Using Machine Learning to understand Bargaining Experiments

Colin F. Camerer
California Institute of Technology
camerer@hss.caltech.edu

Hung-Ni Chen
LMU Munich
hungni.chen@econ.lmu.de

Po-Hsuan Lin
California Institute of Technology
plin@caltech.edu

Gideon Nave
University of Pennsylvania
gnave@wharton.upenn.edu

Alec Smith
Virginia Tech
alecsmith@vt.edu

Joseph Tao-yi Wang
National Taiwan University
josephw@ntu.edu.tw

December 30, 2019

1 Introduction

As all the chapters in this book discuss, bargaining is a fundamental economic activity. This chapter is about a general class of bargaining games in which there is private information about the amount that is being bargained over (often called the "pie size"). This class is most common in everyday bargaining. It is also interesting in both *theory* and *practice*.

Theory is interesting because when there is private information and people are self-interested, theories based on individual rationality typically predict an inevitable loss of efficiency. That is, even when a bargain is mutually beneficial for both sides, they will not always come to agreement.

Private information bargaining is interesting in *practice* because, while inefficiencies are predicted by theory, it also known that if there are observable statistical proxies for the hidden private information, then sets of rules (mechanisms) which use this information can improve efficiency [McAfee and Reny, 1992, Créemer and McLean, 1985, Cremer and McLean, 1988]. Therefore, it is possible that methods for measuring private information can improve efficiency, even when bargainers voluntarily participate in systems using those measures.

There is a long history of using highly controlled laboratory experiments to study bargaining. A brief description of this history helps explain why we are enthusiastic about modern applications of machine learning.

1.1 A brief history of bargaining experiments

The experimental literature on bargaining is vast, so we focus only on those studies closely related to ours.[1] Before theoretical breakthroughs in understanding structured bargaining, most experiments used unstructured communication. The main focus of interest was process-free solution concepts such as the Nash bargaining solution [Nash Jr, 1950], and important extensions to those concepts [e.g. Kalai et al., 1975].

We will refer to the amount of surplus available to share as the "pie". Many bargains [Nydegger and Owen, 1974, Roth and Malouf, 1979] led to an equal split of the pie. Roth suggested that "bargainers sought to identify initial bargaining positions that had some special reasons for being credible... that served as *focal points* that then influenced the subsequent conduct of negotiation" [Roth, 1985]. Under informational asymmetries, disagreements may arise due to coordination

[1] For longer reviews, see Kennan and Wilson [1993], Ausubel et al. [2002], Thompson et al. [2010]

difficulties. Several papers by Roth and colleagues then explored what happens when players bargain over points which have different financial value to players [Roth and Malouf, 1979, Roth et al., 1981, Roth and Murnighan, 1982, Roth, 1985]. In theory, there should be no disagreements in these games but a modest percentage of trials (10-20%) did result in disagreement, which seems to involve differences about which "focal points" are acceptable.[2] Roth et al. [1988], also drew attention to the fact that the large majority of agreements are made just before a (known) deadline, an observation called the "deadline effect."

Two pioneering papers, Ståhl [1972] and Rubinstein [1982] showed how noncooperative game theory might be used to improve the apparent precision of bargaining theories. Since then, almost all experimental studies have tested what happens in highly structured settings using variants of those early game structures; for a review see Ausubel et al. [2002]. In these theories and experiments, "structure" means that the rules of how bargaining proceeds are clearly specified in the theory; put differently, bargainers have no freedom to time their offers or use natural language. The rules typically define when bargaining must be completed (either a deadline or an infinite horizon), who can offer or counteroffer and at what time, when offers are accepted, whether communication is allowed (and in what form), and so on.

Theoretical predictions of outcomes and payoffs depend sensitively on these structural features.[3] That structural-sensitivity proved to be enticing, because it created a cornucopia of interesting experiments testing whether bargaining was sensitive to structured features as theory predicted. This led to a burst of progress in experimental literature testing these theories.[4]

Many other experiments have observed what happens in semi-structured bargaining in which there is *two-sided* private information [Valley et al., 2002]. The term "semi-structured" means that there is structure about bargainers' valuations and beliefs, but players may make offers at any time, and offers can be accompanied by natural language. The typical finding is that in face-to-face and unstructured communication via message-passing, there are *fewer* disagreements than predicted by theory. (A comparable finding in sender-receiver games is that senders willingly share more private information than is selfishly rational; see Crawford [2003], Cai and Wang [2006], Wang et al. [2010a].) However, in the highly structured case in which the bargainers can only make a single offer and no natural language is allowed, disagreements are more common, and the key predictions of theory hold surprisingly well [Radner and Schotter, 1989, Rapoport et al., 1995, Rapoport and Fuller, 1995, Daniel et al., 1998].

Since the rise of structured bargaining theories, experimentation in economics on unstructured bargaining has all but disappeared. Our paper returns to this less popular route, exploring unstructured bargaining with one-sided private information in an experiment.

There are three good reasons to revive the study unstructured bargaining now.

First, most natural two-player bargaining is *not* highly structured. Conventional methods for conducting bargaining do emerge in natural settings, but these methods are rarely constrained, because there are no penalties for deviating from conventions. Studying unstructured bargaining is of particular importance, as strategic behavior may substantially differ between static and dynamic environments that allow continuous-time interaction [Friedman and Oprea, 2012]. There may also be clear empirical regularities in unstructured bargaining—such as deadline effects [Roth et al., 1988, Gächter and Riedl, 2005]—that are evident in the data but not predicted by most theories (though see [Fuchs and Skrzypacz, 2013]). Establishing these regularities can *lead* theorizing, rather than test theory.

Second, even when bargaining is unstructured, theory can still be applied to make clear interesting predictions. A natural intuition is that when bargaining methods are unstructured, no clear predictions can be made, as if the lack of structure in the bargaining protocol must imply a lack of structure (or precision) in predictions. This intuition is just not right. In the case we study, clear predictions about unstructured bargaining do emerge, thanks to the wonderful "revelation principle" [Myerson, 1979, 1984]. This principle has the useful property of implying empirical predictions for noncooperative equilibria, independently of the bargaining protocol, based purely on the information structure. For example, the application of the revelation principle in our setting

[2] See Schelling [1960], Roth [1985], Kristensen and Gärling [1997], Janssen [2001], Binmore and Samuelson [2006], Janssen [2006], Bardsley et al. [2009], Isoni et al. [2014, 2013], Hargreaves Heap et al. [2014].

[3] See Cramton [1984], Chaussees [1985], Rubinstein [1985], Grossman and Perry [1986], Gul and Sonnenschein [1988], Ausubel and Deneckere [1993].

[4] See Ochs and Roth [1989], Camerer et al. [1993], Mitzkewitz and Nagel [1993], Güth et al. [1996], Kagel et al. [1996], Güth and Van Damme [1998], Rapoport et al. [1998], Kagel and Wolfe [2001], Srivastava [2001], Croson et al. [2003], Johnson et al. [2002], Kriss et al. [2013].

leads to the prediction that strikes will become less common as the amount of surplus the players are bargaining over grows. This type of prediction is non-obvious and can be easily tested. Furthermore, if additional assumptions are made about equilibrium offers, and combined with the revelation principle, then exact numerical predictions about offers and strike rates can be made. That is, even if the bargaining protocol lacks structure, predictions can have plenty of restricted "structure" thanks to the beautiful game theory.

Third, unstructured bargaining creates a large amount of interesting data during the bargaining process. Players can make offers at any time, retract offers, and so on. Natural language can be analyzed, perhaps including vocal properties in verbal communication[5]. Self-reported and biological measures of emotion, cognitive effort, visual attention to display elements, and even neural activity can also be gathered.

