Negotiation as the Art of the Deal*

Matthew O. Jackson, Hugo F. Sonnenschein, Yiqing Xing, Christis G. Tombazos and Omar Al-Ubaydli†

Draft: April 2018

Abstract

Negotiation is a ubiquitous and consequential form of economic interaction. It is deal-making in the absence of a designer. We propose a theory of negotiation in which deals have many aspects. This leads to new results showing that efficient trade is possible even with substantial asymmetric information, which we show via both theory and experiments. In a robust class of settings of asymmetric information, the benefits of identifying areas of mutual gain redirect agents away from posturing and manipulating their share of the pie towards growing the pie. We show that equilibria are efficient, with significant implications for applications.

Keywords: Negotiation, Bargaining, Exchange, Trade, Multiple Items, Linking, Efficiency, Experiment

JEL Classification Numbers: A13, C72, C78, D47, D82, C92

*The theory part of the paper draws on key results from the working paper "A Theory of Negotiation" by Jackson, Sonnenschein, and Xing (2015). On a parallel track, Sonnenschein and Tombazos began investigating that theory experimentally and were subsequently joined by Al-Ubaydli. The realization that a compelling description of negotiation would benefit from the manner in which the projects informed each other, led to increased communication between the teams and eventually to this unified effort. For example, the free-form experiments were a response to our desire to faithfully mirror the possibilities in actual negotiations and our theory. The order in which the authors are listed reflects some of the above history as opposed to their individual contributions.

†Jackson: Stanford University, CIFAR, Santa Fe Institute; Sonnenschein: University of Chicago; Xing: Johns Hopkins University; Tombazos: Monash University; Al-Ubaydli: Bahrain Center for Strategic, Inter-national, and Energy Studies (Derasat), George Mason University, and King Fahd University of Petroleum and Minerals. Financial support for the experiments from Monash University is gratefully acknowledged. We thank seminar participants at Chicago, Keio, Monash, Peking, Penn State, Princeton, Stanford, Tokyo, and Yale, as well as the Decentralization conference, for helpful comments and suggestions.

1 Introduction

Negotiation is a ubiquitous and economically important form of social interaction. Two parties come together to reach an agreement: a "deal". Such deals are generally complex. For instance, an agreement between a firm and a union would typically include: salaries, responsibilities, profit-sharing, work-hours, a pension, a medical plan, vacation-hours, safety-rules, seniority, promotion-schedules, etc.

Despite the prevalence and complexity of such negotiation problems, the only models that economists have that offer insights into negotiations are models of 'bargaining': splitting a pie. Bargaining models are fundamental to predicting how people split a known surplus (e.g., Nash 1953 and Rubinstein 1982), as well as to understanding why people may fail to reach an efficient agreement when bargaining over a price at which to potentially trade some object over which they have privately known values (e.g., Myerson and Satterthwaite 1983). However, bargaining models do not offer much insight into the multiple dimensional negotiation problems that management consultants and practitioners spend much of their time on. For instance, Fisher and Uris celebrated 1983 book on negotiating, "Getting to Yes," which is extensively used by practitioners, is much more about how to find and craft the right deals, than about how to agree on a price. It includes statements like: "Realize that each side has multiple interests," "Broaden the options on the table rather than look for a single answer," and "Search for mutual gain." These statements are not to be dismissed, but instead suggest that we need a theory that explains why these issues are central to negotiations.

Moreover, a gap between observed behaviors and what we might expect from bargaining models also calls out for a model of negotiations. In particular, extrapolating from the results of Myerson and Satterthwaite (1983), one would be tempted to predict that negotiations will be inefficient. Curiously, however, there is relatively little empirical evidence in support of such inefficiency. In fact, between 1948 and 2005 "idleness due to strikes in the United States never exceeded one half of one percent of total working days in any year" (Kennan 2005). Since 1990 average lost time has been about twenty minutes per year per worker in the U.S.; and even in a more strike prone country, like Spain, the number is less than 1/3 of a day per worker per year (again, Kennan 2005).

In this paper, we present a model of negotiation that includes substantial asymmetric information, and yet predicts efficient outcomes. Our model sheds light on this apparent paradox by showing how negotiating over multiple dimensions provides strong incentives for people to reach efficient outcomes. It makes clear why negotiations over multiple dimensions can be efficient while bargaining over a single object for trade can lead to inefficiency. In doing this, the theory offers insights behind the above quotes from 'Getting to Yes'. We also present experiments that illustrate the contrast between bargaining over one dimension and negotiating over several.

There are three key aspects to our theory.

The first is that negotiations involve *multiple dimensions*, rather than just one. There is a deal to be crafted rather than just one good to be traded.

The second is that asymmetric information between the two negotiating parties is primarily about *which deals maximize their gains from trade*, rather than whether there exist gains from trade. Without knowing that there are gains from trade, added dimensions would just complicate the inefficiency. As dimensions expand, so do opportunities to find mutual gains from trade, and agents become increasingly sure that there are gains from trade, and how large they must be.[1] The incentives of the agents turn to finding those gains from trade, rather than posturing to take advantage of uncertainty.

The third aspect of our theory regards *how* people negotiate. One might conjecture that the simple expansion of opportunities and known gains from trade make it obvious that outcomes will be efficient. An important preliminary result is that this intuition is wrong. We provide a simple example that strongly makes this point. There are just two goods and it is common knowledge that both involve positive gains from trade and exactly what the total gains from trade are. The only uncertainty is about which good has the higher gains from trade. This example illustrates very starkly the contrast between 'bargaining' and 'negotiating'. Here we show that if agents *bargain* over the goods, so that they offer and counter-offer prices for each of the two goods separately, then despite the common knowledge of the gains from trade and that both goods provide positive gains, *all* equilibria are inefficient. In contrast, if agents negotiate, so that they can propose deals involving multiple goods, then all equilibria are efficient. The distinction between negotiating and bargaining is very consequential.

This previews our more general theoretical results. When agents understand the richness of the set of deals that they can propose, then their incentives become aligned and instead of concealing information and posturing, it becomes in their interest to share information and find the right deal. To understand why negotiation works, it is useful to distinguish between two varieties of asymmetric information. The first is knowledge of the overall possible gains from trade. The second is knowledge of the particular deals that realize these gains. Even with knowledge of the overall gains from trade, asymmetric information can mean that neither agent alone knows which deals are efficient. They must negotiate to find those deals. An essential insight behind our main results is that the knowledge that there exist efficient deals crowds out inefficient deals. Offering a deal that is inefficient is dominated by offering a set of deals that includes all the ones that could be efficient - even if one is not sure which

[1] Whether approximate gains from trade are known in labor-management negotiations is an empirical question. In Section 4 we comment on a range of situations in which the gains from trade are known.

one it is - as that leads to a higher total surplus to be split and the proposer can ask for a bit more than she was asking for with the inefficient deal in every one of the potentially efficient deals offered. Our general result applies to a broad class of negotiation problems and ways in which agents communicate, but this intuition is at its heart.

In addition to our theoretical results, we present a set of experiments that provide results that illustrate this contrast between bargaining and negotiations. Not only do we find predicted inefficiency in bargaining and efficiency in negotiations, but the ways in which agents negotiate and find their mutual gains of trade line up remarkably well with our theory.

Our paper proceeds as follows. In order to fix some main ideas, Section 2 presents an example of a negotiation problem over multiple goods that, as described above, illustrates the multi-dimensional negotiation setting and shows how it is the combination of knowledge of gains of trade, and understanding of the ability to offer a rich set of deals that leads to efficiency; and that knowledge of the surplus alone with simple dimension-by-dimension bargaining leads to inefficiency. This sets up our experiments and theory.

In Section 3 we present results from experiments involving 340 subjects. The basis for the experiments is exactly the examples from Section 2. Some subjects were put into treatments in which they bargained over individual items, based on the examples; while other subjects were put into treatments in which they negotiated over four objects and were able to offer multi-aspect deals. The experimental results are very much in line with our analysis of the example. The bargaining on one item is inefficient, exhibiting delay, posturing, and failure to trade. The negotiations over four items together leads to highly efficient outcomes. In particular, in the bargaining treatments 67 percent of the surplus is obtained - with the inefficiency coming from both substantial delay and outright failure to trade when there are positive gains from trade (only 87 percent of pairs with positive surplus managed to eventually trade).[2] In contrast, in the negotiation treatments 86 percent of the surplus is reached, and the right trades are eventually made in 98 percent of the matches. Given our substantial sample size, both of these differences are significant. The results quoted above are for an alternating-offers version of the treatments. We also did free-form versions of the treatments in which agents could chat and fill in details of offers until they reached mutual agreement. This variation increased the efficiency of both treatments: in the bargaining (one-item) treatments in free-form the efficiency is 75 percent (still only 87 percent of pairs with positive surplus managed to eventually trade), while in the free-form negotiation (four-item) treatments 92 percent of the surplus was reached (with 99.6 percent of the right trades eventually being made). The chat was cheap talk and did not affect the actual offers

[2]This is in fact even worse than the a theoretical bound on the efficiency that can be reached (by the second-best optimal mechanism) in this example with bargaining on each individual item.

available in the free-form treatments, and so the theory still tells us that bargaining should be inefficient and negotiations should be efficient. Again, this is consistent with our theory and analysis of the example, and the difference is significant.

Thus, the experiments show a highly significant increase in efficiency with negotiation as opposed to bargaining with identical expected surplus and item-by-item uncertainty.

In Section 4, we present our general theory of multi-aspect negotiation, giving efficiency and uniqueness results. There is asymmetric information regarding how the agents view alternative multi-dimensional deals, but the maximum surplus is known. We examine various forms of alternating offer negotiations in which agents discount the future. We show that when agents have rich enough ways of communicating so that they can offer menus of full deals or express their types and demand value, then all weak perfect Bayesian equilibria are fully efficient and lead to a deal in the first period. In fact, the share of the surplus that each agent receives is determined as in Rubinstein's (1982) original full-information bargaining theory. However, now the agents need to negotiate in order to discover the deals that generate maximum surplus. The offers made in a negotiation have a particular form. Directly or indirectly they present a set of alternative deals to a counter-party and give her the opportunity to select one of them or decline. When the set of offers that agents select from is "rich", then the agents are able to offer a menu of deals that includes efficient deals as possible outcomes. They are not required to do so; rather, it is in their self-interest. The efficient ones dominate and drive out the inefficient ones, by being able to offer better terms to both agents. This is the invisible hand that enables efficiency.

We include an explanation of a duality between two different ways of negotiating: one by offering menus of deals to the partner to choose from; and the other by revealing preference information to the partner and then allowing them to craft a deal that delivers a requested payoff. Interestingly, we show evidence of both of these being used in the free-form negotiation treatments of the experiment.

A companion paper (Jackson, Sonnenschein, Xing 2015) provides more detailed theoretical results, including analysis of negotiation in which the overall surplus is unknown, as well as some law of large number results, and additional technical results on definitions of equilibrium for cases in which the surplus is unknown.

2 Two Examples of Negotiation Problems

In order to fix some of the main ideas, we begin with a pair of negotiation problems between a buyer and a seller.

PROBLEM 1 *There is a seller who has cost \$0 or \$160, equally likely. The buyer has value \$40 or \$200, equally likely. The costs and values are private information and are determined* *independently. This is a bargaining problem, the only decision is whether the agents can find a price at which both are willing to trade.*

A "deal" prescribes the transfer of the object from the seller to the buyer and a price paid to the seller. One calculates the profit from trade in the usual manner. Three quarters of the time there exist deals that will benefit both agents, while one quarter of the time there is no mutually beneficial deal.

PROBLEM 2 *The seller begins with four goods, each of which has cost either \$0 or \$40, equally likely, and the buyer has values for these objects that may be either \$10 or \$50, equally likely. They also know that each of the four possible combinations of cost and value, (\$0 \$10); (\$0 \$50); (\$40 \$10) and (\$40 \$50), will manifest in the case of each of the four goods. However, the order in which the four possible pairings occur is random. To be more precise, each of the four possible pairings of costs and values occurs exactly once, and the twenty-four 4-tuples of costs and values with this property are equally likely.*

A "deal" prescribes the transfer of some of the goods from the seller to the buyer and a price paid to the seller for those goods.

Note that the two problems have identical expected potential (maximal) gains from trade of \$70. In fact, each of the four objects in Problem 2 is exactly a one-fourth rescaling of Problem 1. Both Problems involve substantial asymmetric information. In the first problem, which is one-dimensional, agents are even unsure if there are gains from trade. Half of the time, an agent knows that there is a mutually-beneficial deal possible (e.g., a \$0 cost seller is certain that there are positive gains from trade), and the other half of the time an agent is uncertain as to whether such a deal even exists (e.g., a \$160 cost seller anticipates a 50 percent chance that there are gains from trade). Inefficiency is unavoidable under the requirements of (interim) incentive compatibility and individual rationality (the price lies between the two valuations), which is a finite adaption of the results of Myerson and Satterthwaite (1983).[3]

The second problem also involves substantial asymmetric information; however, the agents know that there are gains from trade. It is, in fact, common knowledge that there exists a deal with a surplus of \$70: the seller transfers the three goods to the buyer for which her cost is less than the buyers value. The uncertainty is about which deals generate that surplus.

There is no theory to guide us about this second problem, and so that is developed below. In order to preview our main results, and to set up the experiments, we present some results about this example.

[3]See Segal and Whinston (2016) for a discussion of extensions of Myerson and Satterthwaite's results.

2.1 Inefficiency with Bargaining on Problem 2

We begin by showing that the challenges of reaching efficient outcomes with multiple aspects requires not only having some knowledge that there are gains from trade, but also in how agents negotiate. To do this, we first show that in Problem 2, if agents negotiate by offering prices item-by-item, then all equilibria are necessarily inefficient.

In particular, consider the following alternating-offers bargaining protocol. One of the agents quotes separate prices for each of the four items. The other agent can accept any of the offers. If some items remain untraded, then the roles reverse and the other agent offers prices on those items. If some of those prices are accepted then those items are traded. Agents can continue to alternate in offering on the remaining items, but consume current gains from trade when they are realized. The outcome to this bargaining protocol is inefficient *in all equilibria* for some discount factors ($\delta < 1$). We also prove that all equilibria satisfying a refinement (that rules out fully incorrect beliefs off the equilibrium path) are inefficient regardless of the discount factor. We are not sure whether there exist efficient equilibria for some high discount factors, but know that they must have extreme beliefs (that completely rule out the true state) if they do exist. This is proven in Appendix 5.

To understand why this is the case, let us consider a simplified version of the Problem 2 in which things become quite transparent. Without loss of generality, presume that the seller makes the first offer.

Let us focus on the two items that the seller knows should trade - the ones for which she has value \$0. So, the seller is known to have value \$0 on the two items, while the buyer has value \$10 for one item and \$50 for the other item, with equal probability for which item is preferred by the buyer, and that realization is known only to the buyer. This problem has a known surplus of \$60, which is common knowledge along with the fact that both items should trade immediately. Showing that inefficiency results here where both items should trade, makes it easy to extend the argument to the situation in which the agents are not sure which goods should trade.

The basic logic behind inefficiency is as follows (the full proof is given in Appendix 5). In order to get an efficient outcome, a price of no more than \$10 should be offered on each item in the first period. In doing this, the seller gives much of the surplus to the buyer. By instead offering higher prices and screening to get more on the \$50 item, the seller improves her expected utility. One can use upper and lower bounds on the continuation values to show that the buyer will accept a price of more than \$10 on the \$50 item. Thus, the seller is better off giving up some efficiency on trade - delaying trade on the \$10 item - to get more surplus. With more items, this loss of efficiency is substantial.

This is suggestive of a Myerson and Satterthwaite-type inefficiency result in the context of negotiation, but it is important to emphasize the different reason for the result in our context.

Here the gains from trade are common knowledge and it is even common knowledge that both goods should trade - the only uncertainty is about which good is responsible for which part of the total surplus. When bargaining on each item individually, without knowing which one the buyer values more, the seller is willing to delay trade of the low-surplus item to try to screen the buyer and grab more of the surplus on the higher item. Our proof of inefficiency shows that the full gains from trade can only be realized if the agents have the ability to communicate and negotiate in an appropriate manner.

2.2 Efficiency with Negotiations on Problem 2

Now consider the same example as just analyzed. Suppose instead that the offers that can be made include offering to trade both items at a total price. The screening incentive disappears as the seller can now ask for her fair share of the total surplus, which can be immediately realized. In this case, the agents simply agree to trade both items in the first period. As we show (as a corollary to our main results), efficient trade occurs in all equilibria in the first period and the first person to make offers gets a total utility of $\$60/(1+\delta)$ and the other gets $\delta\$60/(1+\delta)$.

This example is so basic that it over-simplifies things - but it still makes very clear that even in the presence of a known surplus, the reason that negotiations can be efficient requires the richness of the offer space, which helps align the negotiators' incentives. It ensures that the agents can offer deals that realize the full surplus and and any split of that surplus.

Once we move back to the four goods example, it is less clear that the agents should always reach efficiency. There they also have to decide which items to trade, not just on a transfer. In that case, the key is that agents have rich enough offers to suggest a *menu* of deals to the other agent, in which they can include the various possible efficient outcomes and then let the other agent choose the most efficient among them.

To see how this works, consider Problem 2 in which the seller has realized values of $0, 0, 40, 40$ for the four items in order, while the buyer has realized values $10, 50, 50, 10$. Although the seller does not know which goods should trade, she knows what the possibilities are. The seller offers two deals: deal 1 is to trade items 1, 2, 3 at a combined price of $\$60/(1+\delta)+40$, and deal 2 is to trade items 1, 2, 4 at a combined price of $\$60/(1+\delta)+40$. She could include more deals, but they would be inefficient and so would generate less surplus than the better of these two. Given the realization of the buyer's type, he accepts deal 1 in the first period and the game concludes. There is no way that the agents can earn a higher combined utility, and as we show these terms of trade are the outcome in all equilibria of the game.

Observe how the richness of the offers is used in the preceding argument. The seller is able to place two deals on the table, each of which involves the transfer of more than one good. The seller uses her knowledge to identify which deals are potentially efficient, and then the seller allows the buyer to use his knowledge to select the efficient deal. In general, the richer set of offers means that the seller knows that she can offer deals that realize the full surplus and then ask for her split of that surplus. Any offers that get less than the total surplus can only lead to worse utilities for one or both agents. Any offer that is clearly inefficient is dominated by one (or a set that is sure to include one) that is efficient, which can offer better total utilities for both agents. This is the force that pushes agents to efficient negotiations. The proof is longer than this, of course. It is also more involved, since it covers negotiation games that are not direct offerings of menus of deals, but might have other forms of communication - for instance declaring values on different dimensions - or being free-form. We show that the knowledge of the total surplus, and the ability to communicate in ways that allow agents to find that surplus, aligns incentives.

These arguments suggest what lies ahead. In Section 3, we present the result of some laboratory experiments in which agents face Problems 1 and 2. This is followed by the theoretical analysis in Section 4, which concerns the efficiency of negotiation when agents have a good idea of the total gains from exchange but have substantial uncertainty as to which combination of dimensions realizes those gains.

3 Experiments

In this Section we present four experimental treatments based on the two Problems above. In order to understand the roles of both the ability of agents to negotiate over several aspects of a deal simultaneously and "known surplus", our treatments contrast the cases of Problem 1 and Problem 2

We are particularly interested in the case where there is no "market design", so that the agents come together without a third party, who has knowledge of the statistical structure of the problem, and who can force the agents to play a particular game. The experiments provide a controlled exploration of the extent to which unmediated negotiation results in efficient exchange. To make this point most forcefully, we also include free-form versions of the treatments in which negotiations are quite open.

3.1 Experiment Design

The experiment is an across-subjects design, and contains comparisons between the two negotiation Problems 1 and 2 from Section 2.

The first pair of treatments that compare Problems #1 and #2 is one in which there are alternating offers, and we call these the "structured" treatments. The second pair of treatments that compare Problems #1 and #2 is one in which who makes offers at which point in time is open, and agents can also freely chat, and we call these "free-form" treatments. See Table 1 for a summary.

Table 1: Design in Each Treatment

	Problem	Format	Number of Subjects
Treatment 1	1 (1 good)	Structured	94
Treatment 2	2 (4 goods)	Structured	96
Treatment 3	1 (1 good)	Free-form	76
Treatment 4	2 (4 goods)	Free-form	74

In the structured treatments agents take turns in making offers, and at most one offer is made in each period. Discounting applies after each period. In Treatment 1 on Problem 1, the agents make offers via a standard alternating-offer bargaining game. In Treatment 2 on Problem 2 we extend alternating offers to allow the agents to offer a menu of deals. In particular, an offer consists of a list of deals: which items are to be traded and a total price to be transferred from buyer to seller. For instance, the seller might choose to list say three deals: trade items 1, 3, and 4, for a price of $50, trade items 2, 3, and 4, for a price of $44, trade items 1 and 4, for a price of $65. In the case of Treatment 2, the agent making offers could include as many deals as the agent wished. The buyer could then choose to accept one of these deals or to reject them all. If they were all rejected, then the period ended and the buyer made the offers. Typically, two deals were offered - both potentially efficient ones. For a precise statement of the rules and screenshots, see the Experimental Supplement.

Discounting occurred after each period. The stakes shrunk by 10 percent after each alternation. There was a cap on 8 periods, and if they did not agree by then the game ended with no trade.

The free-form treatments were similar, except that there was no alternation. At each instant, each subject could propose an offer that would show up on the counter-party's screen, or accept an offer that was currently on the screen from the counter-party; each subject can also make any edit(s) to the current offer as long as it is not yet accepted. In addition, the subjects can chat by typing text in a chat box. Negotiation ended when an offer was accepted. Also, the same as above discounting applied after each "period". The initial period lasted 40 seconds, and each of the remaining lasted 20 seconds. Again, for precise statements of free-form negotiations, see the Experimental Supplement.

3.2 Administrative Details

We conducted the experiment at the Monash University Laboratory for Experimental Economics (MonLEE), using z-Tree (Fischbacher 2007). Subjects were predominantly undergraduate students from Monash University, recruited from a database maintained by MonLEE. No subject participated in more than one session.

We conducted 18 sessions employing a total of 340 subjects. A detailed summary of the sessions is in the Experimental Supplement. Each session lasted approximately 90 minutes.

Subjects were paid an attendance bonus of Australian $10,[4] in addition to their positive or negative earnings from the experiment. Subjects made $37.06 on average, ranging from $3 to $190, including the attendance bonus.

An important fact is that each session was devoted to one treatment. Thus all subjects in a given session did just one treatment, and were randomly rematched within their session to play that same treatment with a series of different counter-parties. In any given session, subjects experienced 10 rounds (matchings) of negotiating in pairs. The first four matchings were "practice" and the last six rounds were "real". Earnings were based on an ex post random selection of one of the last six rounds. Our approach allows us to work with a very conservative statistical treatment of comparing behavior across treatments, with standard errors clustered at the session level.

3.3 Experimental Results

We begin by comparing the efficiency of outcomes, comparing Treatment 1 with Treatment 2, and 3 with 4. As we have suggested, the fact that the overall gains are known in Problem 2 with 4 goods, and that agents can negotiate over aspects simultaneously, allow the subjects to find deals that will lead to mutual gain and greater efficiency than in Problem 1. This is the case. As presented in Table 2, regardless of whether we look at the structured or free-form formats, the negotiation (four-good) treatments lead to significantly more efficient outcomes than the bargaining (one-good) treatments. The results are significant at above the 98 percent level.

