9	Match 10
P1, R1	P1, R2	P1, R3	P1, R4	P1, R5	P1, R6	P1, R7	P1, R8	P1, R9	P1, R10
P2, R2	P2, R3	P2, R4	P2, R5	P2, R6	P2, R7	P2, R8	P2, R9	P2, R10	P2, R1
P3, R3	P3, R4	P3, R5	P3, R6	P3, R7	P3, R8	P3, R9	P3, R10	P3, R1	P3, R2
P4, R4	P4, R5	P4, R6	P4, R7	P4, R8	P4, R9	P4, R10	P4, R1	P4, R2	P4, R3
P5, R5	P5, R6	P5, R7	P5, R8	P5, R9	P5, R10	P5, R1	P5, R2	P5, R3	P5, R4
P6, R6	P6, R7	P6, R8	P6, R9	P6, R10	P6, R1	P6, R2	P6, R3	P6, R4	P6, R5
P7, R7	P7, R8	P7,R9	P7, R10	P7, R1	P7, R2	P7, R3	P7, R4	P7, R5	P7, R8
P8, R8	P8, R9	P8, R10	P8, R1	P8, R2	P8, R3	P8, R4	P8, R5	P8, R8	P8, R9
P9, R9	P9, R10	P9, R1	P9, R2	P9, R3	P9, R4	P9, R5	P9, R8	P9, R9	P9, R10
P10, R10	P10, R1	P10, R2	P10, R3	P10, R4	P10, R5	P10, R8	P10, R9	P10, R10	P10, R1

Each column corresponds to the matchings for a particular MATCH. For example, in MATCH 3; P1 is matched with R3, P2 is matched with R4, etc. One way to think about this procedure is that it is just like speed-dating; its as if the PROPOSERS are sitting in 10 chairs on one side of a table and the RESPONDERS in 10 chairs on the other side and then after each MATCH the RESPONDERS all move one chair to their left (and the RESPONDER in the left-most chair goes back to the other end of the table).

This matching procedure implies that if you are currently matched with person A, not only will you not be matched with them again in the future, but also if person A is later matched with person B, who is later matched with person C, then you will already have been matched with person C. This means that your choices when matched with person A cannot affect the choices of any person you will later be matched with. You can therefore treat each MATCH as independent, as though it is the only MATCH in which you will participate.

The following is an example of what this means: P1 and R1 are matched In Match 1. Then R1 is matched with P10 in Match 2. Then P10 is matched with R2 in Match 3. But by that point, R2 has already been matched with P1 (in Match 2). Therefore, P1s previous choices cant affect R2s choices when they are matched together. You can check that this is not only true in this example, but always.

Experimental Instructions for 3:1 UG

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. Make sure you pay close attention to the instructions because the choices you make will influence the amount of money you will take home with you today. Please ask questions if any instructions are unclear.

The Experiment

The experiment will consist of 10 identical MATCHES that each proceed as follows. First, you will be matched with one other participant from the room, and one of you will be in the role of PROPOSER while the other will be in the role of RESPONDER. There are 100 CHIPS that may be split between you and the other participant in the following way.

The PROPOSER makes an OFFER between 0 and 100 CHIPS (offering partial CHIPS up to 2 decimal places is allowed) which indicates how many of the 100 CHIPS the RESPONDER will receive if the OFFER is ACCEPTED. Simultaneously, the RESPONDER choose a LOWEST ACCEPTABLE OFFER between 0 and 100 (again, up to 2 decimal places is allowed) which indicates the fewest number of CHIPS they are willing to ACCEPT. If the OFFER is greater than or equal to the LOWEST ACCEPTABLE OFFER, then the OFFER is considered ACCEPTED in which case the RESPONDER receives the number of CHIPS OFFERED while the PROPOSER receives the remaining CHIPS. Alternatively, if the OFFER is less than the LOWEST ACCEPTABLE OFFER, then the OFFER is considered REJECTED and both the PROPOSER and RESPONDER get 0 CHIPS.