Our view is that theoretical and experimental economists regarded these types of data as a nuisance—a "bug" in an experimental design rather than a "feature." Such data seem like a nuisance if one does not have a theory to say anything about them. However, if process variables are systematically associated with outcomes, these empirical regularities both challenge simple equilibrium theories and invite new theory development.

To this end, we use a very limited type of process data in a new way: To predict which bargaining trials will results in deals and strikes. We use a penalized regression approach from machine learning, to select predictive features from a large number of process features. Overfitting is controlled by making out of sample, cross-validated predictions. We find that a machine-learned predictive model based only on process features can predict strikes about as accurately as the pie sizes can. Adding both process and pie size together makes even better predictions.

Process data are also useful because practical negotiation advice often consists of simple heuristics about how to bargain well [Pruitt, 2013]. For example, negotiation researchers have long ago postulated that initial offers might serve as bargaining anchors and that various psychological manipulations, such as perspective taking, could potentially bias bargaining outcomes.[6] Advice like this can be easily tested by carefully controlled experimental designs that allows unstructured bargaining while keeping the process fully tractable, such as our paradigm.

The closest precursor to our design is Forsythe, Kennan, and Sopher (henceforth FKS), who studied unstructured bargaining with one-sided private information about the sizes of two possible pies [Forsythe et al., 1991b]. They used mechanism design to identify properties shared by all Bayesian equilibria of any bargaining game, using the revelation principle [Myerson, 1979, 1984]. This approach gives a "strike condition" predicting when disagreements would be ex-ante efficient. They tested their theory by conducting several experimental treatments, with free-form communication. The results qualitatively match the theory. We generalize their earlier model to capture any finite number of pie sizes. Because there are several different pie sizes, equilibria which maximize efficiency or equality create different predictions, which we test. Our experimental design uses 6 pie sizes with rapid bargaining (10 seconds per trial), where bargaining occurs only through visible offers and counter-offers, with no other restrictions. They also did not analyze their process data at all, whereas we use machine learning analysis of the process features to predict strikes on a trial-by-trial basis.

Another branch of literature that is related to our study is the experimental work investigating how humans resolve tradeoffs between equality and efficiency. While this question is still under a (heated) debate[7], it is largely accepted that people are heterogeneous with respect to how they prioritize these factors.[8]

A few recent papers have investigated highly structured strategic interactions [De Bruyn and Bolton, 2008, Blanco et al., 2011, López-Pérez et al., 2015, Jacquemet and Zylbersztejn, 2014], and some have examined free form bargaining with full information [Herreiner and Puppe, 2004, Galeotti et al., 2018]. We extend this literature by deriving theoretical predictions and test empirically how humans resolve the equality-efficiency trade-off in a dynamic strategic environment with informational asymmetry.

Finally, our study closely relates to negotiation research [Pruitt, 2013], a branch of social psy-

[5] CFC: Add Capra cite later

[6] See Kristensen and Gärling [1997], Galinsky and Mussweiler [2001], Van Poucke and Buelens [2002], Mason et al. [2013], Ames and Mason [2015].

[7] See Kritikos and Bolle [2001], Charness and Rabin [2002], Engelmann and Strobel [2004, 2006], Fehr et al. [2006], Bolton and Ockenfels [2006], Durante et al. [2014], El Harbi et al. [2015].

[8] For example, economics students are inclined to favor efficiency over equality, females are more egalitarian than males, and political preferences do not seem to have an effect [Engelmann and Strobel, 2004, Fehr et al., 2006].

chology and organizational behavior research. In contrast to economic theories that typically describe behavior in equilibrium (i.e., when players best respond to each other's actions), negotiation theories assume that bargainers are not in equilibrium and focus on prescriptive models, in which adopting certain strategies improves negotiation outcomes. Negotiation researchers take into account the process of bargaining by studying psychological constructs such as aspirations, defined as "the highest valued outcome at which the negotiator places a non-negligible likelihood that that value would be accepted by the other party" [White and Neale, 1994]. Aspirations play an important role in determining the bargainers' initial offers, and were shown to influence bargaining outcome variables such as disagreement rates and surplus division.[9]

The reminder of this paper is organized as follows. In section 2.1, we use mechanism design theory to derive general qualitative properties of bargaining in equilibrium. We show that in our setting, no equilibrium satisfies both equality and efficiency in all states of the world, and propose two equilibria that solve this tradeoff by either favoring the former or the latter. We present a novel experimental design in section 2.2, and summarize its general results in section 3.1. We use machine learning to examine how bargaining process data can be associated with bargaining outcome variables in section 3.2, and conclude in section 4.

2 Theory and Experiments

In this paper, we adopt the theoretical framework from Camerer et al. [2019] who extend the two state model developed in Kennan and Wilson [1990] and Forsythe et al. [1991b] to an arbitrary finite number of states. This extension yields comparative statics predictions regarding the frequency of disagreements in each state with only the game structure, incentive compatibility (IC) and individual rationality (IR) constraints. However the mechanism design approach only characterizes the class of possible equilibria rather than predicts specific outcomes. Thus, Camerer et al. [2019] further take advantage of the focal points in this game to obtain testable predictions about both deal rates and payoffs in each state. In this section, we first introduce the theoretical framework and then the experiments.

2.1 Theoretical Framework

In this unstructured bargaining game, two players bargain over an economic surplus or "pie," which is a random variable denoted by π. The finite set of true states by indexed by $k \in \{1, 2, \dots, K\}$, and the pie amount in state k is π_k. Without loss of generality, we assume $\pi_k > \pi_j$ when $k > j$. The informed player knows the true pie amount. The uninformed player does not know the pie amount, but knows the informed player knows it. The probability distribution over pie sizes $\Pr(\pi_k) = p_k$ is common knowledge. The players bargain over the payoff of the uninformed player, denoted by w, by continuously communicating their bids within a certain amount of time T—which is also a common knowledge. If the players agree on a payoff for the uninformed player w, then the informed player gets the rest of the pie $\pi - w$. If they do not agree on an allocation before the deadline, both players get nothing and we refer to this outcome as a disagreement, or in keeping with the motivation of Forsythe et al. [1991b], as a *strike*, while successful bargaining outcomes are *deals*.

From a mechanism design perspective, we can view this bargaining game as a process of transmitting the private information about the pie size from the informed player to the uninformed player. By the revelation principle (Myerson [1979, 1984]), we know that every Nash equilibrium in this bargaining game can be implemented in an incentive compatible direct mechanism where the informed player truthfully reports the actual state to a neutral mediator and the player's payoffs are equal to their payoffs in the original bargaining game.

Following Forsythe et al. [1991b] and Camerer et al. [2019], in the direct mechanism the informed player announces that the state is $j \in \{1, \dots, K\}$. Given the announcement, the neutral mediator determines the deal probability (γ_j) and the payoff to the the uninformed player (x_j). The informed player gets the rest of the pie ($\gamma_j \pi_k - x_j$). Thus a mechanism involves $2K$ parameters, $\{\gamma_k, x_k\}_{k=1}^K$.

A mechanism is incentive compatible (IC) if it is optimal for players to reveal their private information. In our setting, this means that the informed player's expected payoff must be (weakly)

[9] See Yukl [1974], White and Neale [1994], White et al. [1994], Kristensen and Gärling [1997], Galinsky and Mussweiler [2001], Van Poucke and Buelens [2002], Buelens and Van Poucke [2004], Mason et al. [2013], Ames and Mason [2015].

maximized in the direct mechanism when she announces the true size of the pie. This requires

$$\gamma_k \pi_k - x_k \geq \gamma_j \pi_k - x_j, \quad \forall k \text{ and } \forall j \neq k. \tag{IC}$$

An IC-mechanism is individually rational (IR) when both players prefer to participate in it. Assuming the players' payoffs from not participating are zero, this means that for every state k the expected payoff to each player is positive, so that

$$\gamma_k \pi_k - x_k \geq 0, \tag{IR-1}$$

$$x_k \geq 0. \tag{IR-2}$$

Based on the IC, IR-1, and IR-2 and conditions, Camerer et al. [2019] prove the following two lemmas.