The p-values are from a most-conservative statistical analysis in which we consider each session as a single observation. We take this extreme caution since rematching of agents could lead to dependent outcomes across pairs of individuals. This gives us, for instance, only four observations of Treatment 1. The results are still highly significant since the variance across sessions of the same treatment was very low (see the Experimental Supplement). The results are similar if we do regressions clustered at the session level or if we do a Mann-Whitney

[4]$3 out of the $10 is guaranteed, so that if a subject made a loss in excess of $7, that subject would walk away with a show-up fee of $3.

test instead of a t-test (again, see the Experimental Supplement).

Table 2: Percentage Loss of Surplus

	1 good	4 goods	p-value
Structured	32.8%	14.3%	.013
Free-form	25.2%	7.9%	.014
p-value	.232	.043	

It is also interesting to observe that the ability of agents to chat and use a free-form format increases efficiency regardless of whether we are in bargaining or negotiation. This is consistent with, for instance, Charness (2000), Ellingsen and Johannesson (2004ab), Charness and Dufwenberg (2006), and Feltovich and Swierzbinski (2011). It does not change the fact that there is a large difference between bargaining and negotiation, as with one good the forces that push towards efficiency are still not present even with a free-format and open chat.

There are different ways to measure efficiency. Table 2 presents a definition that tracks the total surplus that each pair of subjects realized as a percentage of the total surplus available (so pairs are effectively weighted by how much surplus they could have generated, and pairs who had no possible surplus are ignored, as negative surplus was not observed in the experiment). The results are even more pronounced when we just track the percent of surplus that each pair loses and then average those percentages across all pairs unweighted by their size of surplus. Those numbers are more pronounced since the pairs who have the most delay in the one good case are those with the least surplus (again, see the Experimental Supplement).

The inefficiencies in the 1 good case come from both delay and failure to ever trade even though there are gains from trade. In contrast, for the 4 good negotiation case, there is almost always eventual trade - especially in the free-form treatment in which almost all pairs traded - inefficiency only comes from delay when agents try to find the right deal. This is illustrated in Table 3. The differences across treatments is again significant, even with a most conservative test that treats each session as a single observation.[5]

Table 4 provides the breakdown of how the trades vary by period in the free-form treatments. We see earlier and higher trading in the four good negotiations compared to the

[5] Again, we are very conservative and treat each session as an observation, and so the entry is the average of session averages and the p-values are from a t-test across these averages. The p-values are lower if we work with other less conservative tests - such as regressions with clustered standard errors. The high accuracy in spite of having few observations in each cell again comes from the very low variance across sessions of the same treatment.

Table 3: Fraction of Efficient Pairs Eventually Trading

	1 good	4 goods	p-value
Structured	87.1%	97.6%	.041
Free-form	86.5%	99.6%	.022

one-good bargaining.[6]

Table 4: Percent of Pairs of with Positive Surplus Trading by Rounds

	In period #								
	1	2	3	4	5	6	7	8	Not Trading
1 good Free-Form	24.6	19.8	12.0	10.2	7.2	3.6	2.4	7.2	13.2%
4 goods Free-Form	81.1	11.3	2.3	1.8	1.8	1.4	0.0	0.0	0.5%

As discussed above, and proven below, with the four-good negotiation setting, agents no longer have incentives to misrepresent their preferences, while in the one-good bargaining setting they do. The experiments shed some light on this question. In the free-form treatment, out of the subjects who make any claim about their values in the chat, we can track whether those claims are true. There is some subjectivity in categorizing when people are declaring a value, but most cases are fairly clear and we describe the precise rules we followed in the Experimental Supplement.[7] The results are presented in Table 5 . Misrepresentations routinely occur in the 1 good case throughout the game, but rarely in the four good case (only three times and in the first period).

Table 5: Fraction of declarations that are untrue in free-form treatments

	In period #								
	1	2	3	4	5	6	7	8	Average
1 good	92/350	21/77	11/41	11/26	5/21	1/16	1/6	1/12	26.0%
4 goods	3/148	0/2	0/1	n.a.	n.a.	n.a.	n.a.	n.a.	2.0%

[6]The percent not trading is of all pairs, and so that is why, for instance, the 13.2% number for the 1 good free-form treatment is 0.3% different from the 86.5% eventually trading from Table 3. That previous table considers each session as an observation, and then averages across sessions.

[7]For example, we categorize a declaration to be any instance in which a person communicates one of the possible numbers that they might have on an item or items, and a misrepresentation to be when they communicate a number that differs from their actual value. In a number of instances declarations are indirect. As an example, consider the following exchange, from session 11, round 5, group 2, Seller: "What do you have?", Buyer: "Whats not 200?", Seller: "40?", Buyer: "Yeah". The buyer never says 40 explicitly, but the message seems clear. Further details appear in the Experimental Supplement.

To put the 26 percent misrepresentations in the bargaining treatment in context, note that half of the time, when a buyer has a low value or a seller has a high value then there are no real gains to misrepresentation. Thus, misrepresentations should only occur for half of the types, and equilibrium is in mixed strategies. So, 26 percent represents a rate of more than one half in the cases where the agents should be mixing.

The fact that there was only three misrepresentations out of more than one hundred declarations in the negotiations treatment suggests that the force that aligns subjects incentives and makes the issue about finding the right deal rather than posturing is not subtle: there is no heterogeneity here and so the force is strong and not one that requires high levels of sophistication among the subjects. In the theoretical sections we provide reasons to believe that the issue of finding the right deal is similarly transparent, even when negotiations are of rather general complexity.

Examples of how Subjects Bargain and Negotiate

We present some typical examples of how things work in the free-form bargaining and negotiation treatments, as they give an idea of how the experiments worked, and also dovetail with the theory. Of course, they are anecdotal as we have chosen only a few of hundreds of pairings; but these provide insight into Table 5 and the theory to follow.

We start with the free-form one-good bargaining treatment.

Here is an example of a seller with value \$0 and buyer with value \$40.[8] We see the attempt of the seller to get a high price, but the seller eventually gives in and 50 percent of the surplus is lost to discounting.

- Seller: (Offers price of \$180)
- Buyer: (Offers price of \$0)
- Seller: "If you have 200, please accept it. Quickly."
- Buyer: "I dont"
- Seller: "Split 20 20. What do you have?"
- Buyer : "40"
- Seller : "Make an offer"
- Buyer : (Makes an offer of \$20)
- Seller: (Makes a counteroffer of \$179)
- Buyer: "Why not make a leap of faith"
- Seller: "Accept mine then"

[8] In some cases, subjects sent messages at or near the same time. The ordering is set to make as clear as possible. Chats are edited for clarity and not all chatter is reported in these transcripts. For example in the fifth line where the buyer reports "40" - this is edited from Buyer: "Whats not 200?", Seller: "40?", Buyer: "Yeah".

- Buyer: “Trust that Im 40. I cant make negative profit”
- Seller : “Why should I? ha ha”
- Buyer : “Guess we profit 0 then. Your call.”
- Seller : “Ok I trust you if you offer 21”
- Buyer: (Offers a price of $21)
- Seller: (Accepts in bargain period 6 and profits are subjected to a 50% discount.)

Here is another example of a seller with value 0 and buyer with value $40. Here, the seller explicitly misrepresents, and in this case the good never trades.

- Buyer: “Hi whats your cost”
- Seller: “Im guessing yours is 40”
- Buyer “Stage 1 lets go. Yeah mine is 40 - hahah good guess”
- Seller: “This time I got 160 - Not joking.”
- Buyer: “Oh no. Lets not do anything then.”
- Seller: (Makes an offer of $180 and remains firm throughout bargaining periods. Good never traded.)

Here is an example of a seller with value 0 and buyer with value $200 who both misrepresent. Eventually the seller gives in and they trade in the last period but lose 70 percent of the surplus:

- Seller: “Im 160, u?”
- Buyer: “40”
- Seller: “Damn”
- Seller: (Offers price $180)
- Buyer: “Lets be honest [...]”
- Seller: “Whats ur offer”
- Buyer: (Offers price of $20)
- Seller: “Ofc”
- Buyer: “Because its 40 for me - if you have 0 - then this is even”
- Seller: “Well its 160 for me - so yeah”
- Buyer: “We are in stage 5 - no time to waste - if you have 0 - go for it”
- Seller: “Go for mine”
- Buyer: “I sure would have - but I cannot - stage 8 - go for mine”
- Seller: (Accepts the price of $20 in period 8 and profits are subjected to a 70% discount)

In contrast, the four-item negotiations the chat and negotiation tends to be very short and to the point - either offering the possibly efficient deals or expressing valuations truthfully (as in Table 5) and then reaching an efficient deal in the first period. Here are typical examples.

In the first, the buyer offers what "he" thinks could be efficient deals and they are quickly accepted.

- Buyer: (Offers the two deals that could be efficient given his information, with prices that split surplus evenly)
- Buyer: "Its half-half. Pick one. Quick"
- Seller: (Accepts the offer that maximises joint surplus. 35 seconds elapse in total)

In the next one, instead of starting with the deals that could be efficient, the agents begin by declaring their private information truthfully and then constructing the efficient deal.

- Seller: " 0 40 0 40 "
- Buyer: "My 10s are in 1 n 4 "
- Buyer: "50/50 " (Offers a deal that maximises joint surplus and splits it evenly)
- Seller: (Accepts. 21 seconds elapsed in total)

Here is another example.

- Buyer: "10s in 2 and 4"
- Seller: "40s in 1 and 4"
- Buyer: (Offers a deal that maximises joint surplus and splits it evenly)
- Seller: (Accepts. 18 seconds elapsed in total)

To summarize, the efficiency of the four-good negotiation is consistent with the predictions of our theory, and contrasts significantly with the inefficiency of the one-good bargaining setting. This holds even in the free-form setting, where there is no real structure on the ways in which subjects can negotiate: effectively there is no "mechanism design". Moreover, the dialogs of the experiments are suggestive of the anatomy of win-win negotiation - subjects quickly exchange information truthfully and move to an efficient deal. We now show that this should be expected as a general feature of negotiations, when agents realize and use powerful yet simple strategies which generalize the behaviors used in the experiments. The key to our argument is the precise definition of these strategies and the richness condition.

4 A Theory of Multi-Aspect Negotiation

We begin by describing the general model.

4.1 The Model of Negotiations

4.1.1 Multiple Aspects and Deals

A multi-aspect negotiation problem consists of:

- two agents, Alice a and Bob b;
- a finite number, n, of aspects with generic index k;
- for each k, a space X_k from which aspect k takes its potential outcomes, where $X = X_1 \times \ldots \times X_n$ denotes the space of vectors of outcomes;
- the space of deals, $X \times I\!R$, with a representative deal $(x, p) = (x_1, \ldots, x_k, \ldots, x_n, p)$, where p is a monetary transfer from Bob to Alice.

This formulation is rather permissive. For example, in a labor-management negotiation, the outcome of an aspect can indicate which of a variety of medical plans are included in a deal. In Problem 2 (section 2) there is a seller and a buyer and a deal is a 5-tuple, where the first four coordinates are either "trade this good" or "not trade this good", and the last coordinate is a price.

4.1.2 Types and Preferences

More generally, preferences and uncertainty are captured via

- valuation or type spaces Θ_i, $i \in \{a, b\}$ and a joint type space $\Theta \subset \Theta_a \times \Theta_b$;
- a utility function $u_i : X \times \Theta_i \to I\!R$, for each agent $i \in \{a, b\}$;[9] and
- a common knowledge probability distribution f over Θ, with f_i denoting the marginal of f on Θ_i.

For simplicity in defining beliefs and equilibria, we presume that Θ is finite.[10]

In the four goods example (Problem 2 in Section 2) the type of each agent is a vector of costs or values for each of the goods. For instance, $\Theta_a = \{0, 40\}^4$, which is a generic type having four dimensions. Alice's payoff from a deal $(x, p) \in \{\text{trade}, \text{no trade}\}^4 \times I\!R$ is then captured by the utility function $u_a(x, \theta_a) = -\sum_{k: x_k = trade} \theta_{ak}$ with net payoff from the deal of

$$p - \sum_{k: x_k = trade} \theta_{ak}.$$

[9] We presume private values for our analysis, but clearly the model can be stated so that u_i depends on the full vector of types.

[10] That restriction can be removed, but with a bunch of extra technical care in defining beliefs in weak Perfect Bayesian equilibrium.

Problem 2 then has a prior distribution, f, that places equal probability on each of the type combinations $(\theta_a, \theta_b) \in \Theta = \{0, 40\}^4 \times \{10, 50\}^4$ for which each possible matchup of valuations comes up on exactly one of the aspects: $f(\theta_a, \theta_b) = 1/24$ if $(\theta_{ak}, \theta_{bk}) \neq (\theta_{ak'}, \theta_{bk'})$ for all $k \neq k'$, and $f(\theta_a, \theta_b) = 0$ otherwise. These are equivalently the $(\theta_a, \theta_b) \in \Theta = \{0, 40\}^4 \times \{10, 50\}^4$ for which $\sum_k (\theta_{bk} - \theta_{ak})^+ = 70$.

Time advances in discrete periods $t = 1, 2, \ldots$. If a deal (x, p) is reached at time t, the agents' net (realized) utilities from the deal as discounted with respect to time $t = 1$ are:[11]

- for Alice: $U_a(x, p, t, \theta_a) = \delta^{t-1} [u_a(x, \theta_a) + p_t]$;
- for Bob: $U_b(x, p, t, \theta_b) = \delta^{t-1} [u_b(x, \theta_b) - p_t]$.

Here, $\delta \in [0, 1)$ is the factor which agents discount the future.

The results extend to heterogeneous discount factors with the usual Rubinstein shares for heterogeneous discount factors replacing the price determination below. For simplicity, we assume that agents do not consume until all of the aspects of a deal are determined. To be more specific, our negotiations do not allow for certain aspects to be fixed and consumed and others left open while the negotiation remains open. Even if it were possible to make different decisions at different times, all the theorems will hold in the situations we consider in this paper.

We call the efficient total payoff the *surplus* of the problem:

$$\max_{x \in X} [u_a(x, \theta_a) + u_b(x, \theta_b)] .$$

The outcome of a negotiation is efficient if and only if the outcome of the deal, x, maximizes this surplus, and the deal takes place at $t = 1$.

A negotiation problem is thus a profile, (n, X, Θ, f, u), as defined above.

4.1.3 Known Surplus

In this paper, we focus on the case in which the total surplus is commonly known to the negotiating parties.

We say that a negotiation problem (n, X, Θ, f, u) has a *known surplus*, $\bar{S} > 0$, if

$$\bar{S} = \max_{x \in X} [u_a(x, \theta_a) + u_b(x, \theta_b)] \;\; \forall (\theta_a, \theta_b) : f(\theta_a, \theta_b) > 0.$$

The range and importance of situations in which this assumption is (approximately) satisfied is largely an empirical matter. However, since researchers have not distinguished

[11] When payoffs are additively separable across dimensions, they can be written as $U_a = \delta^{t-1} [\sum_k u_a(x_k, \theta_{ak}) + p_t]$, and $U_b = \delta^{t-1} [\sum_k u_b(x_k, \theta_{bk}) - p_t]$. Additive separability is a strong assumption, inconsistent with many applications of interest, so we avoid using it.

between uncertainty regarding where the gains from trade are to be found and uncertainty regarding the magnitude of the possible gains, there has been a tendency to place them into the same asymmetric information bucket. The results in this paper, together with findings mentioned in our introduction, suggest that this choice should be rethought. For example, Labor might have a much less good idea than Management of the benefits to Management of being permitted to hire outside contractors for certain work, so there is significant asymmetric information. But the profitability of the firm might be public information and a good proxy for the maximal gains from trade. With this motivation in mind, we present two stylized examples where the gains from trade are quite well known; especially, when compared to the negotiators initial understandings of the surplus from the various deals. Once again, our larger point is to illustrate some highly relevant forms of asymmetric information which can be successfully mediated via the institution of negotiation.

Situation 1: I.I.D.: the Law of Large Numbers

Following Problem 2, consider a seller with valuation \$0 or \$40 (equiprobable, per item) who meets a buyer with valuation \$10 or \$50 (equiprobable, per item). But now, instead of four items, let there be a large number and let all draws be independent, so that approximately one quarter of the time \$0 faces a \$10, \$0 faces a \$50, \$40 faces a \$10, and \$40 faces a \$50. The gains from trade are approximately known, and negotiation leads to approximately efficient exchange. Jackson and Sonnenschein (2007) solve the above problem via mechanism design by forcing the seller to represent a high type (\$50 or \$40) and a low type (\$0 or \$10) half of the time and having the players split the gains on each item evenly when there are gains to be had.

Situation 2: Maximum Value from a Menu

A buyer walks into a gift shop. He has a limited demand: for instance, he has a unit demand because he only wants one gift for his spouse. Suppose the seller's cost is either \$0 or \$40, and the buyer's value is either \$10 or \$50, i.i.d. across items with arbitrary nondegenerate distributions. As long as there are enough items available in the shop, it is very likely that there is at least one item on which the cost is \$0 and the value is \$50, which implies an (approximately) known surplus of \$50. The challenge is to find one such item. The buyer knows the ones that gives him a \$50 value each, but does not know which among them cost the seller \$0.

Our results are robust in the sense that if the surplus is "almost known", then all outcomes are "approximately" the same and efficient. The formulation and proof of such results requires machinery and innovations in modeling that go beyond the spirit of the current investigation, and can be found in the companion paper Jackson, Sonnenschein and Xing (2015). Thus, "known surplus" is to be regarded as a useful idealized case.

4.1.4 Two Examples of Negotiation Games and Their Duality

Before presenting the general definition of negotiation games it is useful to consider two examples. These two negotiation games are extensions of Rubinstein-Stahl bargaining: agents take turns in making offers, until one offer is accepted which ends the negotiation. They also introduce the defining property of a rich set of strategies in a framework that is substantially less complicated than is required for the general definition of negotiation games (which is in the spirit of free-form).

The first example of a negotiation game is one in which the Sender offers a menu of deals, and the Receiver has the option to accept exactly one of these deals or to decline all of them. This takes a step towards the general definition of a negotiation game where the relation between offers and "deals that are offered" is implicit and therefore less transparent. We also note that the structural treatment (#2) in the experimental section is an example of a Menu-of-Deals Negotiation.

EXAMPLE **1 (Menu-of-Deals Negotiations)**

Agents take turns in making offers. In each period $t = 1, 2, \ldots$

- *One of the agents, say Alice, proposes an offer which is a menu of deals*

$$o \equiv \{(x^{(1)}, p^{(1)}), \ldots, (x^{(m)}, p^{(m)})\} \subset X \times I\!R$$

- *The other agent, say Bob, either accepts exactly one of the listed deals, or rejects them all.*
  - ◇ *If Bob accepts one of the deals from o, the game ends and that deal is implemented.*
  - ◇ *If Bob rejects all of the deals, the roles of the agents are reversed, one period of discounting ensues, and the procedure repeats itself.*

...

The above process continues until a deal is agreed upon.[12]

Consider the four-good private exchange example. The seller, who makes the first offer, may propose "trade goods 1, 2, 4 at a price of 30", as well as "trade goods 3, 4 at a price of 40". The buyer may "accept" by choosing either one; or he may "decline", in which case

[12]We do not model agents explicitly leaving the bargaining process, but assign 0 utility to a game that continues forever.

it becomes his turn to (counter-)offer with a list of deals for the seller to choose from. One period of discounting applies before every counter-offer is made.

The second example of a negotiation game is one in which an agent declares preference information and a required total utility payoff, and then the other agent must deliver a deal that provides that utility or reject. Implicitly, agents are identifying a certain set of deals. This negotiation captures other typical manners of negotiation thats observed in the free-form four-good treatment (# 4) in the experimental section.

EXAMPLE **2 (Type-Declaring Negotiations)**

- *One of the agents, say Alice, announces (not necessarily truthfully) her type* $\widehat{\theta}_a$ *and demands a payoff of* $v_a \in I\!R$.
- *The other agent, say Bob, accepts or rejects.*
- *If Bob accepts, he must construct a deal consisting of a set of items traded,* $x \in X$, *and a price that delivers a net payoff of* v_a *based on Alice's announced type: so, the total transfer given to Alice by Bob is equal to* $p = v_a - u_a(x, \widehat{\theta}_a)$. *The game ends.*

  (In the case where the roles are reversed and Bob announced $\widehat{\theta}_b$ *and demanded a payoff* $v_b \in V$, *then Alice picks* x *and the transfer made by Bob is* $p = u_b(x, \widehat{\theta}_b) - v_b$*.)*
- *If Bob rejects, the procedure is repeated with the roles of the agents reversed (and one period of discounting ensues).*

Again consider the four goods exchange problem. The seller, who is assigned the first offer, "might" honestly declare that her private costs are (0,0,40,40) and demand a payoff of $70/(1+\delta)$, to account for discounting. This "might" be accepted by the buyer, who with his private information is able to meet the demand of the seller, and for himself secure the remainder of the surplus. Or, it "might" be rejected. If the buyer "rejects" the initial offer, then it becomes the buyers turn to offer, etc.

Although these two forms of negotiation involve very different offer spaces, they are effectively equivalent and are a sort of dual to each other. One is a direct offer of a set of deals that the proposer is willing to agree to, while the other indirectly identifies that same set and then leaves it up to the other agent to deliver one. Depending on the problem, it can save on communication to communicate via one or the other of these two systems.

This equivalence is made precise in the following Theorem.

4.1.5 A Preliminary Efficiency Result

We work with *(weak) perfect Bayesian equilibrium* as our equilibrium notion. It is formally defined for general negotiation games in Section 4.2.

Theorem 1 *If a negotiation problem (n, X, Θ, f, u) has a known surplus $\bar{S} > 0$, and the negotiation game is either Menu-of-Deals Negotiations or Type-Declaring Negotiations, then in* all *weak perfect Bayesian equilibria:*

- *the agreement is reached in the first period,*
- *the full surplus is realized, and*
- *agents' expected net payoffs are uniquely determined. In particular, they are the Rubinstein shares; i.e., $\frac{\bar{S}}{1+\delta}$ for Alice, and $\frac{\delta\bar{S}}{1+\delta}$ for Bob.*

This theorem is a corollary of Theorem 2, which follows.

We describe the equilibrium path under Theorem 1 for Menu-of-Deals Negotiations, again with our four-good exchange example. Suppose that it is the seller, who makes the first proposal, and has private costs $(0, 0, 40, 40)$. On the equilibrium path, the seller offers two deals: "trade goods 1, 2, and 3 at a price of $40 + \frac{70}{1+\delta}$" and "trade goods 1, 2, and 4 at a price of $40 + \frac{70}{1+\delta}$". The buyer accepts, and chooses the first deal if he values the 3rd good at 50, or the second deal if he values the 4th good at 50. Notice that no matter which of the two deals is chosen by the buyer, the seller always get a payoff of $\frac{70}{1+\delta}$, which is her Rubinstein share; recalling that $\bar{S} = 70$ is the total surplus in the four-good example. Off the equilibrium path, agents accept offers that lead to at least their Rubinstein share (picking a best one if such exists), and always offer the two deals that could potentially be surplus maximizing for their type along with the Rubinstein price.

Under Type-Declaring Negotiation negotiations, the unique initial offer on the equilibrium path is for the seller to truthfully announce her type and demand a payoff of $\frac{70}{1+\delta}$, which the offer illustrated after the statement of values negotiation. Note that this path captures the key patterns of agents' behaviors in the free-form, four goods experimental treatment (#4).