The CHIPS are worth 75 cents each for PROPOSERS and 25 cents each for RESPONDERS. At the end of each MATCH, we will show each pair the PROPOSERs OFFER and reveal whether the OFFER was ACCEPTED or REJECTED. We will also show the earnings of both participants. Although you will participate in 10 MATCHES, we will only pay you for your earnings in one MATCH which will be randomly determined by the computer at the end of the experiment. In addition to any earnings from CHIPS you will earn 12 dollars for participating. Since it is impossible to earn partial cents, the computer will run a lottery to determine if we round up or down to the nearest cent according to the fraction of a cent which you earned. For example, if you are a RESPONDER who earns 22.67 CHIPS then you should earn \$17.6675 ($12 + 0.25 \times 22.67 = 17.6675$) but since we cant give you 0.75 cents you will get \$17.67 with probability 0.75 and \$17.66 with probability 0.25.

Details of matching procedure

At the beginning of the experiment you will be assigned the role of PROPOSER or

RESPONDER and you will remain in that role for all 10 MATCHES. Since there are 20 participants this means there will be 10 participants assigned to each role. There are 10 MATCHES and 10 participants in the opposite role to your own. You will be matched with each of those 10 participants in exactly one MATCH. The exact procedure for how you will be matched is described as follows. Lets name the PROPOSERS P1, P2, , and P10 and the RESPONDERS R1, R2, ., and R10 (these names are only for the purpose of explaining the procedure and you wont actually be given one of these names in the experiment). The following table shows who will be matched with who in each MATCH.

Match 1	Match 2	Match 3	Match 4	Match 5	Match 6	Match 7	Match 8	Match 9	Match 10
P1, R1	P1, R2	P1, R3	P1, R4	P1, R5	P1, R6	P1, R7	P1, R8	P1, R9	P1, R10
P2, R2	P2, R3	P2, R4	P2, R5	P2, R6	P2, R7	P2, R8	P2, R9	P2, R10	P2, R1
P3, R3	P3, R4	P3, R5	P3, R6	P3, R7	P3, R8	P3, R9	P3, R10	P3, R1	P3, R2
P4, R4	P4, R5	P4, R6	P4, R7	P4, R8	P4, R9	P4, R10	P4, R1	P4, R2	P4, R3
P5, R5	P5, R6	P5, R7	P5, R8	P5, R9	P5, R10	P5, R1	P5, R2	P5, R3	P5, R4
P6, R6	P6, R7	P6, R8	P6, R9	P6, R10	P6, R1	P6, R2	P6, R3	P6, R4	P6, R5
P7, R7	P7, R8	P7,R9	P7, R10	P7, R1	P7, R2	P7, R3	P7, R4	P7, R5	P7, R8
P8, R8	P8, R9	P8, R10	P8, R1	P8, R2	P8, R3	P8, R4	P8, R5	P8, R8	P8, R9
P9, R9	P9, R10	P9, R1	P9, R2	P9, R3	P9, R4	P9, R5	P9, R8	P9, R9	P9, R10
P10, R10	P10, R1	P10, R2	P10, R3	P10, R4	P10, R5	P10, R8	P10, R9	P10, R10	P10, R1

Each column corresponds to the matchings for a particular MATCH. For example, in MATCH 3; P1 is matched with R3, P2 is matched with R4, etc. One way to think about this procedure is that it is just like speed-dating; its as if the PROPOSERS are sitting in 10 chairs on one side of a table and the RESPONDERS in 10 chairs on the other side and then after each MATCH the RESPONDERS all move one chair to their left (and the RESPONDER in the left-most chair goes back to the other end of the table).

This matching procedure implies that if you are currently matched with person A, not only will you not be matched with them again in the future, but also if person A is later matched with person B, who is later matched with person C, then you will already have been matched with person C. This means that your choices when matched with person A cannot affect the choices of any person you will later be matched with. You can therefore treat each MATCH as independent, as though it is the only MATCH in which you will participate.

The following is an example of what this means: P1 and R1 are matched In Match 1. Then R1 is matched with P10 in Match 2. Then P10 is matched with R2 in Match 3. But by that point, R2 has already been matched with P1 (in Match 2). Therefore, P1s previous choices cant affect R2s choices when they are matched together. You can check that this is not only true in this example, but always.