Lemma 1. *If the bargaining mechanism satisfies the IC, IR-1, and IR-2 conditions, then:*

1. *Deal rates are monotonically increasing in the pie size x_k.*
2. *The uninformed player's payoffs are monotonically increasing in the pie size.*
3. *The uninformed player's payoff is identical for all states in which the deal probability is 1.*

Note that the payoff of the uniformed player x_k in the direct mechanism is equivalent to the expected payoff of the uninformed player in state k of the bargaining game: $x_k = \gamma_k w_k$, where w_k is the uninformed player's payoff conditional on a deal being made in state k.

The direct mechanism is interim-efficient [Holmström and Myerson, 1983] if the payoff profile is Pareto-optimal for the set of $K + 1$ agents: the informed player in each of the K possible states of the world and the uninformed player (in expectation). Interim efficiency implies the following *strike condition*:

Lemma 2. *For any mechanism that satisfies the IR-1, IR-2 and IC conditions, strikes in state k are interim-efficient if*

$$\frac{\pi_k}{\pi_{k+1}} < \frac{\left(1 - \sum_{j=1}^{k} p_j\right)}{\left(1 - \sum_{j=1}^{k-1} p_j\right)} = \frac{\Pr(\pi \geq \pi_{k+1})}{\Pr(\pi \geq \pi_k)}.$$

For the proof, see Camerer et al. [2019].

The IC, IR-1, IR-2, and strike conditions limit the scope of possible bargaining outcomes and predict when strikes are likely to occur. However, they are not sufficient to pin down the strike rates $1 - \gamma_k$ and the equilibrium payoffs w_k in each state. To make a more precise prediction, Camerer et al. [2019] use an equilibrium selection approach that assumes that equal payoff splits are natural focal points. In the experiments, the possible states, π, takes on values that are the integer dollar amounts between \$1 and \$6 with equal probability. Therefore, we can restrict the state space to $\{\$1, \ldots, 6\}$.

The importance of focal points has been well-studied in the literature (Schelling [1960], Roth [1985], Kristensen and Gärling [1997], Janssen [2001], Binmore and Samuelson [2006], Janssen [2006], Bardsley et al. [2009], Isoni et al. [2013, 2014]). Absent other salient features of bargaining, the natural focal point is an equal split (i.e., $w_k = \frac{\pi_k}{2}$). Indeed, equal splits often emerge in bargaining experiments (e.g. Lin et al. [2018]). Based on this tendency that players to tend to coordinate on the equal-split allocation, Camerer et al. [2019] propose that the equilibrium payoff of the uninformed player, conditional on a deal, will equal half of the pie size ($w_k = \frac{\pi_k}{2}$) as long as an equal split satisfies the IR and IC conditions (Lemma 1), and subject to efficiency conditions. By either prioritizing the former or the latter, Camerer et al. [2019] derive two competing equilibrium predictions, which are **the efficient equilibrium** and **the equal split equilibrium**.

2.1.1 The Efficient Equilibrium

The IC conditions and Lemma 1 show that if there exists a cutoff state π_c where the deal rate γ_c, is equal to 1, then strikes are inefficient in all states π_k such that $k \geq c$. Lemma 1 also predicts that the uninformed player's payoff must be the same in all states where no disagreements occur.

The efficient equilibrium prioritizes efficiency over equality. In this equilibrium, the deal rate is assumed to be 1 whenever the strike condition (Lemma 2) does not hold. To obtain a precise prediction about the equilibrium uninformed payoffs w_k^*, we assume that players divide the pie equally in lower-value pie states given this constraint:

$$w_k^* = \begin{cases} \frac{\pi_k}{2}, & \forall \pi_k \le \pi_c, \\ \frac{\pi_c}{2}, & \forall \pi_k > \pi_c. \end{cases}$$

In our experiment, π is an integer dollar between 1 and 6 with equal likelihood. It follows numerically that the strike condition (Lemma 2) holds for pies of size 1 and 2. When $\pi = 3$, the two sides of the inequality are equal, so the strike rate is indeterminate. When $\pi \ge 4$, there should be no strikes. Therefore, we can first pin down that the deal rates in large pies would satisfy $\gamma_4 = \gamma_5 = \gamma_6 = 1$. Furthermore, as we plug w_k^* into the IC constraint, we can obtain the following predictions regarding payoffs and deal rates in the efficient equilibrium:

$$(w_1, w_2, w_3, w_4, w_5, w_6) = \left(\tfrac{1}{2}, 1, \tfrac{3}{2}, 2, 2, 2\right),$$
$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = \left(\tfrac{2}{5}, \tfrac{3}{5}, \tfrac{4}{5}, 1, 1, 1\right).$$

2.1.2 The Equal Split Equilibrium

On the other hand, some efficiency must be sacrificed in order to achieve equality for each pie size. In the equal split equilibrium, we first assume that players split the pie equally in all states, and only then maximize efficiency, so that

$$w_k^{**} = \frac{\pi_k}{2}.$$

in all states.

Because deal rates are increasing with the pie size and the uninformed player's payoff must be identical in all states where there are no strikes (Lemma 1), full equality implies that efficiency (i.e., no strikes) can only be achieved in the largest pie. Thus, to pin down exact numerical predictions of deal rates in the equal split equilibrium, we set $\gamma_6 = 1$. Inserting $w_k^{**} = \frac{\pi_k}{2}$ into the IC constraint, we derive the predicted payoffs and deal rates in the equal split equilibrium:

$$(w_1, w_2, w_3, w_4, w_5, w_6) = \left(\tfrac{1}{2}, 1, \tfrac{3}{2}, 2, \tfrac{5}{2}, 3\right),$$
$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = \left(\tfrac{2}{7}, \tfrac{3}{7}, \tfrac{4}{7}, \tfrac{5}{7}, \tfrac{6}{7}, 1\right).$$

Thus the two types of equilibria imply different predictions regarding payoffs and deal rates.

2.2 Experiments

Camerer et al. [2019] developed a novel experimental paradigm of dynamic bargaining that allows both parties to communicate offers whenever they please, while keeping their behavior tractable. This experiment was first conducted by Camerer et al. [2019] (Experiment 1), which is the baseline treatment. We also report results from a follow-up experiment with same design but with different treatments (Experiment 2). In this section, we first introduce the experimental design and then the treatments.

2.2.1 Design

The experimental design is a continuous-time bargaining game with one-sided private information. At the beginning of the experiment, subjects are assigned to one of the two roles: the informed player or the uninformed player. Players' roles are fixed for the session's 120 bargaining rounds.

In each round, each informed player is randomly matched with an uninformed player to bargain over a pie with a size unknown to the uninformed player. The pie size is an integer from 1 to 6, i.e. $\pi \in \{\$1, 2, 3, 4, 5, 6\}$ and drawn from a commonly known discrete uniform distribution. The informed player would know the pie size for that round after the draw is made.

Each pair bargained over the uninformed player's payoff, denoted by w. Both players communicate their offers, in multiples of \$0.1[10], using a mouse click on a graphic interface which was

[10] In Camerer et al. [2019] (Experiment 1), they set the resolution to be in multiples of \$0.2, which is a compromise between \$0.1 (too fine a resolution for coordinating in a short game) and \$0.5 (a resolution that would not allow for testing the use of focal points, as every possible offer would be a half of an integer pie). However, the result in Experiment 1 shows the players are able to coordinate in such a short period, so we increase the resolution to be in multiples of \$0.1 in Experiment 2.

programmed with z-Tree software (Fischbacher [2007]). During the first two seconds, both players can decide their initial bargaining position, without seeing the opponent's position (Figure 1A). The initial cursor location is randomized.