It is far from obvious that these are the only equilibrium outcomes, and that is much of the work of the proof. The general idea is to bound the payoffs using a Shaked and Sutton (1984) argument, and then show that only efficient deals in the first period can reach the mutual bounds, and that these bounds involve the Rubinstein split of the surplus.

The "simplicity" of these negotiation games prevents them from capturing all of the varieties of ways in which agents might actually negotiate. For example, the agents might have both kinds of offers available, together with some additional irrelevant offers.

Thus, in order to build a fully descriptive theory, it is necessary to broaden our notion of negotiation games, and this is done next.

4.2 General Definition of Negotiation Games, and Share-Demanding Negotiations

We now present a general definition of negotiation games and define the richness condition that we have mentioned in an informal manner.

Readers shall keep in mind the Menu-of-Deals Negotiations and the Type-Declaring Negotiations as leading examples of the following definition.

Agents take turns in making offers. In each period t, let $i(t)$ be the proposer and $j(t)$ be the recipient (if the negotiation still continues). We have $(i(t), j(t)) = (a, b)$ at $t = 0, 2, 4, \ldots$, and $(i(t), j(t)) = (b, a)$ at $t = 1, 3, 5, \ldots$.

A negotiation game Γ consists of the following elements: the proposers' action spaces (O_a, O_b), the recipients' action spaces $(R_a(\cdot), R_b(\cdot))$ which may depend on the offer, and the rule to select deals $x(\cdot), p(\cdot)$.

In particular, in each period t such that an agreement has not been reached yet:

- $i(t)$, the current proposer, chooses from a set of possible actions ('offers') $O_{i(t)}$, with a generic offer denoted o^t.
- The other agent, $j(t)$, either "accepts" by choosing one response from the set $R_{j(t)}(o^t)$, with a generic response denoted by r^t, or chooses none of them which means that the agent "declines" the offer (represented by d):
  - ◇ if some $r^t \in R_{j(t)}(o^t)$ is chosen, the deal $x(o^t, r^t), p(o^t, r^t) \in X \times I\!R$ is the outcome;
  - ◇ if the responder "declines", then the agents continue to negotiate with switched roles after one period of discounting ensues.

...

The above continues until an agreement is reached.

This general definition is in the spirit of free-form negotiations in which agents negotiate by freely talking to each other. The only assumption, besides the alternating-offer structure, is that the two agents have some set of language that they can use and commonly understand. We view an agent's action space as the space of language (offers) that are available to that agent. Each offer $o \in O_i$ represents something that agent i communicates to the other agent i, and gives the other agent one or several possible responses, specified by $r \in R_{-i}(o_i)$. The outcome function $x(o, r), p(o, r)$ specifies the the eventual deal that is agreed upon given the way that the agents understand the communication. In the same spirit, the richness, defined below, is a characterization regarding the possibilities for communication available to the agents.

We emphasize the definition of negotiation games is 'universal' in that the above game is not necessarily parameterized by *any* Θ, f, u. This means that our negotiation games can be completely untailored to the particular environment. [13] In this sense, our contrasts with situations in which someone has the power to impose a mechanism or would wish to impose a structure that depends on details of the environment.

4.2.1 Equilibrium

Since sequential equilibria are difficult to define for games with continua of actions, we work with a variant of (weak) perfect Bayesian equilibrium, adapted directly to our setting.

At the beginning of any period $t = 1, 2, \ldots$ agents share a common history of observed actions (offers and reactions) $h^{t-1} \equiv (o^0, d, o^1, d, \ldots, o^{t-1}, d)$ (and additionally they privately know their types); recall d stands for "declines". In addition, after the current proposer $i(t)$ moves, the common history becomes (h^{t-1}, o^t). We denote the set of all possible histories by H, including $h^{-1} \equiv \emptyset$ which is the initial node.

A *belief system* for agent i is a function $\widetilde{f}_i : H \times \Theta_i \to \Delta(\Theta_{-i})$ that maps each history and own type to a distribution over the other agent's type space. In particular, $\widetilde{f}_i(E_{-i} \mid h, \theta_i)$ denotes i's belief over an event (i.e., a collections of the opponent's types) E_{-i}, conditional on a history h and the agent's own type θ_i. To capture the idea that these beliefs apply to nodes in the game, we require that a belief system only place positive probability on θ_{-i} for which $f(\theta_i, \theta_{-i}) > 0$.

Let $H_i \subset H$ be the set of histories at which agent i chooses an action.

At each node $(h, \theta_i) \in H_i \times \Theta_i$, agent i's *strategy*, σ_i, specifies a distribution over the current action space, i.e. $\sigma_i(h, \theta_i) \in \Delta(O_i)$ as the proposer, or $\sigma_i(h, \theta_i) : \Delta(R_i(o^t) \cup \{d\})$ as the recipient (where o^t is the current offer that they have received from the other agent in this period).

Beliefs are *consistent* if for each i and θ_i they correspond to a conditional distribution

[13] Some readers might think of the negotiation games that we consider as "mechanisms", (see, e.g., Jackson (2001,2003), Segal and Whinston (2016), and the references therein for background on mechanism design as well as bargaining inefficiency; and Skrzypacz and Toikka (2014) for the dynamic case). It is important to note that we do not use any distributional information about agents' types. Thus, our work can be thought of in the *broader* spirit of Wilson's (1987) criticisms of mechanisms that depend on agents' (higher-order) beliefs. Satterthwaite, Williams, and Zachariadis (2014) also view such mechanisms as "impractical" as "[the agents'] beliefs are not a datum that is practically available for defining economic institutions" (p.249). Here we go beyond that by not requiring that the negotiation games depend on any payoff or type information at all - it can be a very free-form and open process.

(relative to the common prior f) at almost every h in the support of $\sigma_{-i}, \sigma_i(\theta_i)$.[14]

Let $U_i(\sigma, \widetilde{f}_i, h, \theta_i)$ denote i's expected utility under the strategies σ, conditional on being of type θ_i and history h given the belief system $\widetilde{f}_i$.

A strategy profile σ satisfies sequential rationality (relative to a belief system $\widetilde{f}$) if σ_i maximizes $U_i(\sigma_i, \sigma_{-i}, \widetilde{f}_i, h, \theta_i)$ for each i, θ_i in the support of f, and every $h \in H$ at which i chooses an action.

A *weak perfect Bayesian equilibrium* is a profile $(\sigma_a, \sigma_b, \widetilde{f}_a, \widetilde{f}_b)$ of a strategy profile and a consistent belief system for which the strategy satisfies sequential rationality.

4.2.2 Rich Negotiations

We have observed that under both Menu-of-Deals Negotiations and Type-Declaring Negotiations, efficiency is reached in all equilibria. We now identify a general class of negotiation games of which these are both special cases. The key property is that agents have an offer that they can make which includes the possibility of an efficient deal for every type that the other party may have, and gives them at least a certain payoff regardless of the other party's choice. Any negotiation game that is rich enough to contain such offers, yields the same efficiency result. The key is that agents simply have to understand the evident availability of certain powerful strategies when they are negotiating, which from our experiments appears to be the case, at least in some settings.

With this in mind we define a general notion of "rich" negotiation games.

Definition 1 (Rich Negotiation Games) *An alternating offer negotiation game Γ includes a share-v demanding offer for some agent i of type $\theta_i \in \Theta_i$, $v \in [0, \bar{S}]$, if (in a period that i is the proposer) there exists $o \in O_i$ such that*

- *for every $r \in R_j(o)$, the realized payoff for θ_i in the current period is at least v, and*
- *for every θ_j such that $f(\theta_i, \theta_j) > 0$: there exists $r(\theta_j) \in R_j(o)$ with which the realized payoff for j in the current period is at least $\bar{S} - v$.*

*An alternating offer negotiation game is **rich** if it includes a share-v demanding offer for every share $v \in [0, \bar{S}]$, agent $i = a, b$ and type $\theta_i \in supp(f_i)$.*

The richness condition allows agents to propose efficient deals and arbitrary splits of the surplus at any point of the game. In particular, a current proposer can find some offer to guarantee herself a certain share, as long as it is accepted, regardless of the other

[14]The usual definitions of consistency apply to finite action spaces, whereas we allow for games with continuum of actions and thus conditional probability measures may need to be defined via Radon-Nikodym derivatives and so are only tied down up to sets of measure 0.

party's choice. Also a share-demanding offer leaves enough room for the current recipient to obtain the rest of the total efficient surplus. The condition characterizes strategies that are available. Thus, the definition of the richness,"including share demanding offers", imposes a "lower bound" on what kinds of offers agents are able to use. In particular, the definition does not require that agents have to use only those strategies; they are free to communicate with additional possible offers and messages.

THEOREM **2** *If a negotiation problem with n items has a known surplus $\bar{S} > 0$ and if the negotiation game Γ is rich, then in all weak perfect Bayesian equilibria:*

- *agreement is reached immediately,*
- *the full surplus is realized, and*
- *the agents' expected payoffs equal to their Rubinstein shares; i.e., $\frac{\bar{S}}{1+\delta}$ for Alice, and $\frac{\delta\bar{S}}{1+\delta}$ for Bob.*

A short outline behind the proof Theorem 2 is as follows. If there were any inefficiency on the anticipated equilibrium path, then since the agents know the potential surplus and can make demands for shares of that total surplus, there is an offer that they each know makes them strictly better off if it is immediately accepted compared to the anticipated path. Since, given the richness, they have offers that ares sure to include at least one such efficient offer and can demand any split of the total surplus, such an offer would be an improving deviation. The existence of efficient offers, under the richness of the negotiation game, then crowds out any inefficient outcome. The full details appear in the appendix, but the basic intuition is the availability of efficient deals crowds out the inefficient ones. The argument for the precise Rubinstein shares is based on an extension of that by Shaked and Sutton (1984).

The result demonstrates that if negotiating parties know the gains from trade, and if they also understand the availability of certain powerful strategies and use them, then they are incentivized to make offers that honestly reveal their private information, and counterparties are incentivized to act on that information in a manner that promotes mutual gain. Rather than posturing or hiding information, and rather than haggling over the values of individual aspects, the agents find ways to efficiently exchange information and determine optimal values for the various aspects of the deal. This is "the art of the deal".

5 Concluding Remarks

Although negotiations frequently involve several aspects of a contract or deal, traditional bargaining theory focuses on a situation in which there is a single aspect to be determined.

We extend that theory to encompass negotiations, in which deals have many aspects. Our model is descriptive. Agents freely negotiate the terms of a deal with offers and counteroffers, and they do so in the absence of any mediation. Despite the fact that they intend to serve only their own self-interest, we define a robust class of meaningful situations in which outcomes are always socially efficient. This leads to a new perspective, which would appear to have some empirical relevance regarding the costs of asymmetric information. It is a tale about the reach of the invisible hand.

In both structure and technique, our theoretical analysis is an extension of Rubinstein (1982) to allow for deals with multiple aspects and asymmetric information. The new ideas concern the way in which we decompose the knowledge structure when deals are multi-aspect, as well as the manner in which we model strategic possibilities when the interactions between agents are more complex than in bargaining theory. The decomposition of knowledge into two parts: knowledge of the possible gains from trade and knowledge of where these gains are to be found, is demonstrated to be productive. Even when the gains from trade are not approximately known, we believe that the distinction between these two forms of knowledge will be useful, and its consequence is explored in further work.[15]

When the gains from trade are known, the manner in which agents negotiate is determined by the presence of powerful strategies, which we argue are available to thoughtful players. These strategies, in a sense dominate less efficient ones. They lead the parties to honestly reveal their private information and, when they possess the private information of a counterparty, to use it in a manner that promotes mutual gain. As a consequence, information is shared truthfully and an efficient deal is reached without delay.

Our experiments complement our theoretical treatment. Beyond providing modest tests of the theory, we regard them as an integral part of our analysis, in that the dialogs which we observe provide some considerable comfort regarding the manner in which we have argued the agents negotiate. In particular, the dialogs suggest the relevance of share demanding strategies. Also, when the gains from trade are not common knowledge and share demanding strategies do not exist, the dialogs demonstrate "posturing" and the paths that lead to inefficiency. We believe that the experiments are also noteworthy in the manner in which they allow for behaviors that are both free-form as well as structured and compare the resulting outcomes. Further work will explore outcomes when the gains from trade are

[15]We remind the reader that our goal has been to present the main ideas of our research in a manner that will serve a broad audience. This has led us to lighten notation whenever possible, and in particular in the cases where the consequences of generalizations are straightforward. However, it has also forced us to exclude certain explorations which are central to a deeper understanding of the negotiation process. Two examples of these are our analysis of situations in which the gains from trade are "approximately common knowledge" and our extension of negotiations to include the case where the terms settle over time. These matters, and others, are considered in Jackson, Sonnenschein and Xing (2015).

only approximately known and test the sensitivity of our results to the assumption that discounting is the same for both agents.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1991): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58 (5), 1041-1065.

Armstrong, Mark (1999): "Price Discrimination by a Many-Product Firm," *Review of Economic Studies*, 66 (1), 151-168.

Ausubel, Lawrence M. and Ray J. Deneckere. (1989): "Reputation in bargaining and durable goods monopoly." *Econometrica*, 57 (3), 511-531.

Ausubel, Lawrence M., Peter Cramton and Ray J. Deneckere. (2002): "Bargaining with incomplete information." *Handbook of game theory with economic applications* 3: 1897-1945.

Banks, Jeffrey S. and Joel Sobel (1987), "Equilibrium Selection in Signaling Games, *Econometrica*, 55 (3), 647-661.

Bergemann, Dirk and Stephen Morris (2005): "Robust mechanism design," *Econometrica*, 73 (6), 1771-1813.

Billingsley, P. (1968): *Convergence of Probability Measures*, Wiley, New York, New York.

Charness, Gary. 2000. "Self-Serving Cheap Talk: A Test of Aumann's Conjecture." Games and Economic Behavior, 33(2), 177-94.

Charness, Gary and Martin Dufwenberg. 2006. "Promises and Partnership." Econometrica, 74(6), 1579-601

Cohn, Z. (2010): "A Note on Linked Bargaining," *Journal of Mathematical Economics*, 46 (2), 238-247.

Dvoretzky, A., Kiefer, J., and Wolfowitz, J. (1956): "Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator," *The Annals of Mathematical Statistics*,, 27 (3) 642-669.

van Damme, E., Selten, R., and Winter, E. (1990): "Alternating bid bargaining with a smallest money unit," *Games and Economic Behavior*, 2 (2), 188-201.

Ellingsen, Tore, and Magnus Johannesson (2004a) "Promises, threats and fairness," *Economic Journal* 114, 397-420.

Ellingsen, Tore, and Magnus Johannesson (2004a) "Is there a hold-up problem?" *Scandinavian Journal of Economics* 106, 475-494.

Fang, H. and P. Norman (2003): "An Efficiency Rationale for Bundling of Public Goods," Cowles Foundation Discussion Paper no 1441.

Feltovich, Nick, and Joe Swierzbinski (2011) "The role of strategic uncertainty in games: An experimental study of cheap talk and contracts in the Nash demand game," *European Economic Review* 55(4): 554-574.

Fudenberg, Drew and David Levine (1983): "Subgame-perfect equilibria of finite-and infinite-horizon games," *Journal of Economic Theory*, 31 (2), 251-268.

Fudenberg, Drew and David Levine (1994): "Efficiency and Observability in Games with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62 (1), 103-135.

Fudenberg, Drew and David Levine and Jean Tirole (1985), "Infinite-Horizon Models of Bargaining with One- Sided Incomplete Information," in Alvin Roth (ed), *Game Theoretic Models of Bargaining*, Cambridge, England: Cambridge University Press.

Green, Edward J. and Robert Porter (1984): "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52 (1), 87-100.

Grossman, S. and M. Perry (1986): "Sequential Bargaining under Asymmetric Information," *Journal of Economic Theory*, 39 (1), 120-154

Inderst, R. (2003): "Alternating-offer bargaining over menus under incomplete information," *Economic theory*, 22 (2), 419-429.

Jackson, Matthew O. (2001) "A Crash Course in Implementation Theory," Social Choice and Welfare, Vol. 18, No. 4, pp 655-708.

Jackson, Matthew O. (2003) "Mechanism Theory," in the *Encyclopedia of Life Sciences*, EOLSS Publishers: Oxford UK. http://ssrn.com/abstract=2542983

Jackson, Matthew O., Leo K. Simon, Jeroen M. Swinkels and William R. Zame (2002): "Communication and equilibrium in discontinuous games of incomplete information," *Econometrica*, 70 (5), 1711-1740.

Jackson, Matthew O. and Hugo F. Sonnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75 (1), 241 - 258.

Jackson, Matthew O., Hugo F. Sonnenschein, and Yiqing Xing (2015): "A Theory of Negotiation," mimeo: Stanford.

Kennan, John and Robert Wilson (1993): "Bargaining with private information," *Journal of Economic Literature*, 31 (1), 45-104.

Matsushima, H. (2008): "Detail-free mechanism design in twice iterative dominance: Large economies," *Journal of Economic Theory*, 141 (1), 134-151.

McAfee, R.Preston, (1992): "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," *Journal of Economic Theory*, 56 (2), 266-293.

McAfee, R.Prestion, John McMillan and Michael D. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling and the Correlation of Values," *Quarterly Journal of Economics*, 104 (2), 371-383.

Myerson, Roger and Mark Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265-281.

Nash, John (1953). "Two-Person Cooperative Games," *Econometrica* 21 (1), 128-140.

Radner, Roy (1981): "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica* 49, 1127-1148.

Roughgarden, T. and I. Talgam-Cohen (2013): "Optimal and near-optimal mechanism design with interdependent values," *Proceedings of the fourteenth ACM conference on Electronic commerce* 767-784.

Rubinstein, Ariel (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica* 50 (1): 97109.

Rubinstein, Ariel (1985): "A Bargaining Model with Incomplete Information about Time Preferences," *Econometrica*, 53 (5), 1151-1172.

Satterthwaite, M., S. Williams and K. Zachariadis (2014): "Optimality versus practicality in market design: a comparison of two double auctions," *Games and Economic Behavior*, 86, 248-263.

Sen, Arunava (2000): "Multidimensional bargaining under asymmetric information." *International Economic Review*, 41 (2): 425-450.

Segal, Ilya and Michael Whinston (2016). "Property rights and the efficiency of bargaining," *Journal of the European Economic Association*, 14(6): 1287-1328.

Shaked, Avner, and John Sutton (1984) "Involuntary unemployment as a perfect equilibrium in a bargaining model," *Econometrica*, 1351-1364.

Stahl, I. (1972): "Bargaining theory," Doctoral thesis, Stockholm School of Economics.

Skrzypacz, A. and J. Toikka (2015): "Mechanisms for Repeated Trade," *American Economic Journal: Microeconomics*, 7(4): 252293.

Townsend, Robert (1982): "Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information," *Journal of Political Economy*, 90 (6), 1166-1186.

Wilson, Robert (1987) "Game theoretic analysis of trading processes," in T. Bewley (Ed.), *Advances in Economic Theory*, Cambridge University Press (1987)

Appendix: Proofs

Proof of inefficiency in the example from Section 2.1:

Consider the two-good problem, in which the seller's cost is $(0, 0)$ and the buyer's value is $(10, 50)$ or $(50, 10)$, equally likely and known only to the buyer. The surplus is 60 with certainty.

We first show that all equilibria are inefficient if the two parties negotiate item-by-item (as described in 2.1) and $\delta < .63$ and the seller makes the first offer.

First note that to guarantee trading in the first period of both goods, neither price can exceed 10, and so the seller's payoff would be at most 20. We show that the seller has another strategy whose expected payoff is strictly more than 20, and so efficiency is impossible.

In particular, let L_s be the seller's worst continuation payoff in any seller-offer period in any equilibrium (with both items remaining). We only need to show that $L_s > 20$.

The fact that L_s is the seller's worst continuation payoff implies that when the buyer makes an offer, he gets a continuation payoff of at most $\delta(60 - \delta L_s)$ since the seller can always reject on both items and get at least L_s in the continuation, which leaves at most $(60 - \delta L_s)$ for the buyer in terms of a continuation value.

So, consider seller offering some (p, p) in the first period, with some $p > 10$. The buyer rejects p on the value-10 item, and accepts p on the value-50 item for sure if $p < \tilde{p}$, s.t.

$$50 - \tilde{p} = \delta(60 - \delta L_s). \tag{1}$$

Therefore, with an offer of $(\tilde{p}-\epsilon, \tilde{p}-\epsilon)$ for any $\epsilon > 0$, the seller can always get an acceptance on the value-50 item, and so a payoff of at least

$$\tilde{p}-\epsilon = 50 - \delta(60 - \delta L_s) - \varepsilon.$$

Thus, since L_s is the lowest possible seller's continuation payoff, it must exceed the payoff from the above action $(\tilde{p}-\epsilon, \tilde{p}-\epsilon)$. This requires that

$$L_s \geq 50 - \delta(60 - \delta L_s) - \varepsilon. \tag{2}$$

Since ϵ can be arbitrarily small, it follows that

$$L_s \geq \frac{50}{1+\delta} - \frac{10\delta}{1-\delta^2} \tag{3}$$

As a result, $L_s > 20$ for any $\delta < 0.63$, as claimed. ▮

Now we follow a similar proof for the 4-good problem (Problem #2).

The surplus is 70 instead of 60. Again let L_s be the seller's worst continuation payoff in any seller-offer period in any equilibrium (with all items remaining). Again, we bound L_s and show that it has to be more than the seller could get by offering prices that are sure to lead to efficient trade.

Suppose the seller has $(0, 0, 40, 40)$, wlog. Consider the seller's offer of some $(p, p, 50, 50)$ with some $p > 10$ (which will not be efficient). The buyer accepts at least one price for sure $p < \tilde{p}$, s.t.

$$50 - \tilde{p} = (70 - \delta L_s)\delta, \tag{1'}$$

Accordingly, the analog to equation (2) becomes: [16]

$$L_s \geq \tilde{p} - \epsilon. \tag{2'}$$

The above is resulted from the fact that one of the p's must be accepted. If some of the other prices are accepted, then L_s only becomes higher. Therefore

$$L_s \geq \frac{50}{1+\delta} - \frac{20\delta}{1-\delta^2} \tag{3'}$$

Therefore, $L_s > 30$ for any $\delta < 0.33$.

In any offer that guarantees efficient trade in the first period, the seller's payoff is at most 30, which is the payoff from an offer of $(10, 10, 50, 50)$. Therefore every equilibrium is inefficient for $\delta < 0.33$. ▮

[16] We note that $\tilde{p}$ is constructed so that "rejecting all the four prices" is strictly dominated by accepting one of the first two items (on which his value is 50). Therefore, the buyer should accept at least one of the four prices. Inequality (2') holds even if the buyer accepts only either item-3 or item-4, in which case the seller gets 50 which is higher than $\tilde{p}$, by construction. We intentionally state this argument in a conservative manner, so as to avoid the discussion of off-path beliefs.

Inefficient equilibrium with δ close to 1.

We show that there exists some inefficient equilibrium for *any* $\delta < 1$. We do this for the simplified two-good example, in which the seller has $(0, 0)$ and the buyer has $(10, 50)$ or $(50, 10)$, equally likely. Our goal is not only to show the existence, but to show that every equilibrium that satisfies an intuitive refinement of belief updating is inefficient.