After initial locations are set, the players start a 10-second bargaining round. They communicate the offers with mouse clicks (Figure 1B). As both players' positions match, a green vertical stripe would appear on the screen (Figure 1C), and this position would become the final deal if there is no change on the position in the following 1.5 seconds (or if the period ends, which ever came first)[11]. If no deal is reached within 10 seconds, both players earn nothing. After each round, the players would be notified their payoffs and the actual pie size (Figure 1D).

INSERT FIGURE 1 HERE

2.2.2 Experiment 1

Camerer et al. [2019] conduct eight experimental sessions, five in the Social Science Experimental Laboratory (SSEL) at Caltech and three in the California Social Science Experimental Laboratory (CASSEL) at UCLA. In the beginning of each session, subjects are randomly assigned to isolated computer workstations and are handed printed versions of the instructions (see Appendix D in Camerer et al. [2019]). The instructions are also read aloud by the experimenter (who is the same person in all sessions). All of the participants complete a short quiz to check their understanding of the task. Subjects play 15 practice rounds to become familiar with the game and the interactive interface before the actual play of 120 rounds. Participants' payoffs are based on their profits in a randomly chosen 15% of the rounds, plus a show-up fee of $5. Each session lasts between 70 and 90 minutes (including check-in, reading of instructions, experimental task, and payment).

2.2.3 Experiment 2

The follow-up experiment is conducted in the Taiwan Social Science Experimental Laboratory (TASSEL) at National Taiwan University. We conduct eight experimental sessions. Three sessions are female-informed sessions where female subjects take the role of informed players and played against uninformed male subjects. Another three sessions are male-informed sessions which have the opposite design to the female-informed sessions. In the female-informed and male-informed sessions, we require an equal number of male and female subjects. Subjects are only notified of this requirement when entering the experiment. Besides, we conduct one experienced session and one high-stakes session in order test whether the experimental results are robust to the experience and the stake. In the experienced session, we recruit subjects who have participated one of the six previous sessions. In the high-stakes session, we multiply the stakes by 5. Notice that there is no gender constraint in the experienced session and high-stakes session.

The experimental procedures are the same in Experiment 1 and Experiment 2. In Experiment 2, participants' payoffs are based on their profits in a randomly chosen 10% of the rounds, plus a show-up fee of NT$ 100. Payoffs in the experiments are converted into NT$ according to a preset exchange rate (1 ECU = NT$15) specified in the instructions. Notice that in the high-stakes session, the exchange rate is 1 ECU = NT$75 while the exchange rate is 1 ECU = NT$30 in the experienced session.

After 120 rounds of the bargaining game, we measure subjects' risk preferences and loss aversion by Dynamically Optimized Sequential Experimentation (DOSE) developed by Wang et al. [2010b]. In each round, subjects are asked to choose from 2 lotteries. Lottery 1 is a risky asset, while lottery 2 yields a fixed amount. There are 3 practice rounds and 40 paid rounds. At the end of the experiment, 12 rounds from the bargaining game and 1 round from DOSE would be drawn and play out. Before undergoing DOSE, all subjects have to rate their willingness to take risk on a scale of 0 to 10. (0 means not willing to take any risk at all, and 10 means willing to take any risk.) The evaluation would not affect the payoff. Each session lasts around 2.5 hours.

[11] In Experiment 1, the offers have to match for 1.5 seconds in order to make a deal. In other words, the latest time where the players' bids can match is $t = 8.5$ seconds.

3 Experimental Results

3.1 Basics

In this section, we focus on analyzing the deal rates across different treatments. See Camerer et al. [2019] and Appendix for further analysis on the payoffs and the bargaining dynamics.

Table 1 provides the summary statistics of average bargaining outcomes in different treatments. From this table, we can observe that the average bargaining outcomes are pretty similar in different treatments. The range of the average payoffs across treatments is less than \$0.1. Also, the range of the deal rates is only 5%. However, there are still some differences on the aggregate level. The experienced session has the lowest surplus loss and the highest information value, while the male-informed sessions have lowest surplus loss and the baseline treatment has the lowest information value. This observation implies the bargaining outcomes are generally robust from different treatments (on the aggregate level), but may vary conditional on different pies, resulting in the differences in surplus loss and information value.

Table 1: Summary Statistics for Different Treatments

Treatment	Baseline	Female	Male	Experienced	High-Stake
Informed Payoff[a]	2.01	2.08	2.09	2.10	2.04
	(0.03)	(0.02)	(0.06)	–	–
Uninformed Payoff[a]	1.49	1.42	1.41	1.40	1.46
	(0.03)	(0.02)	(0.06)	–	–
Deal Rate	0.61	0.66	0.62	0.66	0.65
	(0.03)	(0.02)	(0.02)	–	–
Surplus Loss[b]	1.13	1.02	1.18	0.96	1.11
	(0.08)	(0.09)	(0.09)	–	–
Information Value[c]	0.40	0.51	0.49	0.54	0.42
	(0.03)	(0.05)	(0.06)	–	–

Means and standard errors (which are shown in parentheses) are calculated by treating each session's mean as a single observation. Since there is only one session for experienced and high-stake treatment, the standard errors for these two treatments are not computable.
[a] Averages are calculated for deal games only.
[b] Surplus loss = the mean expected loss of pie due to strikes.
[c] Information value = the mean difference between the informed and uninformed payoffs.

In Figure 2 and Figure 3, we break down the deal rates into different pie sizes for different treatments. From the figures, we can see the pattern in all treatments that the deal rates are increasing with the pie size, which confirms the prediction in Lemma 1. As we further analyze the figures, we can observe that the deal rates in female-informed sessions and the experienced session are higher than the baseline sessions in all pies (except the largest pie). On the other hand, the deal rates in male-informed sessions and the high-stake session are higher than the baseline in small pies ($\pi \leq 3$), but lower in large pies ($\pi \geq 4$).

INSERT FIGURE 2 AND 3 HERE

In Appendix, we discuss further results from Experiment 2. These results include analysis of the bargaining dynamics (see Appendix A) and testing predictions in the Lemma 1 (see Appendix B). Basically, the results in Camerer et al. [2019] are replicated by Experiment 2. Besides the monotone increase of deal rates and payoffs, we also observe that the equal-split allocation is the most salient focal point. Regarding the dynamics, we can observe the informed players' offers increase, and the uninformed players' demands decrease with time (within a trial). There is also a strong deadline effect—most of the deals are reached close to the deadline. Lastly, we analyze the differences in equilibrium selections using regression.

3.2 Outcome Prediction via Machine Learning

The unstructured paradigm established by Camerer et al. [2019] records, in addition to initial demands and offers, a large amount of bargaining process data that may be used to predict disagreements before the deadline has arrived. For instance, suppose that at the five-second mark, neither player has changed her offer for more than three seconds. This mutual stubbornness might be associated with an eventual strike. We consider a large number of such candidate observable features in search of a small set that is predictive, using cross-validation (Stone [1974]) to control for overfitting. This machine learning approach has been used in many applications in computer science and neuroscience, and is beginning to be more widely used in economics (e.g. Krajbich et al. [2009]) and other social sciences (e.g. Dzyabura and Hauser [2011]).

In this paper, we treat Experiment 2 as the lockbox test for the predictive model built in Camerer et al. [2019]. Therefore, in this section we report the results from directly feeding the data from Experiment 2 into the model. First of all, we briefly introduce the algorithm here. We follow Camerer et al. [2019] to choose 35 behavioral features recorded during bargaining. Examples of features are the current difference between the offer and demand, the time since the last position change, and an indicator denoting which of the players had changed his or her position in the game first. The full list is in Camerer et al. [2019]. For each of the eight sessions in Experiment 2, we trained a model to classify trials into disagreements or deals, using the data of the remaining seven sessions, by estimating a logistic regression with a least absolute shrinkage and selection operator (LASSO) penalty (Tibshirani [1996]). By applying these trained models, we then made out-of-sample predictions of the binary bargaining outcomes for each of the experimental session.

To assess the predictive power of process data, we estimate three strike prediction models at eight different points in the bargaining process, separated by one-second intervals (i.e., $1, 2, \ldots, 8$ seconds after bargaining starts). One model relies only on the pie size, the second uses only process features, and the third uses both pie size and process features. For each time stamp, predictions are carried out using the following nested cross-validation procedure: For each of the eight sessions, we train a linear model to predict the outcome (a deal or a strike), by fitting a logistic LASSO regression using the seven other sessions. The tuning parameter, λ, is optimized via ten-fold cross-validation, performed within each training set. Finally, using that trained model, we conduct out-of-sample predictions for the holdout sessions.

To evaluate the results from three different models at once, we use the "receiver operating characteristic" (ROC) curves (Hanley and McNeil [1982], Bradley [1997]), which is a standard tool in signal detection theory, used for quantifying the performance of a binary classifier under different trade-offs between type I and type II errors. For a random classifier, the true positive and false positive rates are identical (the 45-degree line in Figure 4). A good classifier increases the true positive rate (moving up on the y axis) and also decreases the false positive rate (moving left on the x axis). The difference between the ROC and the 45-degree line, in the upper-left direction, also known as the "area under the curve" (AUC) is an index of how well the classifier does.

INSERT FIGURE 4 AND 5 HERE

In Figure 4, we plot the ROC curve at $t = 2, 5, 7$ seconds for both Experiment 1 and 2. The ROC analysis shows that process data do better than random for every time stamp in both experiments. As we plot the AUC for both experiments together in Figure 5, we first observe there is some reduction in fit from old dataset to new dataset. However, except for the pie size only model, the other two models' fitness increases all the time and they even perform better in Experiment 2 than Experiment 1 at $t = 8$ seconds. This suggests that the predictive model built on the data from the original data is robust from different datasets.

While the pattern of AUC is similar in Experiment 1 and 2, there are still some subtle differences. In Experiment 1, we can observe that the model with pie and process feature always has the best predictive power and the other two models are not so distinguishable in later seconds. On the other hand, even though the model with pie and process feature is the best model among three, its predictive power is not significantly better than the model with process feature only.

To further investigate which behavioral process features predict strikes, we follow Camerer et al. [2019] to use a "post-LASSO" procedure proposed by Belloni et al. [2013, 2012]. Figure 6 summarizes the marginal effects of all process features (z-scored for every time point) in both experiments. From the figure, we can observe that the general pattern in Experiment 2 is consistent with the original result. The current informed player's offer (positively correlated with a deal) and the current difference between the players' bargaining positions (positively correlated with a strike) are the most predictive process features. On the other hand, one surprising finding in Camerer et al. [2019] is that the initial bargaining positions contain predictive information regarding the chance of reaching a deal, even as the deadline approaches, and even after controlling for current offers. Here we find the effect of initial positions is even stronger in Experiment 2. Furthermore, we also find the negative interaction between initial offer and initial demand and the negative interaction between initial and current offer, which again confirms the argument in Camerer et al. [2019].

INSERT FIGURE 6 HERE

4 Conclusion

In this chapter we described some new evidence of how machine learning can be used to predict whether bargaining trials end in agreements or in disagreements ("strikes").

Experimentation in economics on bargaining mostly abandoned unstructured bargaining because there is too little experimental control over all the things that bargainers can do. Unstructured bargaining seems to hand over the reins of endogenous "treatments" to the experimental subjects.

If the goal is prediction rather than theory-testing, however, having a large amount of data is terrific. For machine learning applications, there is (almost) no such thing as too much data.

Furthermore, a kind of theory-testing can still be done in a machine learning framework. In our example, the revelation principle, along with other restrictions, still delivers predictions about what will happen in equilibrium which are highly independent of the unstructured behavior. Everything depends on pie sizes. A lean predictive "machine" using only pie sizes should therefore predict as accurately as one with many more features.

Our main finding, using National Taiwan University subjects and some small design changes, is a close replication of earlier results using Caltech subjects. Agreements are often equal splits, even though the exact pie size is only known to one side. Deal rates do increase with pie size, but there is a lot of inefficiency—deal rates are too low—compared to revelation principle predictions. However, theory also predicts a break for uninformed offers for pies of $4-6 compared to lower pie amounts, and this break is evident in the data. There are some experience effects (deal rates go up across trials in an experimental session). One session with twice-experienced subjects—repeating the entire experimental session—did not produce results much closer to equilibrium (to our surprise).

There are only weak effects of gender. When females are informed, the deal rate is a bit higher and uninformed (males) get a little less, but the evidence is not statistically strong. While gender effects in bargaining are interesting, a lot more statistical power is probably needed to establish whether there are differences are not [12].

Finally, we hope these data and method inspire other experimenters from a range of social sciences to measure a lot more about what goes in the bargainers' bodies and brains, and results from their typing or talking, on during bargaining. For example, Forsythe et al. [1991a] At the time, methods of analyzing natural language processing (NLP) were so primitive that they did not do any sophisticated analysis of those rich data. What they wrote at the time was:

> Because of the unconstrained messages which pass between the players, our bargaining game is too complicated to allow a detailed strategic analysis. However, by the revelation principle, any Nash equilibrium of this game is equivalent to some direct mechanism, which specifies whether a strike should occur and how much each bargainer should get as a function of the informed bargainer's announcement of the size of the pie.

That is, While they allowed messages and recorded them, they did not analyze them at all because the resulting game—treating messages as strategy choices—is too complicated to solve. Using the messages as data in machine learning does not test a theory either, but it provides preliminary evidence of how features of messages influence agreement rates. Such evidence could provide inspiration for theory.

It is also notable that recording messages is very easy technically. NLP is one area of machine learning which is now hugely successful and improving by leaps and bounds every year. In general, machine learning methods are hungry for any such choice process data. And now we know what to do with them.

[12] And of course, gender differences are likely to vary wildly across the globe, so a serious attempt to understand such differences must look at the influences of developmental life cycle, biological factors such as hormones, and cultural variation.

References

D. R. Ames and M. F. Mason. Tandem anchoring: Informational and politeness effects of range offers in social exchange. *Journal of personality and social psychology*, 108(2):254, 2015.

L. M. Ausubel and R. J. Deneckere. Efficient sequential bargaining. *The Review of Economic Studies*, 60(2):435–461, 1993.

L. M. Ausubel, P. Cramton, and R. J. Deneckere. Bargaining with incomplete information. *Handbook of game theory with economic applications*, 3:1897–1945, 2002.

N. Bardsley, J. Mehta, C. Starmer, and R. Sugden. Explaining focal points: cognitive hierarchy theory versus team reasoning. *The Economic Journal*, 120(543):40–79, 2009.

A. Belloni, D. Chen, V. Chernozhukov, and C. Hansen. Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429, 2012.

A. Belloni, V. Chernozhukov, et al. Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547, 2013.

K. Binmore and L. Samuelson. The evolution of focal points. *Games and Economic Behavior*, 55 (1):21–42, 2006.

M. Blanco, D. Engelmann, and H. T. Normann. A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, 72(2):321–338, 2011.

G. E. Bolton and A. Ockenfels. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *American Economic Review*, 96(5):1906–1911, December 2006.

A. P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.

M. Buelens and D. Van Poucke. Determinants of a negotiator's initial opening offer. *Journal of Business and Psychology*, 19(1):23–35, 2004.

H. Cai and J. T.-Y. Wang. Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1):7–36, 2006.

C. F. Camerer, E. Johnson, T. Rymon, and S. Sen. Cognition and framing in sequential bargaining for gains and losses. *Frontiers of game theory*, 104:27–47, 1993.