ASSUMPTION 1 *If the initial offer is (p, p) with some $50 \geq p > 10$, and the buyer accepts on one item and rejects the other, then the posterior belief is such that the buyer has value 50 on the accepted item and 10 on the rejected item, and the continuation game that follows is as if it had commonly known valuations 0 and 50.*

LEMMA **1** *In the two-good example under the bargaining protocol defined in 2.1, every equilibrium that satisfies Assumption 1 is inefficient, for $\forall \delta < 1$. In particular, at most one of the two goods are traded in the initial period.*

Proof of Lemma 1:

This proof is very similar to that for the not-too-high-δ case. Again let L_s be the seller's worst continuation payoff in any seller-offer period in any equilibrium (with both items remaining). When the buyer makes an offer, he gets a continuation payoff of at most $\delta(60 - \delta L_s)$.

Consider a seller's offer of (p, p) with some $p > 10$. The buyer rejects p on the value-10 good, and accepts p on the value-50 good and rejects p on the value-10 item for sure if $p < \tilde{p}$, s.t.

$$50 - \tilde{p} + \frac{10\delta}{1+\delta} = (60 - \delta L_s)\delta. \tag{1''}$$

Here, there is an additional term $\frac{10\delta}{1+\delta}$ which is the continuation payoff from the value-10 good. This is the corresponding Rubinstein share, following Assumption 1, the remaining game is a complete information bargaining game of alternating offers.

Therefore, with an offer of $(\tilde{p} - \epsilon, \tilde{p} - \epsilon)$ for any $\epsilon > 0$, the seller can always get an acceptance on the value-50 item, and a payoff of at least $\tilde{p} - \epsilon$, plus a continuation payoff of $\frac{10\delta^2}{1+\delta}$ from the other good. By definition of L_s, it follows that

$$L_s \geq \tilde{p} - \epsilon + \frac{10\delta^2}{1+\delta}. \tag{2''}$$

Combining (1″) and (2″) implies that

$$L_s \geq 50 + \frac{10\delta}{1+\delta} - (60 - \delta L_s)\delta - \epsilon + \frac{10\delta^2}{1+\delta},$$

and since ϵ can be arbitrarily small, it follows that

$$L_s \geq \frac{50}{1+\delta} > 25, \forall \delta < 1. \tag{3''}$$

Recall that an offer that guarantees trade of both items in the initial period requires that the seller's expected payoff is at most 20, which is smaller than L_s. Hence efficiency is not possible in any equilibrium that satisfies Assumption 1. ∎

Proof of Theorem 1:
This Theorem is a corollary to Theorem 2 given the fact that "direct negotiations" includes share-demanding offers. In particular, for any v, a share-v demanding offer is to announce the truth θ_i^n and demand a payoff of v. ∎

Proof of Theorem 2:

We presume that a moves first, but the proof holds with roles reversed if it is b.

We begin with some notation. Let $\Theta_i(\theta_j) \equiv \{\theta_i \mid (\theta_i, \theta_j) \in \Theta\}$ be the set of i's types that are "possible" given the other's type being θ_j, and let $\Theta_i \equiv \bigcup_{\theta_j} \Theta_i(\theta_j)$ be the set of all possible types of i in the negotiation game.

Let M_i^t $[L_i^t]$ be the supremum [infimum] of the expected continuation payoff for agent i, starting at the beginning of period t over all equilibrium continuations, all histories that arrive at this period, and all i's types in $\theta_i \in \Theta_i$.

We now establish the upper and lower bounds of the two parties' payoff in any equilibrium, and show that they all correspond to a unique equilibrium payoff that corresponds to immediate and efficient deal, and the Rubinstein shares.

We first show that $M_a^t \leq \frac{\bar{S}}{1+\delta}$ for any odd t.

At $t+1$ (t odd), Bob makes the offers. We argue that Bob with any type $\theta_b \in \Theta_b$ can guarantee a payoff arbitrarily close to

$$\ell_b^{t+1} \equiv \bar{S} - \delta M_a^{t+2}.$$

Bob does so by offering a share-v_{t+1} demanding offer with $v_{t+1} \equiv \ell_b^{t+1} - \eta$ for $\eta > 0$ arbitrarily small. Such an offer is accepted for sure for Alice with any type $\widehat{\theta}_a \in \Theta_a(\theta_b)$, by the known-surplus assumption and the definition of δM_a^{t+2} and share-demanding offer: According to the definition of share-demanding offers, Alice can accept the offer and find a deal that gives herself at least $\bar{S} - v_{t+1} = \delta M_a^{t+2} + \eta$, which exceeds δM_a^{t+2}, the present value of the payoff from the continuation of the game by rejecting the offer. In addition, Bob gets a payoff of $\ell_b^{t+1} - \eta$ if the above offer is accepted, regardless which terms Alice picks. This implies that

$$L_b^{t+1} \geq \ell_b^{t+1} = \bar{S} - \delta M_a^{t+2}.$$

At t, Alice makes offer. We argue that Alice with any type $\theta_a \in \Theta_a$ can get a payoff at most $\bar{S} - \delta L_b^{t+1}$: With any type $\theta_b \in \Theta_b$, by rejecting an offer at t, Bob's payoff from the continuation of the game has a present value of at least $\delta L_b^{t+1} - \delta\eta$ for $\forall \eta > 0$. Hence the payoff left to Alice with $\theta_a \in \Theta_a$ is at most $\bar{S} - \delta L_b^{t+1}$, as the (expected) surplus is $\bar{S}$ by construction. Thus, it follows that

$$M_a^t \leq \bar{S} - \delta L_b^{t+1} \leq (1-\delta)\bar{S} + \delta^2 M_a^{t+2}.$$

The above is true for any $t = 1, 3, 5, ...,$. Iteratively substituting on right hand side, the above leads to

$$M_a^t \leq (1-\delta)\bar{S} \sum_{k=0}^{\infty} (\delta^2)^k,$$

or

$$M_a^t \leq \frac{\bar{S}}{1+\delta}$$

for any $t = 1, 3, 5, ...,$, as claimed.

Also, note that from the above argument we had $L_b^{t+1} \geq \bar{S} - \delta M_a^{t+2}$, and so it follows that

$$L_b^{t+1} \geq \frac{\bar{S}}{1+\delta}.$$

By a parallel argument to the one establishing that $M_a^t \leq \frac{\bar{S}}{1+\delta}$, but reversing the roles, it follows that for any $t = 1, 3, 5, ...$:

$$M_b^{t+1} \leq \frac{\bar{S}}{1+\delta}.$$

Therefore $M_b^2 = L_b^2 = \frac{\bar{S}}{1+\delta}$. Repeating the argument with the η's then implies that $L_a^1 \geq \frac{\bar{S}}{1+\delta}$. In addition, $L_b^1 \geq \delta L_b^2$ since Bob can always reject offers at $t = 1$, and therefore $L_b^1 \geq \delta \frac{\bar{S}}{1+\delta}$. As a result,

$$L_a^1 + L_b^1 \geq \frac{\bar{S}}{1+\delta} + \delta \frac{\bar{S}}{1+\delta}.$$

Note that the right hand side is S, the total surplus from efficient trade. Hence the negotiation outcome must be efficient, which requires immediate trade with the efficient deal. The payoff divisions correspond to the Rubinstein shares. ∎

EXPERIMENTAL SUPPLEMENT

To "Negotiation as the Art of the Deal" by
Jackson, Sonnenschein, Xing, Tombazos, and Al-Ubaydli

(for online publication)

This document details the experimental design, provides comprehensive analysis of the key empirical findings, and documents the experimental instructions.

NOTE: The term "period" used in the paper corresponds to the term "bargaining stage" used in this supplement as well as in the instructions that were distributed to experimental participants (which are provided in the Appendix of this supplement). The use of different terms for the same concept is guided by the different audiences for which they are intended.

1. EXPERIMENTAL DESIGN

The goal of the experiments is to gain insight into how subjects negotiate in different settings. Of particular interest is a comparison of negotiation problems 1 and 2 discussed in the paper.

1.1. STRUCTURE

The experiments use two bargaining protocols. The first is an implementation of alternating offers. The second is a free form setup where subjects can engage in text-based chat and, simultaneously, issue offers. Within each protocol, we have two treatments that follow the basic structure of problems 1 and 2 discussed in the paper. Treatment assignment is at the session level and, in each treatment, all rules are common knowledge.

Each treatment requires an even number of participants who experience 10 rounds of bargaining in pairs. We denote the first four rounds "practice" and the last six rounds "real". Earnings are based on an *ex post* random selection of a real round. Each round, participants are assigned the role of either buyer or seller. During the practice rounds, participants experience both roles. Once the practice rounds are completed, participants are assigned the role of either buyer or seller and remain in that role throughout the real rounds. At the beginning of each round sellers and buyers are anonymously and randomly paired (random stranger). Upon role assignment, sellers and buyers are privately informed of their cost(s) and value(s), respectively. Cost(s) and value(s) are equiprobable across rounds and participants.

For each pair, each bargaining round is composed of up to 8 stages of bargaining. A financial discounting penalty is applied based on the number of stages required to reach agreement: 0% in the first stage, increasing arithmetically by 10% every stage until 70% in stage 8. If a pair of traders fails to reach agreement within 8 stages, bargaining is suspended and traders earn zero profit for that round.

In what follows, financial rewards are expressed in Australian dollars. In the 1-good treatments, seller costs are drawn from {$0, $160} and buyer values from {$40, $200} with equal probability. Hence, three out of the four possible combinations of cost and value imply a strictly positive surplus, whereas one combination implies a strictly negative surplus. The expected surplus of a pair of traders is $70.

In the 4-good treatments, the cost and value pairs are {($0, $10), ($0, $50), ($40, $10), ($40, $50)} in every round, but the assignment of each of these combinations on each of the four goods is randomized. Consequently, by way of an example, a seller with a cost of $0 for the first good does not know whether the buyer's value for that good is $10 or $50. Therefore, total surplus per round is fixed at $70, equal to the expected surplus in the 1-good treatments.

In 1-good treatments, offers are unidimensional trading-prices. In 4-good treatments, offers are a list of the goods included in a bundle that is on offer in addition to a price for the entire bundle. A bundle may include anywhere from one to four goods. There are no restrictions on which goods may be included.

In any given round of alternating offer treatments and for any given pair of traders, the computer selects randomly one of the two traders to make the first offer. The responder can accept or reject such an offer. If a trader accepts the first offer, the round is concluded for the pair that this trader belongs. If the first offer is rejected, the other trader has the opportunity to make a counter-offer. This alternating process continues until one of the pair of traders makes an offer that the other trader accepts. An important difference between the 1-good and the 4-good alternating offer treatments is that in the case of the former a proposer can issue a single offer whereas in the latter a trader can issue multiple (up to six) offers in any given stage. Of course, a responder can accept, at most, a single offer. A time limit is given for proposing and for responding. In stage 1 a proposer is given 30 seconds and a responder is given 10 seconds. By contrast, in stages 2-8 a proposer is given 15 seconds and a responder is given 5 seconds. In the case of any given pair of traders, a bargaining stage encompasses the proposer's offer and the responder's decision to accept or reject.

Unlike the alternating offers setup, free form treatments are considerably less structured. In any given round, traders are allowed to exchange messages and, simultaneously, offers with their partners. Traders can issue as many messages or offers to their partners as they like without waiting for their partner to respond. Offers do not expire across bargaining stages and may be accepted at any time. Time limits are comparable to those that apply in the case of the alternating offers treatments: 40 seconds for stage 1, and 20 seconds for all remaining stages. Bargaining automatically advances to the next stage once a stage concludes without agreement.

1.2. Outcome Variables and Hypothesis

Let $i \in \{1,2, \dots\}$ denote session, $t \in \{1,2, \dots, 10\}$ denote bargaining round, $j \in \{1,2, \dots J_{it}\}$ denote bargaining pair (J_{it} is the number of pairs in bargaining period t of session i), and $r \in \{1,2, \dots, 8\}$ denote bargaining stage.

Let $s_{itj} \in \{\$0, \$40, \$70, \$200\}$ denote the potential surplus for pair j in round t of session i. Let d_{itjr} be a dummy variable that takes the value "1" if pair j in round t of session i successfully trades in stage r, and "0" otherwise. Let $D_{itj} = \sum_r d_{itjr}$. Trade takes place when $D_{itj} = 1$, and fails to take place otherwise. Finally, let R_{itj} denote the stage in which a trade occurred for pair j in round t of session i, and let it equal "0" in the event that no trade occurred.

At this stage, we presage our results by pointing out that no negative-surplus trades occurred. Conditional on this fact, and given the goal of the experiment, we had no interest in the actual trade price when a trade occurred, despite the availability of such data.[1]

Our primary outcome variable of interest is efficiency, e_{itj}, which we define as follows for any pair j of round t of session i:

$$e_{itj} = \left(\frac{11 - R_{itj}}{10}\right) D_{itj}\, s_{itj} \; if \; s_{itj} > 0$$

It is the percentage of potential surplus realized once the discount factor has been applied, and taking into account the possibility that the traders may fail to trade. Note that e_{itj} is undefined when the potential surplus is zero.

Hypothesis: Efficiency will be higher in 4-good treatments than in 1-good treatments that use the same bargaining protocol (i.e., alternating and free-form).

Looking ahead to the econometric testing, under certain specifications, we will need to test our primary hypothesis by comparing efficiency that has been aggregated at the round level or session level. In these situations, there are two ways of calculating average efficiency: weighted and unweighted means.

To explain the difference, we begin by arbitrarily reordering the pairs in each round of each session so that the first $\tilde{J}_{it} \leq J_{it}$ pairs are the ones for whom the surplus is strictly positive, $s_{itj} > 0$, and therefore for whom efficiency is defined. This is without loss of generality.

[1] It should be noted that in 5 out of a total of 1,011 incentivized subject interactions an efficient trade took place at a price that implied negative earnings for one of the two traders. In 4 of these 5 cases this appeared to be the result of human error, whereas in the reamining case the choice to make a loss fit a pattern of random choices throughout the experiment. In three of the 5 cases the losses were sufficiently large (i.e., greater than $7) so as to bankrupt participants. A bankrupt participant is suspended from subsequent rounds of the same session. This requires another to be forced to sit out as the session continues with an even number of players. This second participant was randomly selected at the beginning of each round of any given session.

In the case of round-level aggregation:

$$\bar{e}_{it}^{u} = \frac{1}{\tilde{J}_{it}} \sum_{j=1}^{\tilde{J}_{it}} e_{itj}\,, \qquad \bar{e}_{it}^{w} = \frac{1}{\sum_{j=1}^{\tilde{J}_{it}} s_{itj}} \sum_{j=1}^{\tilde{J}_{it}} s_{itj} e_{itj}$$

The unweighted average assigns each pair in the round an equal weight. The weighted average assigns each pair a weight that is equal to the potential surplus for the pair. This is equivalent to defining average efficiency for the round as the total surplus realized across all pairs divided by the total potential surplus across all pairs.

Equivalently, in the case of session-level aggregation:

$$\bar{e}_{i}^{u} = \frac{1}{\sum_{t=1}^{6} \tilde{J}_{it}} \sum_{t=1}^{6} \sum_{j=1}^{\tilde{J}_{it}} e_{itj}\,, \qquad \bar{e}_{it}^{w} = \frac{1}{\sum_{t=1}^{6} \sum_{j=1}^{\tilde{J}_{it}} s_{itj}} \sum_{t=1}^{6} \sum_{j=1}^{\tilde{J}_{it}} s_{itj} e_{itj}$$

In each case, both weighted and unweighted averages are valid; we therefore report both.

1.3. Procedure

The experimenter script and the experimental instructions are in the experimental materials appendix which follows section 3. Here, we briefly describe the most salient features of the recruitment and implementation procedures.

The experimental sessions took place at the Monash University Laboratory for Experimental Economics (MonLEE) on the Clayton campus of Monash University in Melbourne, Australia. Subjects were predominantly undergraduate students from Monash University. They were recruited from a database of individuals who expressed interest in participating in economic experiments that is maintained by MonLEE. No subject participated more than once. We conducted 18 sessions employing a total of 340 subjects. Each session involved between 8 and 24 subjects as summarized in Table 2.1.1a of the next section.

Within each cluster of four consecutive sessions within the first 16 sessions, assignment of each of the four treatments to each of the sessions in the cluster was random. By the conclusion of the 16th session, we determined that we had sufficient observations for statistical inference. Sessions 17 and 18 were added to introduce parity between the number of subjects involved in alternating offer treatments.

Upon arrival at the lab subjects earned an attendance fee of $10. This has two components, a show-up fee of $3 and a participation fee of $7. Any positive earnings that subjects made by trading one or more goods during an experiment were added to this $7, whereas any negative earnings were subtracted from this $7. If a subject made a loss in excess of $7 in any given round, then their participation in subsequent rounds of the session was suspended and the participant only received their show-up fee of $3.

At the beginning of a session subjects were given written instructions. These instructions were also read aloud in an effort to make the rules of the game common knowledge. Following this step, subjects were given the opportunity to ask questions in private. The experiment then commenced. Throughout a session, subjects interacted only through the MonLEE computer network running an application written using the *z-Tree* experimental package (Fischbacher, 2007). Following the 10th round of any given session; subject earnings were determined, subjects were paid privately, and their participation in the experiment was then concluded.

2. EMPIRICAL RESULTS

2.1. DESCRIPTIVE STATISTICS

We ran 18 sessions involving 340 subjects. Tables 2.1.1a and 2.1.1b provide relevant details, including a breakdown of the allocation across treatments. All sessions were conducted during Monash University's teaching semester 2, 2017. Variation in the number of subjects across sessions was largely driven by subject availability.

Table 2.1.1a: Sessions: Details

Session	Date	Goods	Structure	Subjects
1	1-Aug-17	1	Open chat	8
2	1-Aug-17	4	Open chat	10
3	12-Sep-17	4	Alternating	20
4	19-Sep-17	1	Alternating	24
5	21-Sep-17	4	Open chat	22
6	25-Sep-17	1	Alternating	24
7	25-Sep-17	1	Open chat	20
8	26-Sep-17	4	Alternating	24
9	3-Oct-17	4	Alternating	14
10	3-Oct-17	1	Alternating	24
11	5-Oct-17	1	Open chat	24
12	5-Oct-17	4	Open chat	22
13	12-Oct-17	4	Open chat	20
14	25-Oct-17	1	Open chat	24
15	26-Oct-17	4	Alternating	16
16	21-Oct-17	1	Alternating	22
17	22-Oct-17	4	Alternating	12

18	27-Oct-17	4	Alternating	10

Table 2.1.1b: Sessions: Summary

Treatment	Sessions	Subjects
1 good, alternating	4	94
1 good, open chat	4	76
4 goods, alternating	6	96
4 goods, open chat	4	74

The primary summary statistics are shown in Table 2.1.2, and in Figure 2.1.1. Note that these summary statistics define an observation as a specific pair in a specific period in a specific session, which we refer to as a unique matching.

Table 2.1.2: Means and SDs for Primary Variables where Observations are Matchings

Treatment	Surplus	# Obs	% Trading	Bargain Stages	Efficiency: Unweighted	Efficiency: Weighted
1 good, alternating	0	71	0%	N/A	N/A	N/A
	40	142	81%	5.1 (2.1)	47% (30%)	N/A
	200	69	100%	3.4 (2.2)	76% (22%)	N/A
	Pooled	282	87%	4.5 (2.3)	57% (30%)	68%
4 goods, alternating		284	98%	1.9 (1.4)	85% (21%)	85% (21%)
1 good, open chat	0	56	0%	N/A	N/A	N/A
	40	113	81%	3.3 (2.2)	62% (37%)	N/A
	200	54	100%	2.8 (2.0)	82% (20%)	N/A
	Pooled	223	87%	3.1 (2.2)	68% (33%)	76%
4 goods, open chat		222	100%	1.4 (0.9)	91% (18%)	91% (18%)
All		1011	94%	2.5 (2.1)	77% (29%)	80%

Figures in parentheses denote standard deviations. In pooled data, “% Trading” shows the percentage trading conditional on a strictly positive surplus. “Bargain Stages” refers to the stages taken to agree on a trade conditional on a trade occurring.

Figure 2.1.1: Means of Key Variables where Observations are Matchings

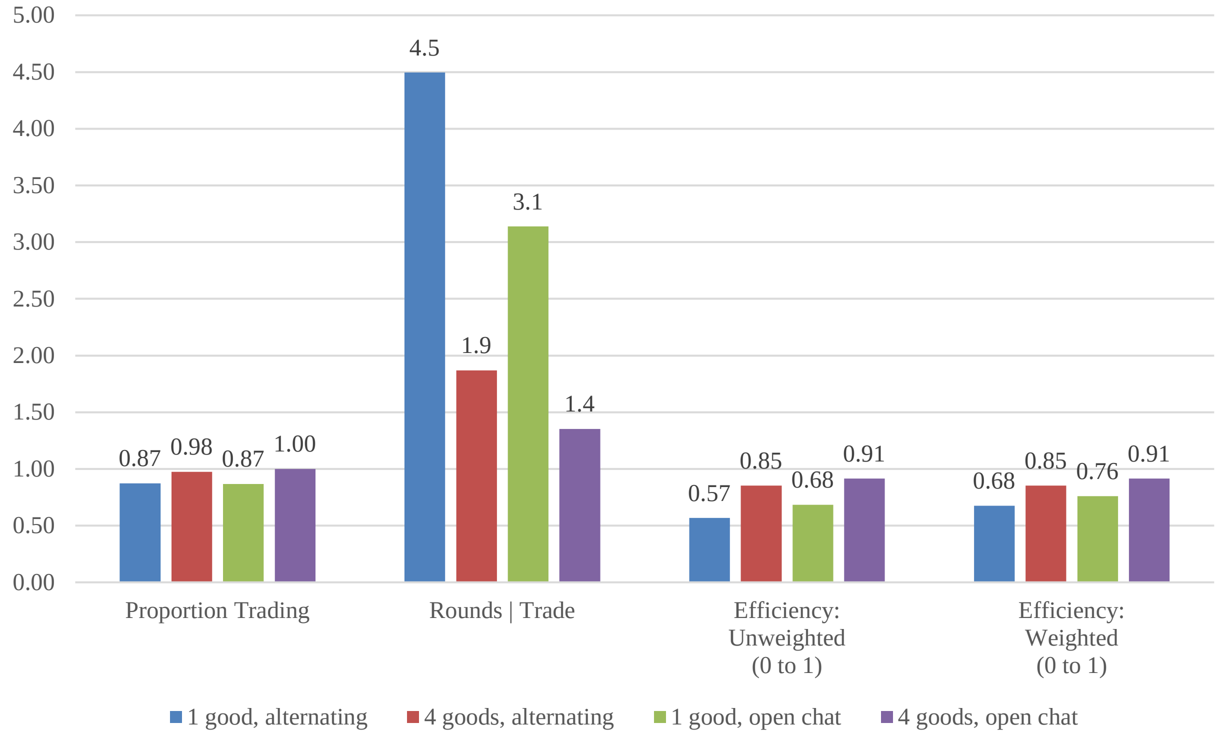

Percentages are expressed as proportions to maintain a single vertical axis. “Proportion trading” is conditional on the potential surplus being positive.

The data have several notable features.

First, 4-good treatments result in higher levels of average efficiency as compared to their 1-good counterparts (28% higher for alternating, unweighted, 17% higher for alternating, weighted; 23% higher for open chat, unweighted, 15% higher for open chat, weighted).