C. F. Camerer, G. Nave, and A. Smith. Dynamic unstructured bargaining with private information: theory, experiment, and outcome prediction via machine learning. *Management Science*, 65(4): 1867–1890, 2019.

G. Charness and M. Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

E. Chaussees. Infinite-horizon models of bargaining with one-sided incomplete information. *Game-theoretic models of bargaining, Cambridge University Press, West Nyack, NY*, pages 73–98, 1985.

P. C. Cramton. Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *The Review of Economic Studies*, 51(4):579–593, 1984.

V. P. Crawford. Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *American Economic Review*, 93(1):133–149, 2003.

J. Créemer and R. P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53:345–361, 1985.

J. Cremer and R. P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.

R. Croson, T. Boles, and J. Murnighan. Cheap talk in bargaining experiments: Lying and threats in ultimatum games. *Journal of Economic Behavior & Organization*, 51(2):142–159, 2003.

T. E. Daniel, D. A. Seale, and A. Rapoport. Strategic play and adaptive learning in the sealed-bid bargaining mechanism. *Journal of Mathematical Psychology*, 42(2-3):133–166, 1998.

A. De Bruyn and G. E. Bolton. Estimating the influence of fairness on bargaining behavior. *Management Science*, 54(10):1774–1791, 2008.

R. Durante, L. Putterman, and J. Weele. Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086, 2014.

D. Dzyabura and J. R. Hauser. Active machine learning for consideration heuristics. *Marketing Science*, 30(5):801–819, 2011.

S. El Harbi, I. Bekir, G. Grolleau, and A. Sutan. Efficiency, equality, positionality: What do people maximize? experimental vs. hypothetical evidence from tunisia. *Journal of Economic Psychology*, 47:77–84, 2015.

D. Engelmann and M. Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94(4):857–869, September 2004.

D. Engelmann and M. Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply. *American Economic Review*, 96(5):1918–1923, December 2006.

E. Fehr, M. Naef, and K. M. Schmidt. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *American Economic Review*, 96(5):1912–1917, December 2006.

U. Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

R. Forsythe, J. Kennan, and B. Sopher. Dividing a shrinking pie: an experimental study of strikes in bargaining games with complete information. *Research in experimental economics*, 4:223–68, 1991a.

R. Forsythe, J. Kennan, and B. Sopher. An experimental analysis of strikes in bargaining games with one-sided private information. *The American Economic Review*, pages 253–278, 1991b.

D. Friedman and R. Oprea. A continuous dilemma. *American Economic Review*, 102(1):337–63, 2012.

W. Fuchs and A. Skrzypacz. Bargaining with deadlines and private information. *American Economic Journal: Microeconomics*, 5(4):219–43, 2013.

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

F. Galeotti, M. Montero, and A. Poulsen. Efficiency versus equality in bargaining. *Journal of the European Economic Association*, 17(6):1941–1970, 2018.

A. D. Galinsky and T. Mussweiler. First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of personality and social psychology*, 81(4):657, 2001.

S. J. Grossman and M. Perry. Sequential bargaining under asymmetric information. *Journal of Economic Theory*, 39(1):120–154, 1986.

F. Gul and H. Sonnenschein. On delay in bargaining with one-sided uncertainty. *Econometrica: Journal of the Econometric Society*, pages 601–611, 1988.

W. Güth and E. Van Damme. Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *Journal of mathematical Psychology*, 42(2-3):227–247, 1998.

W. Güth, S. Huck, and P. Ockenfels. Two-level ultimatum bargaining with incomplete information: An experimental study. *The Economic Journal*, 106(436):593–604, 1996.

J. A. Hanley and B. J. McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.

S. Hargreaves Heap, D. Rojo Arjona, and R. Sugden. How portable is level-0 behavior? a test of level-k theory in games with non-neutral frames. *Econometrica*, 82(3):1133–1151, 2014.

D. K. Herreiner and C. Puppe. Equitable allocations in experimental bargaining games: Inequality a version versus efficiency. Technical report, Bonn econ discussion papers, 2004.

B. Holmström and R. B. Myerson. Efficient and durable decision rules with incomplete information. *Econometrica*, pages 1799–1819, 1983.

A. Isoni, A. Poulsen, R. Sugden, and K. Tsutsui. Focal points in tacit bargaining problems: Experimental evidence. *European Economic Review*, 59:167–188, 2013.

A. Isoni, A. Poulsen, R. Sugden, and K. Tsutsui. Efficiency, equality, and labeling: An experimental investigation of focal points in explicit bargaining. *American Economic Review*, 104(10):3256–87, 2014.

N. Jacquemet and A. Zylbersztejn. What drives failure to maximize payoffs in the lab? a test of the inequality aversion hypothesis. *Review of Economic Design*, 18(4):243–264, 2014.

M. C. Janssen. Rationalizing focal points. *Theory and Decision*, 50(2):119–148, 2001.

M. C. Janssen. On the strategic use of focal points in bargaining situations. *Journal of Economic Psychology*, 27(5):622–634, 2006.

E. J. Johnson, C. Camerer, S. Sen, and T. Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1): 16–47, 2002.

J. H. Kagel and K. W. Wolfe. Tests of fairness models based on equity considerations in a three-person ultimatum game. *Experimental Economics*, 4(3):203–219, 2001.

J. H. Kagel, C. Kim, and D. Moser. Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior*, 13(1):100–110, 1996.

E. Kalai, M. Smorodinsky, et al. Other solutions to nash's bargaining problem. *Econometrica*, 43 (3):513–518, 1975.

J. Kennan and R. Wilson. Can strategic bargaining models explain collective bargaini. *The American Economic Review*, 80(2):405, 1990.

J. Kennan and R. Wilson. Bargaining with private information. *Journal of Economic Literature*, 31(1):45–104, 1993.

I. Krajbich, C. Camerer, J. Ledyard, and A. Rangel. Using neural measures of economic value to solve the public goods free-rider problem. *Science*, 326(5952):596–599, 2009.

P. H. Kriss, R. Nagel, and R. A. Weber. Implicit vs. explicit deception in ultimatum games with incomplete information. *Journal of Economic Behavior & Organization*, 93:337–346, 2013.

H. Kristensen and T. Gärling. The effects of anchor points and reference points on negotiation process and outcome. *Organizational Behavior and Human Decision Processes*, 71(1):85–94, 1997.

A. Kritikos and F. Bolle. Distributional concerns: equity-or efficiency-oriented? *Economics Letters*, 73(3):333–338, 2001.

P.-H. Lin, A. L. Brown, T. Imai, J. T.-y. Wang, S. Wang, and C. Camerer. General economic principles of bargaining and trade: Evidence from 2,000 classroom experiments. *Available at SSRN 3250495*, 2018.

R. López-Pérez, Á. Pintér, and H. J. Kiss. Does payoff equity facilitate coordination? a test of schelling's conjecture. *Journal of Economic Behavior & Organization*, 117:209–222, 2015.

M. F. Mason, A. J. Lee, E. A. Wiley, and D. R. Ames. Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *Journal of Experimental Social Psychology*, 49(4):759–763, 2013.

R. P. McAfee and P. J. Reny. Correlated information and mecanism design. *Econometrica: Journal of the Econometric Society*, pages 395–421, 1992.

M. Mitzkewitz and R. Nagel. Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22(2):171–198, 1993.

R. B. Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, pages 61–73, 1979.

R. B. Myerson. Two-person bargaining problems with incomplete information. *Econometrica*, pages 461–487, 1984.

J. F. Nash Jr. The bargaining problem. *Econometrica: Journal of the Econometric Society*, pages 155–162, 1950.

R. V. Nydegger and G. Owen. Two-person bargaining: An experimental test of the nash axioms. *International Journal of game theory*, 3(4):239–249, 1974.

J. Ochs and A. E. Roth. An experimental study of sequential bargaining. *The American Economic Review*, pages 355–384, 1989.

D. G. Pruitt. *Negotiation behavior*. Academic Press, 2013.

R. Radner and A. Schotter. The sealed-bid mechanism: An experimental study. *Journal of Economic Theory*, 48(1):179–220, 1989.

A. Rapoport and M. A. Fuller. Bidding strategies in a bilateral monopoly with two-sided incomplete information. *Journal of Mathematical Psychology*, 39(2):179–196, 1995.

A. Rapoport, I. Erev, and R. Zwick. An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Science*, 41(3):377–394, 1995.

A. Rapoport, T. E. Daniel, and D. A. Seale. Reinforcement-based adaptive learning in asymmetric two-person bargaining with incomplete information. *Experimental Economics*, 1(3):221–253, 1998.

A. E. Roth. Toward a focal point theory of bargaining. *Game-Theoretic Models of Bargaining*, 259:265–67, 1985.

A. E. Roth and M. W. Malouf. Game-theoretic models and the role of information in bargaining. *Psychological review*, 86(6):574, 1979.

A. E. Roth and J. K. Murnighan. The role of information in bargaining: An experimental study. *Econometrica: Journal of the Econometric Society*, pages 1123–1142, 1982.

A. E. Roth, M. W. Malouf, and J. K. Murnighan. Sociological versus strategic factors in bargaining. *Journal of Economic Behavior & Organization*, 2(2):153–177, 1981.

A. E. Roth, J. K. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *The American Economic Review*, 78(4):806–823, 1988.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109, 1982.

A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica: Journal of the Econometric Society*, pages 1151–1172, 1985.

T. Schelling. The strategy of conflict harvard university press cambridge, 1960.

J. Srivastava. The role of inferences in sequential bargaining with one-sided incomplete information: Some experimental evidence. *Organizational behavior and human decision processes*, 85(1):166–187, 2001.

I. Ståhl. Bargaining theory. *PhD diss. Economic Research Institute at the Stockholm School of Economics (EFI)*, 1972.

M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133, 1974.

L. L. Thompson, J. Wang, and B. C. Gunia. Negotiation. *Annual review of psychology*, 61:491–515, 2010.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

K. Valley, L. Thompson, R. Gibbons, and M. H. Bazerman. How communication improves efficiency in bargaining games. *Games and Economic Behavior*, 38(1):127–155, 2002.

D. Van Poucke and M. Buelens. Predicting the outcome of a two-party price negotiation: Contribution of reservation price, aspiration price and opening offer. *Journal of Economic Psychology*, 23(1):67–76, 2002.

J. T.-y. Wang, M. Spezio, and C. F. Camerer. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, 100(3):984–1007, 2010a.

S. W. Wang, M. Filiba, and C. F. Camerer. Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters. *Manuscript submitted for publication*, 2010b.

S. B. White and M. A. Neale. The role of negotiator aspirations and settlement expectancies in bargaining outcomes. *Organizational Behavior and Human Decision Processes*, 57(2):303–317, 1994.

S. B. White, K. L. Valley, M. H. Bazerman, M. A. Neale, and S. R. Peck. Alternative models of price behavior in dyadic negotiations: Market prices, reservation prices, and negotiator aspirations. *Organizational Behavior and Human Decision Processes*, 57(3):430–447, 1994.

G. Yukl. Effects of the opponent's initial offer, concession magnitude and concession frequency on bargaining behavior. *Journal of Personality and Social Psychology*, 30(3):323, 1974.

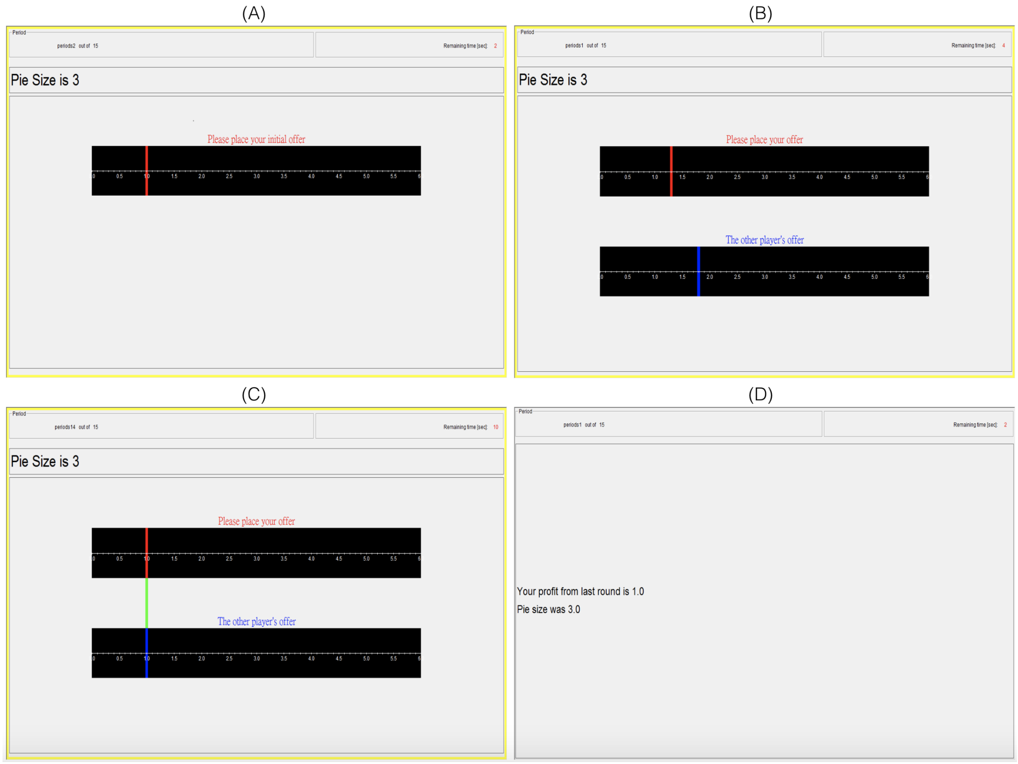

Figure 1: (A) Initial offer screen: in the first two seconds of bargaining, both players can set their initial positions without revealing to the opponent. The pie size is on the top left corner and it only appears on informed player's screen. (B) Players communicate their offers using mouse click on the interface. (C) When two players' positions match, the green vertical stripe appears and this would be the deal if there is no change in the following 1.5 seconds. (D) After the bargaining round, both players would be notified about their payoffs and the pie size.

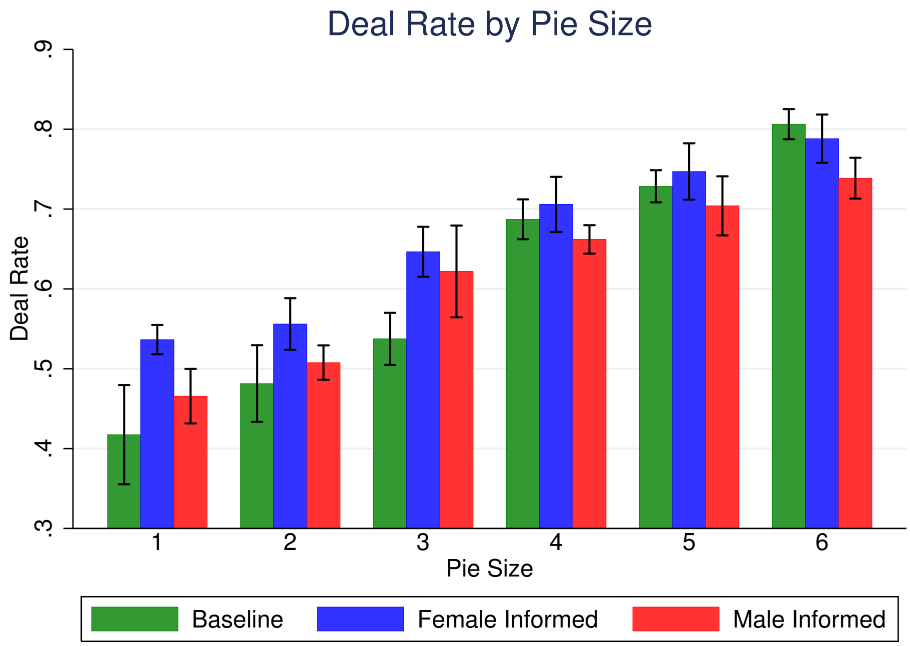

Figure 2: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for female-informed sessions and male-informed sessions, respectively. The standard errors (overlaid on the bars) are calculated at the session level.