Second, the difference in efficiency between corresponding 1-good and 4-good treatments is partially due to failure to trade in 1-good treatments when surplus is positive; and partially due to the fact that subjects typically take a greater number of bargaining stages in 1-good than 4-good treatments to reach an agreement, therefore incurring a higher discount factor.

Third, the increase in efficiency that is observed as the protocol switches from alternating offers to free form bargaining is consistent with an expansive literature documenting the role of cheap talk in promoting coordination (e.g., Charness 2000, Charness and Dufwenberg 2006). Of course, such matters extend beyond the scope of inquiry of our study.

These remarks are all based on comparisons of the point values of sample means. In the next section, we conduct formal inference.

In the main paper, in anticipation of the statistical inference, we report sample means *where observations are at the session level*. For the readers' convenience, we reproduce Table 2 of the paper as Table 2.1.3, and Table 3 of the paper as Table 2.1.4.

Naturally, the *matching* level results of Table 2.1.2 and the *session* level results of the following Tables will exhibit differences.

Table 2.1.3: Percentage Loss of Surplus

	1 good	4 goods	p-value
Structured	32.8%	14.3%	0.013
Free-form	25.2%	7.9%	0.014
p-value	0.232	0.043	

These sample means treat each session as an observation. This corresponds to Table 2 in the main paper. P-values are from a t-test.

Table 2.1.4: Fraction of Efficient Pairs Eventually Trading

	1 good	4 goods	p-value
Structured	87.1%	97.6%	0.041
Free-form	86.5%	99.6%	0.022

These sample means treat each session as an observation. This corresponds to Table 3 in the main paper. P-values are from a t-test.

As we discuss in the paper, incentives for posturing in 1 good but not in 4 good treatments explain the markedly lower efficiency in the former. Posturing can manifest in a number of different ways. In free form treatments, this will lead to delays in reaching a deal and includes misrepresenting one's own cost or value. Tables 4 and 5 of the paper present relevant evidence. We reproduce them below as Tables 2.1.5 and 2.1.6.

Table 2.1.5: Percent of Pairs with Positive Surplus Trading by Rounds

Bargaining stage	1	2	3	4	5	6	7	8	Not Trading
1 good Free-Form	24.6	19.8	12.0	10.2	7.2	3.6	2.4	7.2	13.2
4 goods Free-Form	81.1	11.3	2.3	1.8	1.8	1.4	0.0	0.0	0.5

Note that the percent not trading in Table 2.1.5 is of all subject matchings, and so that is why, for instance, the 13.2% number for the 1 good free-form treatment is 0.3% different from the 86.5% eventually trading of Table 2.1.4. Table 2.1.4 considers each session as an observation, and then averages across sessions.

Table 2.1.6: Fraction of Declarations that Are Untrue in Free-Form Treatments

Bargaining stage	1	2	3	4	5	6	7	8	Average
1 good	92/350	21/77	11/41	11/26	5/21	1/16	1/6	1/12	26.0%
4 good	3/148	0/2	0/1	n.a.	n.a.	n.a.	n.a.	n.a.	2.0%

In each fraction, the numerator is the number of untrue declarations, while the denominator is the number of total declarations

As we discuss in the paper, determining what subjects declare in chats is generally clear. However, at times, making such determinations may require subjective judgement. To limit the role of subjectivity we have adopted the following rules in how we process exchanges between subjects:

Rule 1: We define a message to begin when a subject starts typing in the chat box and to end when the subject hits "ENTER" and the message is transmitted to their trade partner.

We adopt this rule because, on occasion, a participant will send multiple messages before their partner responds. Under the circumstances, this rule promotes consistency in how exchanges are processed.

Rule 2: A declaration, or a potential root of an indirect declaration, may reside within a message (as defined by Rule 1) if this message includes at least one number. The number may be written using either Arabic numerals or text. A declaration may either be explicit or implicit. The former state the declared value or cost explicitly. The latter imply the declared value or cost.

Examples of explicit declarations: "Honestly sitting on a $40 value" (session 1, round 1, group 1), "value is 40" (session 7, round 2, group 5), "40" (session 7, round 3, group 2), are all interpreted as direct declarations that the subject's value is 40.

Example of an implied declaration: "if you can't go lower than 40 then don't bother" (session 11, round 4, group 3). Here we consider that the buyer is declaring a value of 40.

Rule 3: In instances where messages that satisfy rule 2 (by containing a number) are unclear, the message is considered as a potential root of an indirect declaration and the adjoining messages are also considered to decipher what message a subject is intended to convey.

Example (from session 11, round 5, group 2): Consider the buyer's declaration root "whats not 200". What information this may convey is unclear. In this light, we consider this message a potential root of an indirect declaration and examine the adjoining messages in conjunction with the declaration root. The complete sequence is: Seller: "what do you have", Buyer: "whats not 200", Seller "40?", Buyer: "yea". The indirect declaration by the buyer here is 40 even though the number 40 was not explicitly stated by this buyer.

2.2. Formal Inference

2.2.1. Inferential Strategy

These data are organized at three levels: session, round, and pair (of traders). Unlike conventional hierarchical data, these data involve overlapping layers, because the pairs are repeatedly drawn randomly from the same people in each session. The dimensionality of possible pairs means that there is no parsimonious or tractable way to capture the statistical dependence between the data when presented at the level of the pair.

This means that the conventional suite of models, such as clustering and multi-level modeling, requires arbitrary additional assumptions if they are to be deployed, and the results—specifically the implied standard errors—are likely to be highly sensitive to the precise arbitrary assumption settled upon.

In light of this, we consider three inferential strategies.

Strategy 1: Only use first-round data. In principle, the first round involves pairing with no history-based dependence in the data. In fact, this is not true, since the subjects play four practice rounds prior to the first real round, presenting us with a quandary: either we use unincentivized data (first practice round), or we use partially-contaminated data (first real round). We take the latter option, and simply ignore the contamination resulting from practice rounds. To account for the within-session dependence, we use session-level clusters, which is likely to be a reasonable approximation of the data-generating process since we only sample one round.

Setting aside the issue of the practice round contamination, the main advantage of this approach is that it requires minimal assumptions on the nature of dependence across observations. However, it involves ignoring 83% of the data, which lowers power. In addition, it brings into question the representativeness of the data from the first round. This is an important concern because as subjects experience successive inventivised rounds over time they are likely to benefit from learning effects that will result in the kind of mature behaviour that is of interest to this study.

Strategy 2: Session-level aggregation. This is a hyper-conservative approach that involves using one observation per session, resulting from averaging variables across all rounds and pairs in that session.

The main advantage of this approach is that it requires no arbitrary dependence assumptions. However, it involves an even greater decline in the number of observations than strategy 1, severely compromising power. Yet, since it incorporates data from all rounds, it potentially yields more representative insights.

Strategy 3: Round-level aggregation with a simple dependence model. This approach entails calculating the sample means for each round in each session (i.e., averaging across pairs in the round), and then clustering standard errors at the session level as a rudimentary model of dependence across rounds. This yields an intermediate number of observations (since most sessions involve more pairs per round than the total number of rounds).

The main advantage of this approach is that, if the dependence model is a reasonable approximation of the data-generating process, it results in six times as much data as strategy 2. Moreover, it uses data from all rounds, potentially allowing us to correct for learning effects. However, the drawback is that the dependence model is potentially a poor approximation of reality.

None of the three strategies is without flaws.

2.2.2. Choice of Statistical Tests

In all tests, we have two treatments. The data are unpaired. We therefore deploy three types of statistical tests in the *ex post* inference.

1. A conventional unpaired t-test.
2. A Matt-Whitney test.
3. A linear regression with session-level clusters.

In the case of the round-level data (strategy 3), we included round dummies.

Tests 1 and 3 are parametric, allowing us to conduct ex ante power calculations. We did not conduct pilots designed to inform power calculations, and so we targeted an equal number of observations across treatment pairs, and then conducted power calculations after conventional inference to evaluate the design, rather than to assist in the design.

The theory predicts the sign but not the magnitude of the treatment effect. Therefore, in our power calculations, we arbitrarily use a 20% treatment effect as our benchmark.

For the unpaired t-test, we used the following formula for power calculations:

$$\hat{\rho}_{\mu} = \Pr\left(t_{m_0+m_1-2} > t^{2.5\%}_{m_0+m_1-2} - \frac{0.2}{\sqrt{\frac{\hat{\sigma}_0^2}{m_0} + \frac{\hat{\sigma}_1^2}{m_1}}} \right)$$

where (m_0, m_1) are the sample sizes of the control and treatment, respectively, $(\hat{\sigma}_0, \hat{\sigma}_1)$ are the respective sample standard deviations, and $t^{2.5\%}$ is the 2.5% critical value from a t-table.

For the clustered regressions, we used the following formula:

$$\hat{\rho}_{\beta} = \Pr\left(t_{dof} > t^{2.5\%}_{dof} - \frac{0.2}{\sigma_{\hat{\beta}}} \right)$$

Where dof denotes degrees of freedom, and $\sigma_{\hat{\beta}}$ is the standard error of the regression coefficient of the treatment dummy.

2.2.3. Results

The primary results are shown in Table 2.2.3.1. They paint a highly homogenous picture that strongly supports the main hypothesis in the paper.

All treatment effects are roughly comparable to those reported in the descriptive statistics (Table 2.1.2), meaning that they are economically large: 17% and over in all cases. The estimated treatment effects are also all statistically significant at conventional levels and beyond. Note that tests 9 and 12 correspond to the treatment effects reported and tested in Table 2 of the main paper.

We also estimate the power of each test in detecting a 20% treatment effect, using the standard errors estimated from the sample. All exceed 92%, except for the round 1 data for alternating offers, where the regression implied a power of 79%.

The reason that the p-values are so small and the power is so high even in the session-level tests (2, 3, 6, 7, 9, 10, 12, 13), which involve as few as 8 observations, is that the between-session variation (across treatments) is extremely large compared to the within-session variation, i.e., the estimated standard errors are very small.

In results that we omitted in the interests of parsimony, we note that the coefficients on the round dummies in the round-level clustered regressions are all small in magnitude and statistically insignificant. This is likely because the four practice rounds allowed the subjects to refine their strategy sufficiently prior to the real rounds.

Table 2.2.3.1: Primary Statistical Tests: Unweighted (Top) and Weighted (Bottom)

Test	Bargaining	Data Type	Test type	# Obs	Treat. Effect	P-Value	Power
1	Alt. Offer	Period 1 only	Clustered reg.	83	26%	0.003	0.79
2	Alt. Offer	Session-level	t-test	10	29%	0.002	0.98
3	Alt. Offer	Session-level	MW test	10	-	0.011	-
4	Alt. Offer	Period-level	Clustered reg.	60	29%	< 0.001	0.98
5	Open Chat	Period 1 only	Clustered reg.	66	17%	0.003	0.99
6	Open Chat	Session-level	t-test	8	25%	0.001	0.99
7	Open Chat	Session-level	MW test	8	-	0.021	-
8	Open Chat	Period-level	Clustered reg.	48	26%	< 0.001	0.99

Test	Bargaining	Data Type	Test type	# Obs	Treat. Effect	P-Value	Power
9	Alt. Offer	Session-level	t-test	10	19%	0.013	0.96
10	Alt. Offer	Session-level	MW test	10	-	0.011	-
11	Alt. Offer	Period-level	Clustered reg.	60	19%	0.002	0.97
12	Open Chat	Session-level	t-test	8	17%	0.014	0.97
13	Open Chat	Session-level	MW test	8	-	0.021	-
14	Open Chat	Period-level	Clustered reg.	48	19%	0.006	0.92

Round-level models include round dummies. Power refers to estimated power of detecting a 20% treatment effect.

3. SUMMARY

The data strongly support the main hypothesis, both in terms of simple comparisons of the unconditional means, and in terms of a wide range of parametric and non-parametric tests that tackle the dependence in the data in a variety of ways.

REFERENCES

Charness, Gary. 2000. "Self-Serving Cheap Talk: A Test of Aumann's Conjecture." *Games and Economic Behavior*, 33(2), 177-94.

Charness, Gary and Martin Dufwenberg. 2006. "Promises and Partnership." *Econometrica*, 74(6), 1579-601.

Fischbacher, U. 2007. "Z-Tree – Zurich Toolbox for Readymade Economic Experiments." *Experimental Economics*, 10, 171-78.

APPENDIX: EXPERIMENTAL MATERIALS

EXPERIMENTER SCRIPT

Announcements appear in quotes

1. Turn on computers. Prepare z-Tree code and leaves (start fewer leaves than needed. Others can be opened if required.) Distribute documents (explanatory statement, consent form, and receipt) and pens to cabins.
2. Prepare cabin cards. When expecting 24 students order cards so as to maximize physical space between subjects that stand in line next to each other (i.e., 1, 8, 2, 9, 3, 10 etc.)
3. Ask the students to form line.
4. Check ids, mark attendees, assign to cabin using cabin cards.
5. On the hour, post signs that "Experiment is in Progress" and lock door.
6. "Thank you for coming to today's experiment. You have so far received a total of three documents. One of them is a receipt – please set this aside for now. I will ask you to fill it out at the end of today's session. The explanatory statement is for your information, and you are welcome to keep that after today's session. Please read the consent form, tick the various boxes, sign, and date it. Once I collect all the consent forms we will be ready to start today's experiment. Today's date is ______. I will give you a few minutes to do so and then I will come around to collect it. If you have any questions please raise your hand and I will come to you."
7. Give students sufficient time and then collect consent forms.
8. "We are now ready to begin today's session. I will pass around the instructions. Once I distribute them to everyone I will read them out loud. I do this for two reasons. First, to make sure that everyone has identical instructions. Second, to give you sufficient opportunity to understand these instructions. Please note that the instructions are printed on both sides of the document."
9. Distribute the instructions.
10. In 1.3 and 1.4 after reading the "Summary of key points" I return to the second table of the instructions and ask subjects to reflect on the various points listed in the summary:
    a. I re-read the information in the table
    b. We make clear that one will never know their trade partner's values or costs. Each trader knows their own values or costs not their partners.
    c. Present the following examples:
        i. Assume that you are a seller with a cost of C (0 or 160). Assume that your partner has value V (40 or 200) and that you settle on a price P in stage 1. How much money do you make? (P-C). How much money does your partner make? (V-P).
        ii. Assume that you are a seller with a cost of C (0 or 60). Assume that your partner has value V (40 or 200) and that you settle on a price P in stage 3. How much money do you make? 0.8(P-C) How much money does your partner make? 0.8 (V-P).
11. In 2.1 and 2.2 after reading the "Summary of key points" I return to the second table of the instructions and ask subjects to reflect on the various points listed in the summary:
    a. I re-read the information in the table

b. We make clear that one will never know their trade partner's values or costs for any one unit. Each trader knows their own values or costs not their partners.
c. Present the following examples:
    i. Assume that you agree to exchange a combination of goods that have a total cost C and a total value of V for a price P in stage 1. How much money do you make? (P-C). How much money does your partner make? (V-P).
    ii. Assume that you agree to exchange a combination of goods that have a total cost C and a total value of V for a price P in, say, stage 3. How much money do you make? 0.8(P-C). How much money does your partner make? 0.8(V-P).

12. Calibrate z-Tree code with the correct number of participants.
13. Run z-Tree code.
14. "We have now concluded the experiment. Please enter the total amount that you earned on your receipt. Once you have completed your receipts I will administer a short questionnaire and then call you to come and receive payment. Please remain seated until your cabin number is called."
15. "Has everyone completed the receipts? If yes, I will now run the questionnaire. Please note that once I begin running the questionnaire your total payoff will disappear from your screen. Has everyone completed their receipts?"
16. As I call your cabin please come forward and bring your
    a. Receipts and
    b. Instructions
    c. The laminated cabin number card
    d. The pen

    You can keep the explanatory notes.

INSTRUCTIONS: SINGLE GOOD, ALTERNATING OFFER[2]

Welcome and Introduction

Welcome to our study of decision making. If you read these instructions carefully and make good decisions, you can earn a considerable amount of money.

Kindly refrain from talking with other participants during the session. Also, it is an important requirement of this experiment that you **please turn your mobile phones to silent and abstain from using them during the experiment**.

By coming to this session, you have earned an attendance fee of $10. This has two components. A show up fee of $3 and a participation fee of $7. The participation fee is discussed further below.

The $10 attendance fee will be in addition to any amount that you earn based on your decisions. Once your earnings are determined you will be paid privately and your participation in this experiment will then be concluded.

[2] This title is, of course, not part of the instructions that were distributed to experimental participants which begin with the line "Welcome and Introduction" below.

If you have a question at any time, please raise your hand and I will approach you so that you can ask your question in private.

Setup

Today, we are going to set up a market in which some of you will be buyers and some of you will be sellers. In this market you will be given the opportunity to trade a commodity. We will not specify a name for the commodity; we will simply refer to it as a "unit".

Trading will occur in a sequence of trading rounds. The prices that you negotiate in each round will determine your earnings.

The experiment will consist of 10 rounds: 4 practice rounds followed by 6 real rounds.

The first 4 rounds will be practice and will not affect your earnings for the experiment.

The final 6 rounds will be real and will affect your earnings. At the end of the experiment, the computer will select one of the 6 real rounds at random and you will be paid based on your earnings in that round in cash.

Matching Rules

In every round, the computer will tell you whether you are a buyer or a seller for that round. During the practice rounds, you will experience both the role of a buyer as well as the role of a seller. Once we have completed the practice rounds, you will be assigned the role of either buyer or seller and will remain in that role throughout the remainder of the session. You have a 50% chance of being a buyer, and a 50% chance of being a seller.

Each round, every buyer will be randomly matched with a seller. That means that each buyer is equally likely to be matched with each seller. The matching is independent every round, which means that being matched with a specific trader in one round has no effect on the likelihood of being matched with the same trader in a future round.

All matching is anonymous, meaning that you will never know the identity of whom you are matched with in any round.

Participation fee

At the start of each round, you will be given a $7 participation fee. Any positive earnings that you may make by trading the "unit" will be added to this $7, whereas any negative earnings will be subtracted from this $7. If you do not end up trading a unit, then you just keep your $7 participation fee.

Profit from trading

In each round, sellers and buyers will have the opportunity to exchange their unit.

Prior to the start of each round, sellers will be given a number known as their "cost" and buyers will be given a number known as their "value". The cost represents the minimum amount for which a seller can sell a unit without making a loss. The value represents the maximum amount for which a buyer can purchase a unit without making a loss.

Sellers earn money by selling a unit at a price that is above their cost. Seller earnings from the sale of a unit are the difference between the sale price and the cost. For example, if a seller has a cost of $100 and sells their unit for $140, the seller earns $140 – $100 = $40.

Buyers earn money by buying a unit at a price that is below their value. Buyer earnings from the purchase of a unit are the difference between the value and the purchase price. For example, if a buyer has a value of $150 and buys a unit for $100, the buyer earns $150 – $100 = $50.

If a seller sells a unit at a price that is less than their cost, they will make a loss. If a buyer buys a unit at a price that is greater than their value, they will make a loss. If you are at risk of making a loss, the computer will notify you and ask you to confirm.

If the seller and the buyer do not exchange the unit, they each earn a profit of $0 for that round.

How the seller's cost and the buyer's value for a unit are determined is described in the next section.

How are Costs and Values Determined?

Each round, for each seller, there is:

- a 50% probability that their cost will be $0
- a 50% probability that their cost will be $160

Each round, for each buyer, there is:

- a 50% probability that their value will be $40
- a 50% probability that their value will be $200

Given the probabilities above, there are 4 possible ways that a seller's cost will match with a buyer's value:

	Probability	Seller's cost	Buyer's value
a.	25%	0	40
b.	25%	0	200
c.	25%	160	40
d.	25%	160	200

Traders' costs/values are determined independently, meaning that knowing an individual trader's cost/value during a round tells you nothing about the cost/value of any other trader in that round, or any other round.

Sellers are only informed of their cost, and buyers are only informed of their value, meaning that neither side knows about the other.

So, if you are a seller and your cost is, say, $0 you will not know whether the buyer that you are matched with has a value of $40 or $200.

Similarly, if you are a buyer and your value is, say, $200 you will not know whether the seller that you are matched with has a cost of $0 or $160.

Naturally, you will never be compelled to either buy or sell a "unit". If a given pair of seller and buyer do not reach agreement on a price for a "unit" after 8 bargaining stages, bargaining is suspended. In this case each trader only receives their $7 participation fee.

A screenshot below shows a trader discovering that they are a seller, and learning their cost:

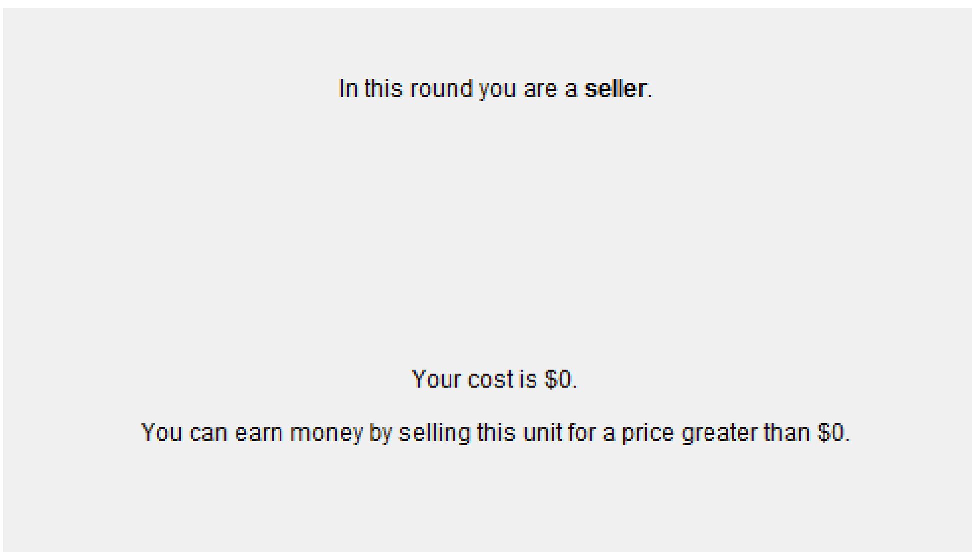

A screenshot below shows a trader discovering that they are a buyer, and learning their value:

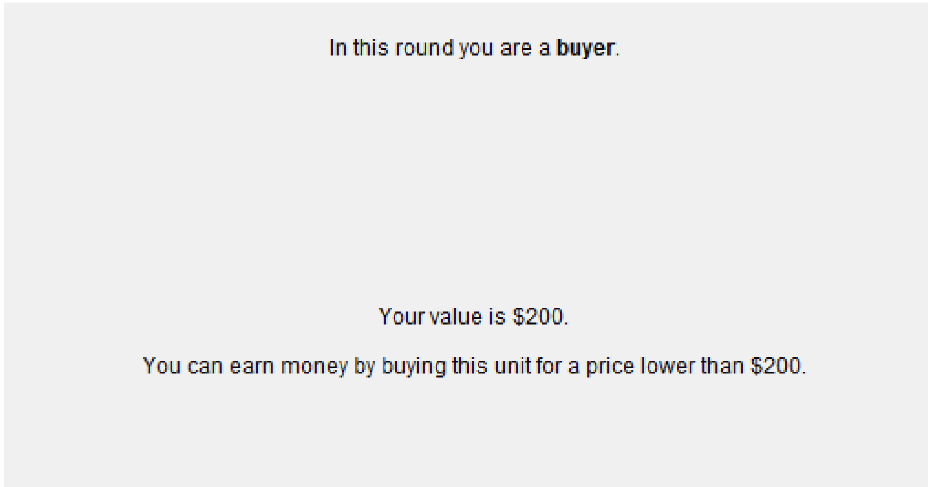

How Much Money Can You and Your Partner Make?

When you and your partner are bargaining over a unit, the total money that you can make between you from successfully trading that unit is equal to the buyer's value minus the seller's cost. This follows from the way in which buyer and seller profits are calculated.

To understand this, we will examine one possible example. Consider a buyer with value $200 and a seller with cost $160. If they agree on a price $175 then the buyer will make $200 - $175 = $25 and the seller $175 - $160 = $15. Together they make $25 + $15 = $40. Alternatively, they could agree on a price $190 and then the buyer will make $200 - $190 = $10 and the seller $190 - $160 = $30. Again, together they make $10 + $30 = $40. Clearly, the total profit is the difference between the cost and the value: $200 - $160 = $40. The trade price merely determines how the total profit of $40 is divided between the buyer and seller.

Since in every round, each of the following 4 possibilities is equally likely, there are 4 possibilities regarding the total amount of profit that you and your partner can make:

	Probability	Seller's cost	Buyer's value	Combined buyer and seller profit if trade occurs	Combined buyer and seller profit if no trade occurs
a.	25%	0	40	$40 - $0 = $40	$0
b.	25%	0	200	$200 - $0 = $200	$0
c.	25%	160	40	$40 - $160 = -$120	$0
d.	25%	160	200	$200 - $160 = $40	$0

Any amount that you and your partner make is shared between you on the basis of the trading agreement that you make with your partner, which is done according to the mechanism described in the next section.

Trading Mechanism

After being matched with another trader and seeing your cost/value, you will have the opportunity to trade with your partner in a sequence of 8 bargaining stages. Bargaining stages consist of you and your partner making offers of trading prices for the unit. The computer will select one of you at random to make the first offer to their partner.

If the first offer is accepted, the round is over.

If the first offer is rejected, the other trader has the opportunity to make a counter-offer.

This alternating process continues until one of the pair of traders makes an offer that the other trader accepts.

If your partner makes you an offer, you can either accept it or you can reject it and counter-offer.

If you are randomly selected to make an offer in the first bargaining stage, you will have 30 seconds to complete your offer and submit it to your partner.

If you do not submit an offer within this 30 second window, bargaining advances to the next stage, and the right to make an offer is given to your partner.

If you are randomly selected to respond to an offer in the first bargaining stage, you will have 10 seconds to accept or reject your partner's offer.

If you accept their offer, bargaining ends.

If you reject their offer, then you are given the opportunity to make a counter-offer.

If you do not respond to your partner (accept or reject their offer) within this 10 second window, your inaction counts as a rejection but you are not given the right to make a counter-offer. Instead, bargaining advances to the next stage, and the right to make an offer reverts to your partner.

The procedure that we describe above for bargaining stage 1 works the same way for all subsequent bargaining stages (stages 2-8) but with one key difference. As we note above, in stage 1 a proposer is given 30 seconds and a responder is given 10 seconds. By contrast, in stages 2-8 a proposer is given 15 seconds and a responder is given 5 seconds.

If a given pair of seller and buyer do not reach agreement on a price for a "unit" after 8 bargaining stages, bargaining is suspended. In this case each trader only receives their $7 participation fee.

If it is your turn to make an offer during the bargaining stage, you can choose any price from $0 to $200, inclusive. The offer must be a whole number, so no fractions of dollars are allowed.

The screenshot below shows you a buyer considering what offer to make:

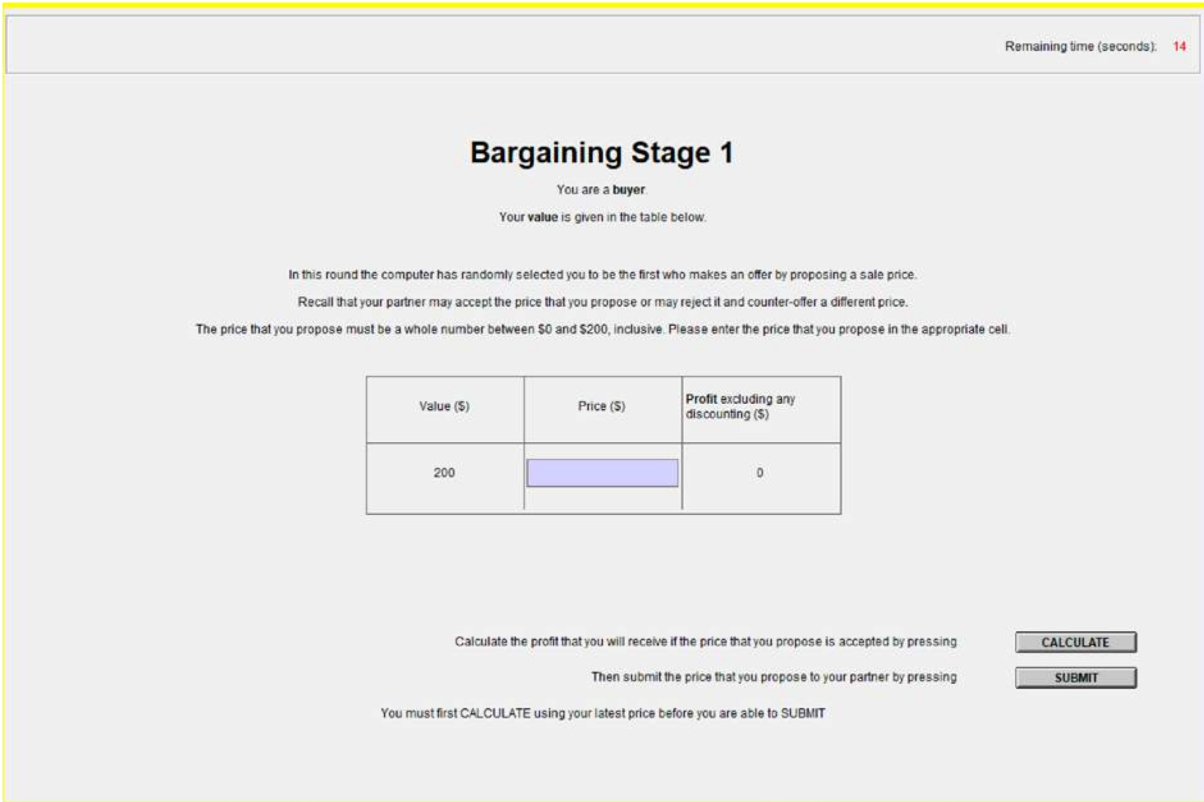

Once you have finished composing your offer, press the CALCULATE button to calculate how much profit you will make if your offer is accepted and then press the SUBMIT button to submit the offer to your trading partner.

Making a Loss and Exhausting Your $7

You will never be compelled to make a loss. If you make a decision that puts you at risk of making a loss, the computer will warn you with a pop-up message and ask for confirmation. The pop-up window will have two buttons: “OK” and “CANCEL”. Press **“OK” to accept to make a loss** or **“CANCEL” to revise your decision.** If you agree to make a loss up to $7, this loss will simply be subtracted from your participation fee of $7. If you agree to make a loss that is greater than your $7 participation fee for that round, your participation in the experiment for all subsequent rounds will be suspended, and you will earn nothing beyond your show-up fee of $3, which you will receive after the conclusion of all the rounds in the experiment.

Time Limits and Trading Reductions

Each bargaining round is split into 8 stages. Once the 8 stages are complete, the round is over, and if you and your partner failed to accept an offer, each of you earns no profit from trading in that round.

Stage 1 lasts 40 seconds (30 seconds for the proposer, and 10 seconds for the responder). Each of stages 2 to 8 last 20 seconds (15 seconds for the proposer and 5 seconds for the responder). The top of the screen will indicate to you what stage you are in, and how much time you have left in that stage. Once a stage is complete, the next one will start immediately, until the last (8^{th}) stage is complete.

If you or your partner accept an offer during the first bargaining stage, your earnings for that round will be equal to the numbers shown on the screen.

If you or your partner accept an offer during the second bargaining stage or later, there will be a reduction on the profit that both players receive from trading (this reduction does not apply to the $7 participation fee, or the $3 show-up fee). This reduction will take the form of a percentage that will be deducted from your earnings and those of your partner. The table below specifies the reductions.

Bargaining stage	% of Your Earnings that You Lose
1	0%
2	10%
3	20%
4	30%
5	40%
6	50%
7	60%
8	70%
End of round	Round canceled for you and your partner (and you both earn no profit from trading)

For example, if it took until bargaining stage 2 for one of you to accept an offer, then each of you will lose 10% of your earnings.

Sequence of events

In each round, the sequence of decisions is as follows.

1. The computer randomly matches you with another trader.
2. You are randomly assigned your cost/value:

- 50% chance of a cost of $0 and 50% chance of a cost of $160 if you are a seller.
- 50% chance of a value of $40 and 50% chance of a value of $200 if you are a buyer.

3. The computer randomly selects one member of each pair of traders to make the first offer.
4. That trader makes an offer.
5. The other trader can accept or reject.
6. If the trader rejects, it becomes that trader's turn to make an offer.
7. The opportunity to make an offer keeps alternating until:
    a. one trader accepts an offer
    b. or 8 offers are rejected, which will suspend bargaining in that round
8. Earnings for that round are calculated, and a reduction may be applied depending on which bargaining stage the acceptance came in.
9. At the conclusion of the last real round, one of the 6 real rounds is selected at random and used to pay the participants.

Summary of Key Points

By way of providing a summary of the key points, we suggest that you return to the table of page 6[3] and consider:

1. What information a buyer/seller will have in each round
2. What are the possible combinations of costs and values
3. How, given a particular combination of cost and value, a price determined in a certain bargaining stage, that may involve a reduction, determines the profit that a buyer and a seller will make

INSTRUCTIONS: SINGLE GOOD, OPEN CHAT[4]

Welcome and Introduction

Welcome to our study of decision making. If you read these instructions carefully and make good decisions, you can earn a considerable amount of money.

Kindly refrain from talking with other participants during the session. Also, it is an important requirement of this experiment that you **please turn your mobile phones to silent and abstain from using them during the experiment**.

[3] The pages have been renumbered in this supplement. This is the second table that appears in the instructions for this treatment.

[4] This title is, of course, not part of the instructions that were distributed to experimental participants which begin with the line "Welcome and Introduction" below.

By coming to this session, you have earned an attendance fee of $10. This has two components. A show up fee of $3 and a participation fee of $7. The participation fee is discussed further below.

The $10 attendance fee will be in addition to any amount that you earn based on your decisions. Once your earnings are determined you will be paid privately and your participation in this experiment will then be concluded.

If you have a question at any time, please raise your hand and I will approach you so that you can ask your question in private.

Setup

Today, we are going to set up a market in which some of you will be buyers and some of you will be sellers. In this market you will be given the opportunity to trade a commodity. We will not specify a name for the commodity; we will simply refer to it as a “unit”.

Trading will occur in a sequence of trading rounds. The prices that you negotiate in each round will determine your earnings.

The experiment will consist of 10 rounds: 4 practice rounds followed by 6 real rounds.

The first 4 rounds will be practice and will not affect your earnings for the experiment.

The final 6 rounds will be real and will affect your earnings. At the end of the experiment, the computer will select one of the 6 real rounds at random and you will be paid based on your earnings in that round in cash.

Matching Rules

In every round, the computer will tell you whether you are a buyer or a seller for that round. During the practice rounds, you will experience both the role of a buyer as well as the role of a seller. Once we have completed the practice rounds, you will be assigned the role of either buyer or seller and will remain in that role throughout the remainder of the session. You have a 50% chance of being a buyer, and a 50% chance of being a seller.

Each round, every buyer will be randomly matched with a seller. That means that each buyer is equally likely to be matched with each seller. The matching is independent every round, which means that being matched with a specific trader in one round has no effect on the likelihood of being matched with the same trader in a future round.

All matching is anonymous, meaning that you will never know the identity of whom you are matched with in any round.

Participation fee

At the start of each round, you will be given a $7 participation fee. Any positive earnings that you may make by trading the "unit" will be added to this $7, whereas any negative earnings will be subtracted from this $7. If you do not end up trading a unit, then you just keep your $7 participation fee.

Profit from trading

In each round, sellers and buyers will have the opportunity to exchange their unit.

Prior to the start of each round, sellers will be given a number known as their "cost" and buyers will be given a number known as their "value". The cost represents the minimum amount for which a seller can sell a unit without making a loss. The value represents the maximum amount for which a buyer can purchase a unit without making a loss.

Sellers earn money by selling a unit at a price that is above their cost. Seller earnings from the sale of a unit are the difference between the sale price and the cost. For example, if a seller has a cost of $100 and sells their unit for $140, the seller earns $140 – $100 = $40.

Buyers earn money by buying a unit at a price that is below their value. Buyer earnings from the purchase of a unit are the difference between the value and the purchase price. For example, if a buyer has a value of $150 and buys a unit for $100, the buyer earns $150 – $100 = $50.

If a seller sells a unit at a price that is less than their cost, they will make a loss. If a buyer buys a unit at a price that is greater than their value, they will make a loss. If you are at risk of making a loss, the computer will notify you and ask you to confirm.

If the seller and the buyer do not exchange the unit, they each earn a profit of $0 for that round.

How the seller's cost and the buyer's value for a unit are determined is described in the next section.

How are Costs and Values Determined?

Each round, for each seller, there is:

- a 50% probability that their cost will be $0
- a 50% probability that their cost will be $160

Each round, for each buyer, there is:

- a 50% probability that their value will be $40
- a 50% probability that their value will be $200

Given the probabilities above, there are 4 possible ways that a seller's cost will match with a buyer's value:

	Probability	Seller's cost	Buyer's value
a.	25%	0	40
b.	25%	0	200
c.	25%	160	40
d.	25%	160	200

Traders' costs/values are determined independently, meaning that knowing an individual trader's cost/value during a round tells you nothing about the cost/value of any other trader in that round, or any other round.

Sellers are only informed of their cost, and buyers are only informed of their value, meaning that neither side knows about the other.

So, if you are a seller and your cost is, say, $0 you will not know whether the buyer that you are matched with has a value of $40 or $200.

Similarly, if you are a buyer and your value is, say, $200 you will not know whether the seller that you are matched with has a cost of $0 or $160.

Naturally, you will never be compelled to either buy or sell a "unit". If a given pair of seller and buyer do not reach agreement on a price for a "unit" after 8 bargaining stages, bargaining is suspended. In this case each trader only receives their $7 participation fee.

A screenshot below shows a trader discovering that they are a seller, and learning their cost:

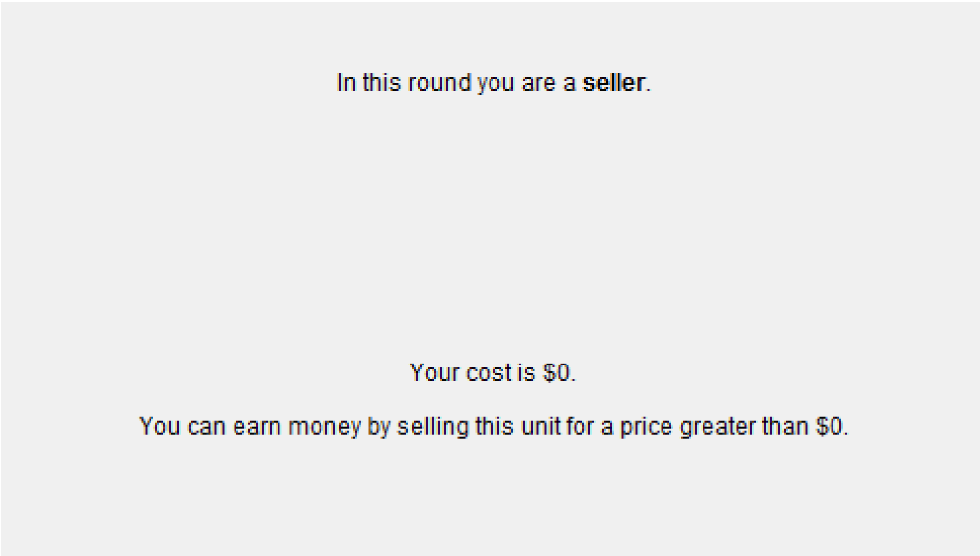

A screenshot below shows a trader discovering that they are a buyer, and learning their value:

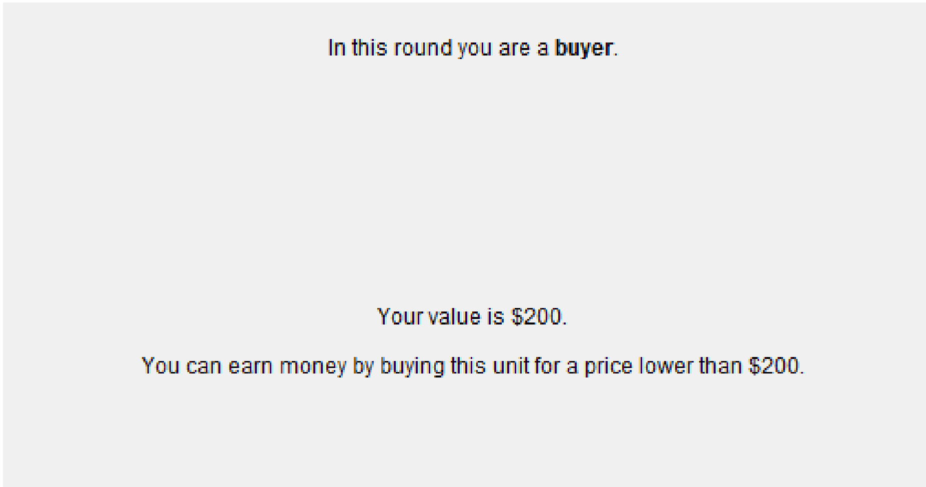

How Much Money Can You and Your Partner Make

When you and your partner are bargaining over a unit, the total money that you can make between you from successfully trading that unit is equal to the buyer’s value minus the seller’s cost. This follows from the way in which buyer and seller profits are calculated.

To understand this, we will examine one possible example. Consider a buyer with value $200 and a seller with cost $160. If they agree on a price $175 then the buyer will make $200 - $175 = $25 and the seller $175 - $160 = $15. Together they make $25 + $15 = $40. Alternatively, they could agree on a price $190 and then the buyer will make $200 - $190 = $10 and the seller $190 - $160 = $30. Again, together they make $10 + $30 = $40. Clearly, the total profit is the difference between the cost and the value: $200 - $160 = $40. The trade price merely determines how the total profit of $40 is divided between the buyer and seller.

Since in every round, each of the following 4 possibilities is equally likely, there are 4 possibilities regarding the total amount of profit that you and your partner can make:

	Probability	Seller’s cost	Buyer’s value	Combined buyer and seller profit if trade occurs	Combined buyer and seller profit if no trade occurs
a.	25%	0	40	$40 - $0 = $40	$0
b.	25%	0	200	$200 - $0 = $200	$0
c.	25%	160	40	$40 - $160 = -$120	$0
d.	25%	160	200	$200 - $160 = $40	$0

Any amount that you and your partner make is shared between you on the basis of the trading agreement that you make with your partner, which is done according to the mechanism described in the next section.

Trading Mechanism

After being matched with another trader and seeing your cost/value, you will have the opportunity to trade with your partner in a sequence of 8 bargaining stages. Bargaining stages consist of you and your partner making offers of trading prices for the unit.

During each bargaining stage, you and your partner are free to make offers at any time. Each of you is also free to accept an offer made by the other at any time. Once an offer is accepted, the round is over. You and your partner also have a time limit, which we will explain below.

When making an offer, you can choose any price from $0 to $200, inclusive. The offer must be a whole number, so no fractions of dollars are allowed.

The screenshot below shows you a buyer considering what offer to make:

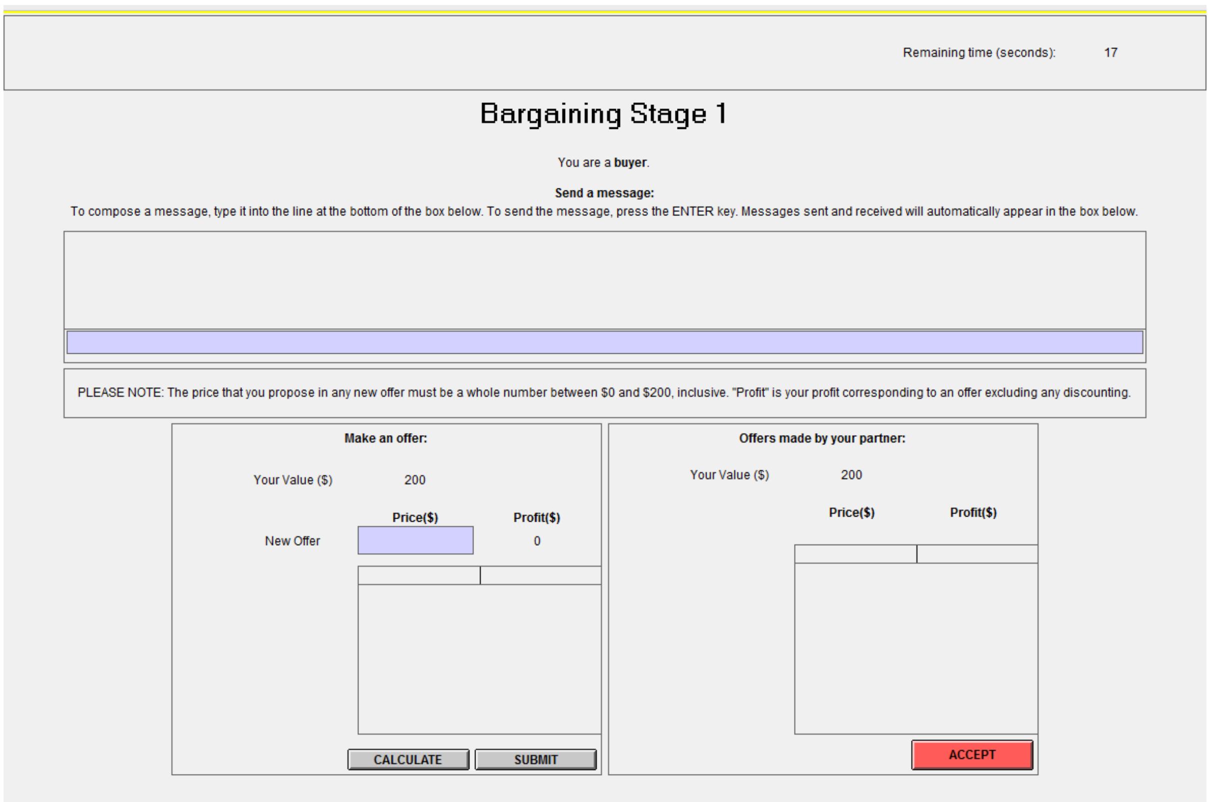

Once you have finished composing your offer, press the CALCULATE button to calculate how much profit you will make if your offer is accepted and then press the SUBMIT button to submit the offer to your trading partner.

The offers that you make will appear on the left side of the screen. The most recent offer will be on top, and the older ones will be at the bottom.

Your partner is free to accept any of your offers at any point during the round. If your partner accepts an offer, the round ends immediately.

While you are making offers, your partner can also make offers. **Both of you can make as many offers as you like without waiting for the other person to respond.** Your partner's offers will appear on the right side of your screen, with the most recent ones at the top, and the older ones at the bottom.

The computer will automatically display the profit that you will earn from each offer if you accept it.

During a round, you and your partner will be able to chat with each other via a chat window. You can see this chat window at the top of the screenshot that appears on the previous page. You can send messages at any time, and can send as many messages as you like. However, we ask that you please observe the following rules:

- No threatening or abusive language
- No messages that reveal your identity, such as your name or your location in the lab.

Accepting an Offer

To accept an offer, highlight the offer by moving the mouse cursor over it, and then click on ACCEPT. The round will end immediately as soon as either you or your partner accepts an offer.

Making a Loss and Exhausting Your $7

You will never be compelled to make a loss. If you make a decision that puts you at risk of making a loss, the computer will warn you with a pop-up message and ask for confirmation. The pop-up window will have two buttons: "OK" and "CANCEL". Press **"OK" to accept to make a loss** or **"CANCEL" to revise your decision.** If you agree to make a loss up to $7, this loss will simply be subtracted from your participation fee of $7. If you agree to make a loss that is greater than your $7 participation fee for that round, your participation in the experiment for all subsequent rounds will be suspended, and you will earn nothing beyond your show-up fee of $3, which you will receive after the conclusion of all the rounds in the experiment.

Time Limits and Trading Reductions

Each bargaining round is split into 8 stages. Once the 8 stages are complete, the round is over, and if you and your partner failed to accept an offer, each of you earns no profit from trading in that round.

Stage 1 lasts 40 seconds, while each of stages 2 to 8 last 20 seconds. The top of the screen will indicate to you what stage you are in, and how much time is left until the end of the stage. Once a stage is complete, the next one will start immediately, until the last (8^{th}) stage is complete.

Offers do not expire and an offer made in one stage can be accepted in that stage or in any subsequent stage.

If you or your partner accept an offer during the first bargaining stage, your earnings for that round will be equal to the numbers shown on the screen.

If you or your partner accept an offer during the second bargaining stage or later, there will be a reduction on the profit that both players receive from trading (this reduction does not apply to the $7 participation fee, or the $3 show-up fee). This reduction will take the form of a percentage that will be deducted from your earnings and those of your partner. The table below specifies the reductions.

Bargaining stage	% of Your Earnings that You Lose
1	0%
2	10%
3	20%
4	30%
5	40%
6	50%
7	60%
8	70%
End of round	Round canceled for you and your partner (and you both earn no profit from trading)

For example, if it took until bargaining stage 2 for one of you to accept an offer, then each of you will lose 10% of your earnings.

Sequence of events

In each round, the sequence of decisions is as follows.

1. The computer randomly matches you with another trader.
2. You are randomly assigned your cost/value:
   - 50% chance of a cost of $0 and 50% chance of a cost of $160 if you are a seller.
   - 50% chance of a value of $40 and 50% chance of a value of $200 if you are a buyer.
3. Bargaining commences; you and your partner are free to make offers, accept offers, and chat.

4. Bargaining continues until:
    a. one trader accepts an offer
    b. or 8 bargaining stages are concluded, which will suspend bargaining in that round
5. Earnings for that round are calculated, and a reduction may be applied depending on which bargaining stage the acceptance came in.
6. At the conclusion of the last real round, one of the 6 real rounds is selected at random and used to pay the participants.

Summary of Key Points

By way of providing a summary of the key points, we suggest that you return to the table of page 6[5] and consider:

1. What information a buyer/seller will have in each round
2. What are the possible combinations of costs and values
3. How, given a particular combination of cost and value, a price determined in a certain bargaining stage, that may involve a reduction, determines the profit that a buyer and a seller will make

INSTRUCTIONS: FOUR GOODS, ALTERNATING OFFER[6]

Welcome and Introduction

Welcome to our study of decision making. If you read these instructions carefully and make good decisions, you can earn a considerable amount of money.

Kindly refrain from talking with other participants during the session. Also, it is an important requirement of this experiment that you **please turn your mobile phones to silent and abstain from using them during the experiment**.

By coming to this session, you have earned an attendance fee of $10. This has two components. A show up fee of $3 and a participation fee of $7. The participation fee is discussed further below.

The $10 attendance fee will be in addition to any amount that you earn based on your decisions. Once your earnings are determined you will be paid privately and your participation in this experiment will then be concluded.

[5] The pages have been renumbered in this supplement. This is the second table that appears in the instructions for this treatment.
[6] This title is, of course, not part of the instructions that were distributed to experimental participants which begin with the line "Welcome and Introduction" below.

If you have a question at any time, please raise your hand and I will approach you so that you can ask your question in private.

Setup

Today, we are going to set up a market in which some of you will be buyers and some of you will be sellers. In this market you will be given the opportunity to trade commodities. We will not specify names for these commodities; we will simply refer to them as “units”.

Trading will occur in a sequence of trading rounds. The decisions that you make in each round will determine your earnings.

The experiment will consist of 10 rounds: 4 practice rounds followed by 6 real rounds.

The first 4 rounds will be practice and will not affect your earnings for the experiment.

The final 6 rounds will be real and will affect your earnings. At the end of the experiment, the computer will select one of the 6 real rounds at random and you will be paid based on your earnings in that round in cash.

Matching Rules

In every round, the computer will tell you whether you are a buyer or a seller for that round. During the practice rounds, you will experience both the role of a buyer as well as the role of a seller. Once we have completed the practice rounds, you will be assigned the role of either buyer or seller and will remain in that role throughout the remainder of the session. You have a 50% chance of being a buyer, and a 50% chance of being a seller.

Each round, every buyer will be randomly matched with a seller. That means that each buyer is equally likely to be matched with each seller. The matching is independent every round, which means that being matched with a specific trader in one round has no effect on the likelihood of being matched with the same trader in a future round.

All matching is anonymous, meaning that you will never know the identity of whom you are matched with in any round.

Participation fee

At the start of each round, you will be given a $7 participation fee. Any positive earnings that you may make by trading one or more “units” will be added to this $7, whereas any negative earnings will be subtracted from this $7. If you do not end up trading one or more units, then you just keep your $7 participation fee.

Profit from trading

In each round, sellers and buyers will have the opportunity to exchange up to 4 units.

For each unit, sellers will be given a number known as their "cost" and buyers will be given a number known as their "value". The cost represents the minimum amount for which a seller can sell that unit without making a loss. The value represents the maximum amount for which a buyer can purchase that unit without making a loss.

Sellers and buyers exchange units in bundles that include 1 to 4 units for one total price paid for the entire bundle.

Sellers earn money when they sell a bundle at a total price that is above the total cost of the units inside that bundle. Seller earnings from the sale of a bundle are the difference between the price of the entire bundle and the total of the costs of the units included in the bundle. For example, if a seller has a unit with cost of $20 and a unit with cost of $80, and sells a bundle that includes both units for a total price of $140, then the total cost of the units in the bundle is $20 + $80 = $100, and the seller earns $140 – $100 = $40.

Buyers earn money when they buy a bundle at a total price that is below the total value of the units inside that bundle. Buyer earnings from the purchase of a bundle are the difference between the total of the values of the units included in the bundle and the price of the entire bundle. For example, if a buyer's value for one unit is $90 and for another unit is $60, and buys a bundle that includes both units for a total price of $100, then the total value of the units in the bundle is $90 + $60 = $150, and the buyer earns $150 – $100 = $50.

EXAMPLE:

Consider a case in which a seller and a buyer have the following costs and values for the 4 units:

	Unit 1	Unit 2	Unit 3	Unit 4
Seller's cost ($)	A	B	C	D
Buyer's value ($)	E	F	G	H

The seller's costs of units 1, 2, 3, and 4 are A, B, C, and D, respectively.
The buyer's values of units 1, 2, 3, and 4 are E, F, G, and H, respectively.

Assume that the buyer and seller agree to exchange a bundle that includes only Unit 2 for a price $X. Their profit is calculated as follows:
Seller's profit=X-B
Buyer's profit=F-X

Alternatively, assume that the buyer and seller agree to exchange a bundle that includes Unit 1 and Unit 2 for a price $X. Their profit is calculated as follows:
Seller's profit=X-(A+B)
Buyer's profit=(E+F)-X

The calculations are similar in the case of bundles that include 3 or 4 goods.

If a seller sells a bundle at a price that is less than the total of the costs of the units inside the bundle, this seller will make a loss. If a buyer buys a bundle at a price that is greater than the total of the values of the units inside the bundle, this buyer will make a loss. If you are at risk of making a loss, the computer will notify you and ask you to confirm.

A seller and a buyer **can exchange at most ONLY ONE BUNDLE per round**. If they do not exchange a bundle, they each earn a profit of $0 for that round.

How the seller's cost and the buyer's value of each of the 4 units are determined is described in the next section.

How are Costs and Values Determined?

For each seller-buyer pair, there is a fixed number of possibilities for the combination of costs and values that can occur for each of their four units:

		Seller's cost	Buyer's value
a.	For one of the four units:	$0	$10
b.	For one of the four units:	$0	$50
c.	For one of the four units:	$40	$10
d.	For one of the four units:	$40	$50

Each of these 4 possibilities, (a), (b), (c), and (d) above, MUST OCCUR EXACTLY ONCE in any round. The order with which they occur within a round is random, and each possibility regarding the order in which (a), (b), (c), and (d) occur is equally likely. Hence, knowing this order in one round tells you nothing about how it may manifest in any other round. A list of all the equally likely possibilities is provided at the end of this document.

Sellers are only informed of their cost for each unit, and buyers are only informed of their value for each unit, meaning that neither side knows about the other. So, for a given unit, a seller cannot distinguish between combinations (a) and (b) above, because the seller is given the same information in each case (cost of $0); or between (c) and (d) (cost of $40). Similarly, for a given unit, a buyer cannot distinguish between (a) and (c), because the buyer is given the same information (value of $10) or between (b) and (d) (value of $50).

In other words:

In any round, sellers will have two units each with cost $0 and two units each with cost $40. However, they will not know which of their two $0s meets a buyer's $10 and which meets a buyer's $50. Similarly, they will not know which of their two $40s meets a buyer's $10 and which meets a buyer's $50.

In any round, buyers will have two units each with value $50 and two units each with value $10. However, they will not know which of their two $50s meets a seller's $0 and which meets a seller's $40. Similarly, they will not know which of their two $10's meets a seller's $0 and which meets a seller's $40.

A screenshot below shows a trader discovering that they are a seller, and learning their costs:

In this round you are a **seller**.

The cost of each of your 4 units is as follows:

Unit	1	2	3	4
Cost ($)	0	0	40	40

You can earn money by selling one or more of these units for a total price that is higher than their total cost.

A screenshot below shows a trader discovering that they are a buyer, and learning their values:

In this round you are a **buyer**.

The value of each of your 4 units is as follows:

Unit	1	2	3	4
Value ($)	10	50	10	50

You can earn money by buying one or more of these units for a total price that is lower than their total value.

How Much Money Can You and Your Partner Make?

When you and your partner are bargaining over a bundle of 1 to 4 units, the total money that you can make between you from successfully trading that bundle of 1 to 4 units is equal to the buyer's total value of the traded units minus the seller's total cost of the traded units. This follows from the way in which buyer and seller profits are calculated.

To understand this, we will examine an example with one unit. Consider a buyer with value $50 and a seller with cost $40. If they agree on a price $43 then the buyer will make $50 - $43 = $7 and the seller $43 - $40 = $3. Together they make $7 + $3 = $10. Alternatively, they could agree on a price $48 and then the buyer will make $50 - $48 = $2 and the seller $48 - $40 = $8. Again, together they make $2 + $8 = $10. Clearly, the total profit is the difference between the cost and the value: $50 - $40 = $10. The trade price merely determines how the total profit of $10 is divided between the buyer and seller.

This example was for one unit. When dealing with 4 units, the same principle applies: the total money that you and your partner can make is the total of how much you can make from each unit.

Since in every round, each of the following 4 possibilities occurs exactly once, there is a **fixed maximum total amount of profit** that you and your partner can make every round, and it is equal to $70. To see why, consider the following table:

		Seller's cost	Buyer's value	Combined buyer and seller profit if traded	Combined buyer and seller profit if not traded
a.	For one of the four units:	$0	$10	$10 - $0 = $10	$0
b.	For one of the four units:	$0	$50	$50 - $0=$50	$0
c.	For one of the four units:	$40	$10	$10 - $40 = -$30	$0
d.	For one of the four units:	$40	$50	$50 - $40 = $10	$0

Hence, the **maximum** amount of combined buyer and seller earnings is $10 + $50 + $0 + $10 = $70.

Therefore, depending on the choices that you and your partner make in every round, you can together make a maximum of $70 between you.

Any amount that you and your partner make is shared between you on the basis of the trading agreement that you make with your partner, which is done according to the mechanism described in the next section.

Trading Mechanism

After being matched with another trader and seeing your costs/values, you will have the opportunity to trade with your partner in a sequence of 8 bargaining stages. Bargaining stages consist of you and your partner making proposals of trading prices for bundles of 1 to 4 units. The computer will select one of you at random to make the first proposal to their partner.

If the first proposal is accepted, the round is over.

If the first proposal is rejected, the other trader has the opportunity to make a counter-proposal.

This alternating process continues until one of the pair of traders makes a proposal that the other trader accepts.

If your partner makes you a proposal, you can either accept it or you can reject it and counter-propose.

If you are randomly selected to make a proposal in the first bargaining stage, you will have 30 seconds to complete your proposal and submit it to your partner.

If you do not submit a proposal within this 30 second window, bargaining advances to the next stage, and the right to make a proposal is given to your partner.

If you are randomly selected to respond to a proposal in the first bargaining stage, you will have 10 seconds to accept or reject your partner's proposal.

If you accept their proposal, bargaining ends.

If you reject their proposal, then you are given the opportunity to make a counter-proposal.

If you do not respond to your partner (accept or reject their proposal) within this 10 second window, your inaction counts as a rejection but you are not given the right to make a counter-proposal. Instead, bargaining advances to the next stage, and the right to make a proposal reverts to your partner.

The procedure that we describe above for bargaining stage 1 works the same way for all subsequent bargaining stages (stages 2-8) but with one key difference. As we note above, in stage 1 a proposer is given 30 seconds and a responder is given 10 seconds. By contrast, in stages 2-8 a proposer is given 15 seconds and a responder is given 5 seconds.

Proposal Structure

A proposal consists of 1 to 6 simultaneous offers of bundles to your trading partner. Your trading partner can either accept exactly 1 of those offers, meaning that the proposal is accepted, or they can reject all of them, meaning that the proposal is rejected. Accepting more than 1 offer is not permissible.

Each offer has two components:

1. A list of the units included in the bundle; a bundle must contain at least 1 unit, and at the most 4 units
2. A total price for **all the units** in the bundle, which must be a whole number between $0 and $200, inclusive.

If it is your turn to make a proposal during the bargaining stage, you can choose to make anywhere between 1 and 6 offers.

To make an offer, select the units that will be included in the bundle using the checkboxes, and determine the price. The screenshot below shows you a buyer considering what offer(s) to make.

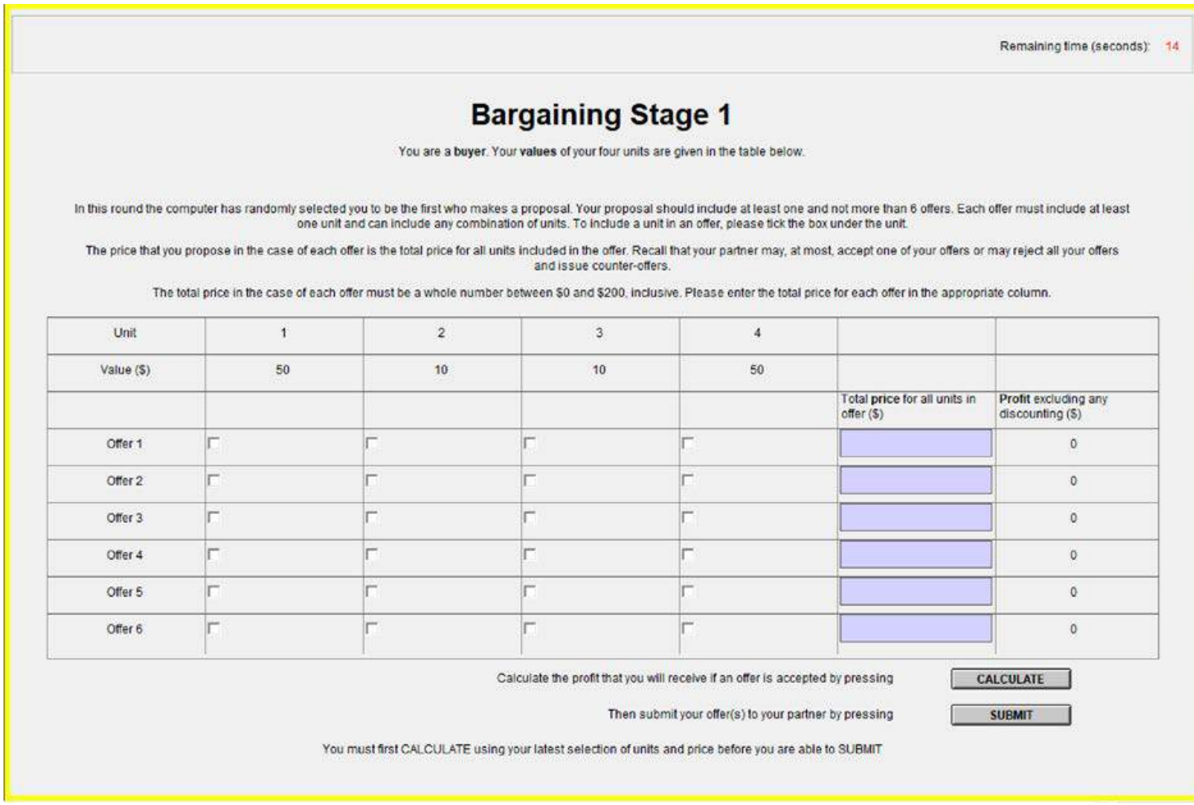

Remaining time (seconds): 14

Bargaining Stage 1

You are a **buyer**. Your **values** of your four units are given in the table below.

In this round the computer has randomly selected you to be the first who makes a proposal. Your proposal should include at least one and not more than 6 offers. Each offer must include at least one unit and can include any combination of units. To include a unit in an offer, please tick the box under the unit.

The price that you propose in the case of each offer is the total price for all units included in the offer. Recall that your partner may, at most, accept one of your offers or may reject all your offers and issue counter-offers.

The total price in the case of each offer must be a whole number between $0 and $200, inclusive. Please enter the total price for each offer in the appropriate column.

Unit	1	2	3	4		
Value ($)	50	10	10	50		
					Total **price** for all units in offer ($)	**Profit** excluding any discounting ($)
Offer 1						0
Offer 2						0
Offer 3						0
Offer 4						0
Offer 5						0
Offer 6						0

Calculate the profit that you will receive if an offer is accepted by pressing CALCULATE

Then submit your offer(s) to your partner by pressing SUBMIT

You must first CALCULATE using your latest selection of units and price before you are able to SUBMIT

Once you have finished composing your proposal, press the CALCULATE button to calculate how much profit you will make if any one of your offers is accepted and then press the SUBMIT button to submit the proposal to your trading partner.

Recall that the responder to a proposal can accept **at most ONLY ONE OFFER**.

Making a Loss and Exhausting Your $7

You will never be compelled to make a loss. If you make a decision that puts you at risk of making a loss, the computer will warn you with a pop-up message and ask for confirmation. The pop-up window will have two buttons: "OK" and "CANCEL". Press **"OK" to accept to make a loss** or **"CANCEL" to revise your decision.** If you agree to make a loss up to $7, this loss will simply be subtracted from your participation fee of $7. If you agree to make a loss that is greater than your $7 participation fee for that round, your participation in the experiment for all subsequent rounds will be suspended, and you will earn nothing beyond your show-up fee of $3, which you will receive after the conclusion of all the rounds in the experiment.

Time Limits and Trading Reductions

Each bargaining round is split into 8 stages. Once the 8 stages are complete, the round is over, and if you and your partner failed to accept an offer, each of you earns no profit from trading in that round.

Stage 1 lasts 40 seconds (30 seconds for the proposer, and 10 seconds for the responder). Each of stages 2 to 8 last 20 seconds (15 seconds for the proposer and 5 seconds for the responder). The top of the screen will indicate to you what stage you are in, and how much time you have left in that stage. Once a stage is complete, the next one will start immediately, until the last (8^{th}) stage is complete.

If you or your partner accept an offer during the first bargaining stage, your earnings for that round will be equal to the numbers shown on the screen.

If you or your partner accept an offer during the second bargaining stage or later, there will be a reduction on the profit that both players receive from trading (this reduction does not apply to the $7 participation fee, or the $3 show-up fee). This reduction will take the form of a percentage that will be deducted from your earnings and those of your partner. The table below specifies the reductions.

Bargaining stage	% of Your Earnings that You Lose
1	0%
2	10%
3	20%
4	30%
5	40%
6	50%
7	60%
8	70%
End of round	Round canceled for you and your partner (and you both earn no profit from trading)

For example, if it took until bargaining stage 2 for one of you to accept an offer, then each of you will lose 10% of your earnings.

Sequence of events

In each round, the sequence of decisions is as follows.

1. The computer randomly matches you with another trader.
2. You are randomly assigned your costs/values for the 4 units. A list of the 24 possible combinations is provided below.
3. The computer randomly selects one member of each pair of traders to make the first proposal.
4. That trader makes a proposal that includes 1 to 6 bundle offers.

5. The other trader can accept, at most, exactly one of the offers, or reject all of them.
6. If the trader rejects, it becomes that trader's turn to make a proposal.
7. The opportunity to make a proposal keeps alternating until
    a. one trader accepts an offer included in a proposal
    b. or 8 proposals are rejected, which will suspend bargaining in that round
8. Earnings for that round are calculated, and a reduction may be applied depending on which bargaining stage the acceptance came in.
9. At the conclusion of the last real round, one of the 6 real rounds is selected at random and used to pay the participants.

[List of all possible cost/value combinations on next page]

Possible cost/value combinations

Combination No.	Unit 1		Unit 2		Unit 3		Unit 4	
	Seller Cost	Buyer Value	Seller Cost	Buyer Value	Seller Cost	Buyer Value	Seller Cost	Buyer Value
1	$0	$10	$0	$50	$40	$10	$40	$50
2	$0	$10	$0	$50	$40	$50	$40	$10
3	$0	$10	$40	$10	$0	$50	$40	$50
4	$0	$10	$40	$10	$40	$50	$0	$50
5	$0	$10	$40	$50	$0	$50	$40	$10
6	$0	$10	$40	$50	$40	$10	$0	$50
7	$0	$50	$0	$10	$40	$10	$40	$50
8	$0	$50	$0	$10	$40	$50	$40	$10
9	$0	$50	$40	$10	$0	$10	$40	$50
10	$0	$50	$40	$10	$40	$50	$0	$10
11	$0	$50	$40	$50	$0	$10	$40	$10
12	$0	$50	$40	$50	$40	$10	$0	$10
13	$40	$10	$0	$10	$0	$50	$40	$50
14	$40	$10	$0	$10	$40	$50	$0	$50
15	$40	$10	$0	$50	$0	$10	$40	$50
16	$40	$10	$0	$50	$40	$50	$0	$10
17	$40	$10	$40	$50	$0	$10	$0	$50
18	$40	$10	$40	$50	$0	$50	$0	$10
19	$40	$50	$0	$10	$0	$50	$40	$10
20	$40	$50	$0	$10	$40	$10	$0	$50
21	$40	$50	$0	$50	$0	$10	$40	$10
22	$40	$50	$0	$50	$40	$10	$0	$10
23	$40	$50	$40	$10	$0	$10	$0	$50
24	$40	$50	$40	$10	$0	$50	$0	$10

This table shows that **the matchings of cost and value are always the same** but they may **occur in the case of different units.** All combinations have a cost $0 meeting a value $10, a cost $0 meeting a value $50, a cost $40 meeting a value $10, and a cost $40 meeting a value $50. However, across the various possible combinations, these matchings of cost and value occur at different units.

As an example, consider combinations (1) and (8) of the table:

- $0 meets a $10 in unit 1 of combination (1), but in unit 2 of combination (8)
- $0 meets a $50 in unit 2 of combination (1), but in unit 1 of combination (8)
- $40 meets a $10 in unit 3 of combination (1), but in unit 4 of combination (8)
- $40 meets a $50 in unit 4 of combination (1), but in unit 3 of combination (8)

Summary of Key Points

By way of providing a summary of the key points, we suggest that you return to the table of page 7[7] and consider:

1. What information a buyer/seller will have in each round
2. What are the combinations of costs and values for the four units
3. How, given a particular combination of traded units, a price determined in a certain bargaining stage, that may involve a reduction, determines the profit that a buyer and a seller will make

INSTRUCTIONS: FOUR GOODS, OPEN CHAT[8]

Welcome and Introduction

Welcome to our study of decision making. If you read these instructions carefully and make good decisions, you can earn a considerable amount of money.

Kindly refrain from talking with other participants during the session. Also, it is an important requirement of this experiment that you **please turn your mobile phones to silent and abstain from using them during the experiment**.

By coming to this session, you have earned an attendance fee of $10. This has two components. A show up fee of $3 and a participation fee of $7. The participation fee is discussed further below.

The $10 attendance fee will be in addition to any amount that you earn based on your decisions. Once your earnings are determined you will be paid privately and your participation in this experiment will then be concluded.

If you have a question at any time, please raise your hand and I will approach you so that you can ask your question in private.

[7] The pages have been renumbered in this supplement. This is the second table that appears in the instructions for this treatment.

[8] This title is, of course, not part of the instructions that were distributed to experimental participants which begin with the line "Welcome and Introduction" below.

Setup

Today, we are going to set up a market in which some of you will be buyers and some of you will be sellers. In this market you will be given the opportunity to trade commodities. We will not specify names for these commodities; we will simply refer to them as “units”.

Trading will occur in a sequence of trading rounds. The decisions that you make in each round will determine your earnings.

The experiment will consist of 10 rounds: 4 practice rounds followed by 6 real rounds.

The first 4 rounds will be practice and will not affect your earnings for the experiment.

The final 6 rounds will be real and will affect your earnings. At the end of the experiment, the computer will select one of the 6 real rounds at random and you will be paid based on your earnings in that round in cash.

Matching Rules

In every round, the computer will tell you whether you are a buyer or a seller for that round. During the practice rounds, you will experience both the role of a buyer as well as the role of a seller. Once we have completed the practice rounds, you will be assigned the role of either buyer or seller and will remain in that role throughout the remainder of the session. You have a 50% chance of being a buyer, and a 50% chance of being a seller.

Each round, every buyer will be randomly matched with a seller. That means that each buyer is equally likely to be matched with each seller. The matching is independent every round, which means that being matched with a specific trader in one round has no effect on the likelihood of being matched with the same trader in a future round.

All matching is anonymous, meaning that you will never know the identity of whom you are matched with in any round.

Participation fee

At the start of each round, you will be given a $7 participation fee. Any positive earnings that you may make by trading one or more “units” will be added to this $7, whereas any negative earnings will be subtracted from this $7. If you do not end up trading one or more units, then you just keep your $7 participation fee.

Profit from trading

In each round, sellers and buyers will have the opportunity to exchange up to 4 units.

For each unit, sellers will be given a number known as their "cost" and buyers will be given a number known as their "value". The cost represents the minimum amount for which a seller can sell that unit without making a loss. The value represents the maximum amount for which a buyer can purchase that unit without making a loss.

Sellers and buyers exchange units in bundles that include 1 to 4 units for one total price paid for the entire bundle.

Sellers earn money when they sell a bundle at a total price that is above the total cost of the units inside that bundle. Seller earnings from the sale of a bundle are the difference between the price of the entire bundle and the total of the costs of the units included in the bundle. For example, if a seller has a unit with cost of $20 and a unit with cost of $80, and sells a bundle that includes both units for a total price of $140, then the total cost of the units in the bundle is $20 + $80 = $100, and the seller earns $140 – $100 = $40.

Buyers earn money when they buy a bundle at a total price that is below the total value of the units inside that bundle. Buyer earnings from the purchase of a bundle are the difference between the total of the values of the units included in the bundle and the price of the entire bundle. For example, if a buyer's value for one unit is $90 and for another unit is $60, and buys a bundle that includes both units for a total price of $100, then the total value of the units in the bundle is $90 + $60 = $150, and the buyer earns $150 – $100 = $50.

EXAMPLE:

Consider a case in which a seller and a buyer have the following costs and values for the 4 units:

	Unit 1	Unit 2	Unit 3	Unit 4
Seller's cost ($)	A	B	C	D
Buyer's value ($)	E	F	G	H

The seller's costs of units 1, 2, 3, and 4 are A, B, C, and D, respectively.
The buyer's values of units 1, 2, 3, and 4 are E, F, G, and H, respectively.

Assume that the buyer and seller agree to exchange a bundle that includes only Unit 2 for a price $X. Their profit is calculated as follows:
Seller's profit=X-B
Buyer's profit=F-X

Alternatively, assume that the buyer and seller agree to exchange a bundle that includes Unit 1 and Unit 2 for a price $X. Their profit is calculated as follows:
Seller's profit=X-(A+B)

Buyer's profit=(E+F)-X

The calculations are similar in the case of bundles that include 3 or 4 goods.

If a seller sells a bundle at a price that is less than the total of the costs of the units inside the bundle, this seller will make a loss. If a buyer buys a bundle at a price that is greater than the total of the values of the units inside the bundle, this buyer will make a loss. If you are at risk of making a loss, the computer will notify you and ask you to confirm.

A seller and a buyer **can exchange at most ONLY ONE BUNDLE per round**. If they do not exchange a bundle, they each earn a profit of $0 for that round.

How the seller's cost and the buyer's value of each of the 4 units are determined is described in the next section.

How are Costs and Values Determined?

For each seller-buyer pair, there is a fixed number of possibilities for the combination of costs and values that can occur for each of their four units:

		Seller's cost	Buyer's value
a.	For one of the four units:	$0	$10
b.	For one of the four units:	$0	$50
c.	For one of the four units:	$40	$10
d.	For one of the four units:	$40	$50

Each of these 4 possibilities, (a), (b), (c), and (d) above, MUST OCCUR EXACTLY ONCE in any round. The order with which they occur within a round is random, and each possibility regarding the order in which (a), (b), (c), and (d) occur is equally likely. Hence, knowing this order in one round tells you nothing about how it may manifest in any other round. A list of all the equally likely possibilities is provided at the end of this document.

Sellers are only informed of their cost for each unit, and buyers are only informed of their value for each unit, meaning that neither side knows about the other. So, for a given unit, a seller cannot distinguish between combinations (a) and (b) above, because the seller is given the same information in each case (cost of $0); or between (c) and (d) (cost of $40). Similarly, for a given unit, a buyer cannot distinguish between (a) and (c), because the buyer is given the same information (value of $10) or between (b) and (d) (value of $50).

In other words:

In any round, sellers will have two units each with cost $0 and two units each with cost $40. However, they will not know which of their two $0s meets a buyer's $10 and which meets a buyer's $50. Similarly, they will not know which of their two $40s meets a buyer's $10 and which meets a buyer's $50.

In any round, buyers will have two units each with value $50 and two units each with value $10. However, they will not know which of their two $50s meets a seller's $0 and which meets a seller's $40. Similarly, they will not know which of their two $10's meets a seller's $0 and which meets a seller's $40.

A screenshot below shows a trader discovering that they are a seller, and learning their costs:

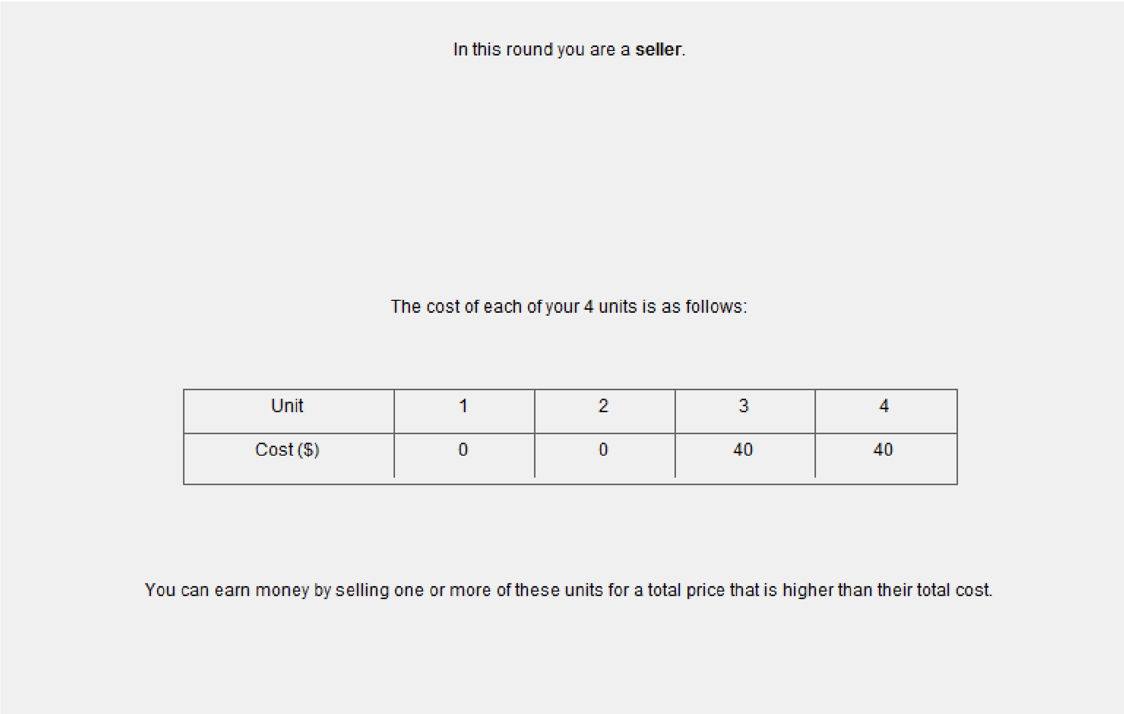

In this round you are a **seller**.

The cost of each of your 4 units is as follows:

Unit	1	2	3	4
Cost ($)	0	0	40	40

You can earn money by selling one or more of these units for a total price that is higher than their total cost.

A screenshot below shows a trader discovering that they are a buyer, and learning their values:

In this round you are a **buyer**.

The value of each of your 4 units is as follows:

Unit	1	2	3	4
Value ($)	10	50	10	50

You can earn money by buying one or more of these units for a total price that is lower than their total value.

How Much Money Can You and Your Partner Make

When you and your partner are bargaining over a bundle of 1 to 4 units, the total money that you can make between you from successfully trading that bundle of 1 to 4 units is equal to the buyer's total value of the traded units minus the seller's total cost of the traded units. This follows from the way in which buyer and seller profits are calculated.

To understand this, we will examine an example with one unit. Consider a buyer with value $50 and a seller with cost $40. If they agree on a price $43 then the buyer will make $50 - $43 = $7 and the seller $43 - $40 = $3. Together they make $7 + $3 = $10. Alternatively, they could agree on a price $48 and then the buyer will make $50 - $48 = $2 and the seller $48 - $40 = $8. Again, together they make $2 + $8 = $10. Clearly, the total profit is the difference between the cost and the value: $50 - $40 = $10. The trade price merely determines how the total profit of $10 is divided between the buyer and seller.

This example was for one unit. When dealing with 4 units, the same principle applies: the total money that you and your partner can make is the total of how much you can make from each unit.

Since in every round, each of the following 4 possibilities occurs exactly once, there is a **fixed maximum total amount of profit** that you and your partner can make every round, and it is equal to $70. To see why, consider the following table:

		Seller's cost	Buyer's value	Combined buyer and seller profit if traded	Combined buyer and seller profit if not traded
a.	For one of the four units:	$0	$10	$10 - $0 = $10	$0
b.	For one of the four units:	$0	$50	$50 - $0=$50	$0
c.	For one of the four units:	$40	$10	$10 - $40 = -$30	$0
d.	For one of the four units:	$40	$50	$50 - $40 = $10	$0

Hence, the **maximum** amount of combined buyer and seller earnings is $10 + $50 + $0 + $10 = $70.

Therefore, depending on the choices that you and your partner make in every round, you can together make a maximum of $70 between you.

Any amount that you and your partner make is shared between you on the basis of the trading agreement that you make with your partner, which is done according to the mechanism described in the next section.

Trading Mechanism

After being matched with another trader and seeing your costs/values, you will have the opportunity to trade with your partner in a sequence of 8 bargaining stages. Bargaining stages consist of you and your partner making offers of trading prices for bundles of 1 to 4 units.

During each bargaining stage, you and your partner are free to make offers at any time. Each of you is also free to accept an offer made by the other at any time. Once an offer is accepted, the round is over. You and your partner also have a time limit, which we will explain below.

In what follows we first describe the offer structure, and then we describe how to accept an offer.

Offer Structure

Each offer has two components:

3. A list of the units included in the offer; an offer must contain at least 1 unit, and at the most 4 units
4. A total price for **all the units** in the offer, which must be a whole number between $0 and $200, inclusive.

To make an offer, select the units that will be included in the offer using the checkboxes, and enter the price.

The screenshot below shows you a buyer considering what offer to make:

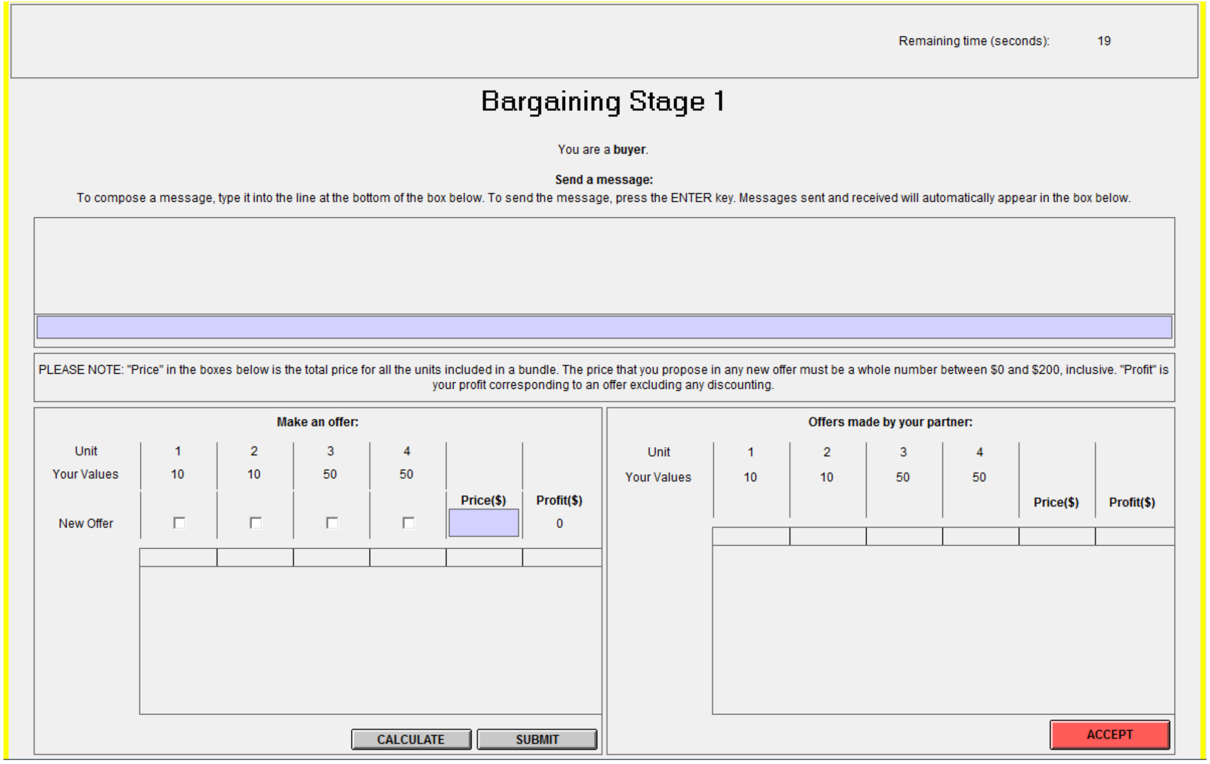

Once you have finished composing your offer, press the CALCULATE button to calculate how much profit you will make if your offer is accepted and then press the SUBMIT button to submit the offer to your trading partner.

The offers that you make will appear on the left side of the screen. The most recent offer will be on top, and the older ones will be at the bottom.

Your partner is free to accept any of your offers at any point during the round. If your partner accepts an offer, the round ends immediately.

While you are making offers, your partner can also make offers. **Both of you can make as many offers as you like without waiting for the other person to respond.** Your partner's offers will appear on the right side of your screen, with the most recent ones at the top, and the older ones at the bottom.

The computer will automatically display the profit that you will earn from each offer if you accept it.

During a round, you and your partner will be able to chat with each other via a chat window. You can see this chat window at the top of the screenshot above. You can send messages at any time, and can send as many messages as you like. However, we ask that you please observe the following rules:

- No threatening or abusive language
- No messages that reveal your identity, such as your name or your location in the lab.

Accepting an Offer

To accept an offer, highlight the offer by moving the mouse cursor over it, and then click on ACCEPT. The round will end immediately as soon as either you or your partner accepts an offer.

Making a Loss and Exhausting Your $7

You will never be compelled to make a loss. If you make a decision that puts you at risk of making a loss, the computer will warn you with a pop-up message and ask for confirmation. The pop-up window will have two buttons: "OK" and "CANCEL". Press **"OK" to accept to make a loss** or **"CANCEL" to revise your decision.** If you agree to make a loss up to $7, this loss will simply be subtracted from your participation fee of $7. If you agree to make a loss that is greater than your $7 participation fee for that round, your participation in the experiment for all subsequent rounds will be suspended, and you will earn nothing beyond your show-up fee of $3, which you will receive after the conclusion of all the rounds in the experiment.

Time Limits and Trading Reductions

Each bargaining round is split into 8 stages. Once the 8 stages are complete, the round is over, and if you and your partner failed to accept an offer, each of you earns no profit from trading in that round.

Stage 1 lasts 40 seconds, while each of stages 2 to 8 last 20 seconds. The top of the screen will indicate to you what stage you are in, and how much time is left until the end of the stage. Once a stage is complete, the next one will start immediately, until the last (8^{th}) stage is complete.

Offers do not expire and an offer made in one stage can be accepted in that stage or in any subsequent stage.

If you or your partner accept an offer during the first bargaining stage, your earnings for that round will be equal to the numbers shown on the screen.

If you or your partner accept an offer during the second bargaining stage or later, there will be a reduction on the profit that both players receive from trading (this reduction does not apply to the $7 participation fee, or the $3 show-up fee). This reduction will take the form of a percentage that will be deducted from your earnings and those of your partner. The table below specifies the reductions.

Bargaining stage	% of Your Earnings that You Lose
1	0%
2	10%
3	20%
4	30%
5	40%
6	50%
7	60%
8	70%
End of round	Round canceled for you and your partner (and you both earn no profit from trading)

For example, if it took until bargaining stage 2 for one of you to accept an offer, then each of you will lose 10% of your earnings.

Sequence of events

In each round, the sequence of decisions is as follows.

1. The computer randomly matches you with another trader.
2. You are randomly assigned your costs/values for the 4 units. A list of all 24 possible combinations is provided below.
3. Bargaining commences; you and your partner are free to make offers, accept offers, and chat.
4. Bargaining continues until:
    a. one trader accepts an offer
    b. or 8 bargaining stages are concluded, which will suspend bargaining in that round
5. Earnings for that round are calculated, and a reduction may be applied depending on which bargaining stage the acceptance came in.
6. At the conclusion of the last real round, one of the 6 real rounds is selected at random and used to pay the participants.

[List of all possible cost/value combinations on next page]

Possible cost/value combinations

Combination No.	Unit 1		Unit 2		Unit 3		Unit 4	
	Seller Cost	Buyer Value	Seller Cost	Buyer Value	Seller Cost	Buyer Value	Seller Cost	Buyer Value
1	$0	$10	$0	$50	$40	$10	$40	$50
2	$0	$10	$0	$50	$40	$50	$40	$10
3	$0	$10	$40	$10	$0	$50	$40	$50
4	$0	$10	$40	$10	$40	$50	$0	$50
5	$0	$10	$40	$50	$0	$50	$40	$10
6	$0	$10	$40	$50	$40	$10	$0	$50
7	$0	$50	$0	$10	$40	$10	$40	$50
8	$0	$50	$0	$10	$40	$50	$40	$10
9	$0	$50	$40	$10	$0	$10	$40	$50
10	$0	$50	$40	$10	$40	$50	$0	$10
11	$0	$50	$40	$50	$0	$10	$40	$10
12	$0	$50	$40	$50	$40	$10	$0	$10
13	$40	$10	$0	$10	$0	$50	$40	$50
14	$40	$10	$0	$10	$40	$50	$0	$50
15	$40	$10	$0	$50	$0	$10	$40	$50
16	$40	$10	$0	$50	$40	$50	$0	$10
17	$40	$10	$40	$50	$0	$10	$0	$50
18	$40	$10	$40	$50	$0	$50	$0	$10
19	$40	$50	$0	$10	$0	$50	$40	$10
20	$40	$50	$0	$10	$40	$10	$0	$50
21	$40	$50	$0	$50	$0	$10	$40	$10
22	$40	$50	$0	$50	$40	$10	$0	$10
23	$40	$50	$40	$10	$0	$10	$0	$50
24	$40	$50	$40	$10	$0	$50	$0	$10

This table shows that **the matchings of cost and value are always the same** but they may **occur in the case of different units.** All combinations have a cost $0 meeting a value $10, a cost $0 meeting a value $50, a cost $40 meeting a value $10, and a cost $40 meeting a value $50. However, across the various possible combinations, these matchings of cost and value occur at different units.

As an example, consider combinations (1) and (8) of the table:

- $0 meets a $10 in unit 1 of combination (1), but in unit 2 of combination (8)
- $0 meets a $50 in unit 2 of combination (1), but in unit 1 of combination (8)
- $40 meets a $10 in unit 3 of combination (1), but in unit 4 of combination (8)
- $40 meets a $50 in unit 4 of combination (1), but in unit 3 of combination (8)

Summary of Key Points

By way of providing a summary of the key points, we suggest that you return to the table of page 7[9] and consider:

1. What information a buyer/seller will have in each round
2. What are the combinations of costs and values for the four units
3. How, given a particular combination of traded units, a price determined in a certain bargaining stage, that may involve a reduction, determines the profit that a buyer and a seller will make

[9] The pages have been renumbered in this supplement. This is the second table that appears in the instructions for this treatment.