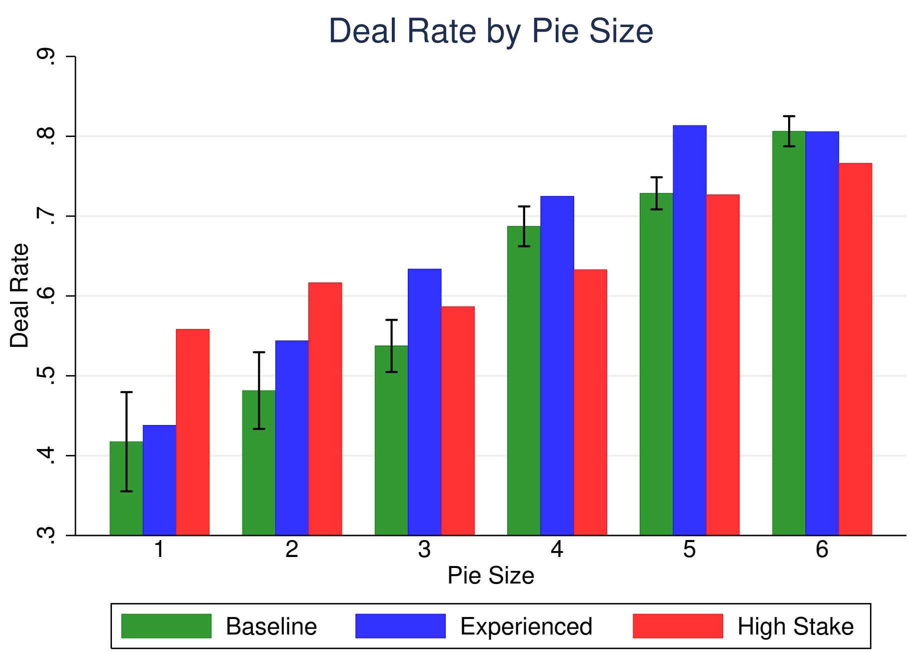

Figure 3: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for the experienced sessions and high-stake sessions, respectively. The standard errors (overlaid on the bars) are calculated at the session level. Since there is only one session for experienced and high-stake treatment, the standard errors are not computable for these two treatments.

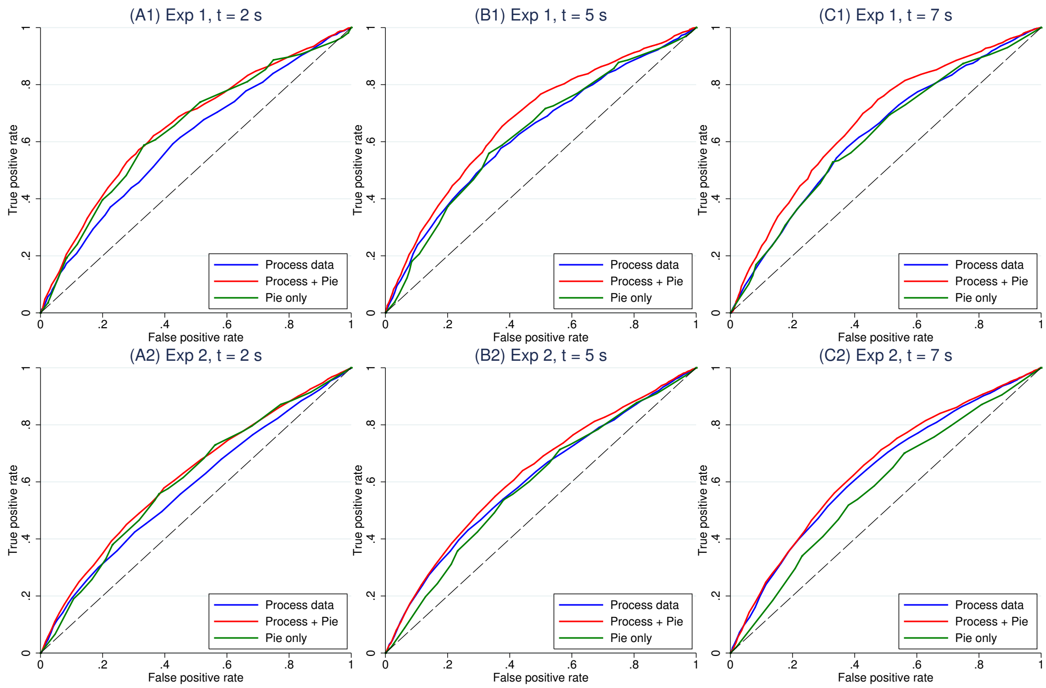

Figure 4: Receiver operating characteristic (ROC) for predicting disagreements, two, five and seven seconds into the bargaining game. The dashed lines represent the false and true positive rates of a random classifier. (A1–C1) show the data from Camerer et al. [2019] (Experiment 1) and (A2–C2) plot the result from Experiment 2.

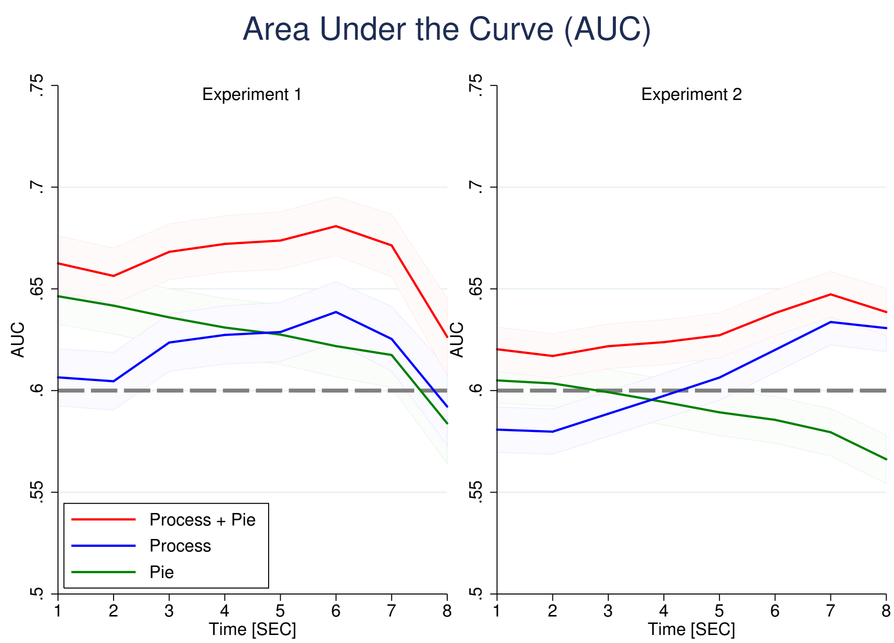

Figure 5: Area under the curve (AUC) of disagreements classifiers using process data, pie size, and the two combined. Note that the classifier's input included only trials that were still in progress (when a deal has not yet been achieved), and excluded trials in which the offers and demand were equal at the relevant time stamp. The left figure is the original result from Camerer et al. [2019] (Experiment 1) and the right one is the result from Experiment 2.

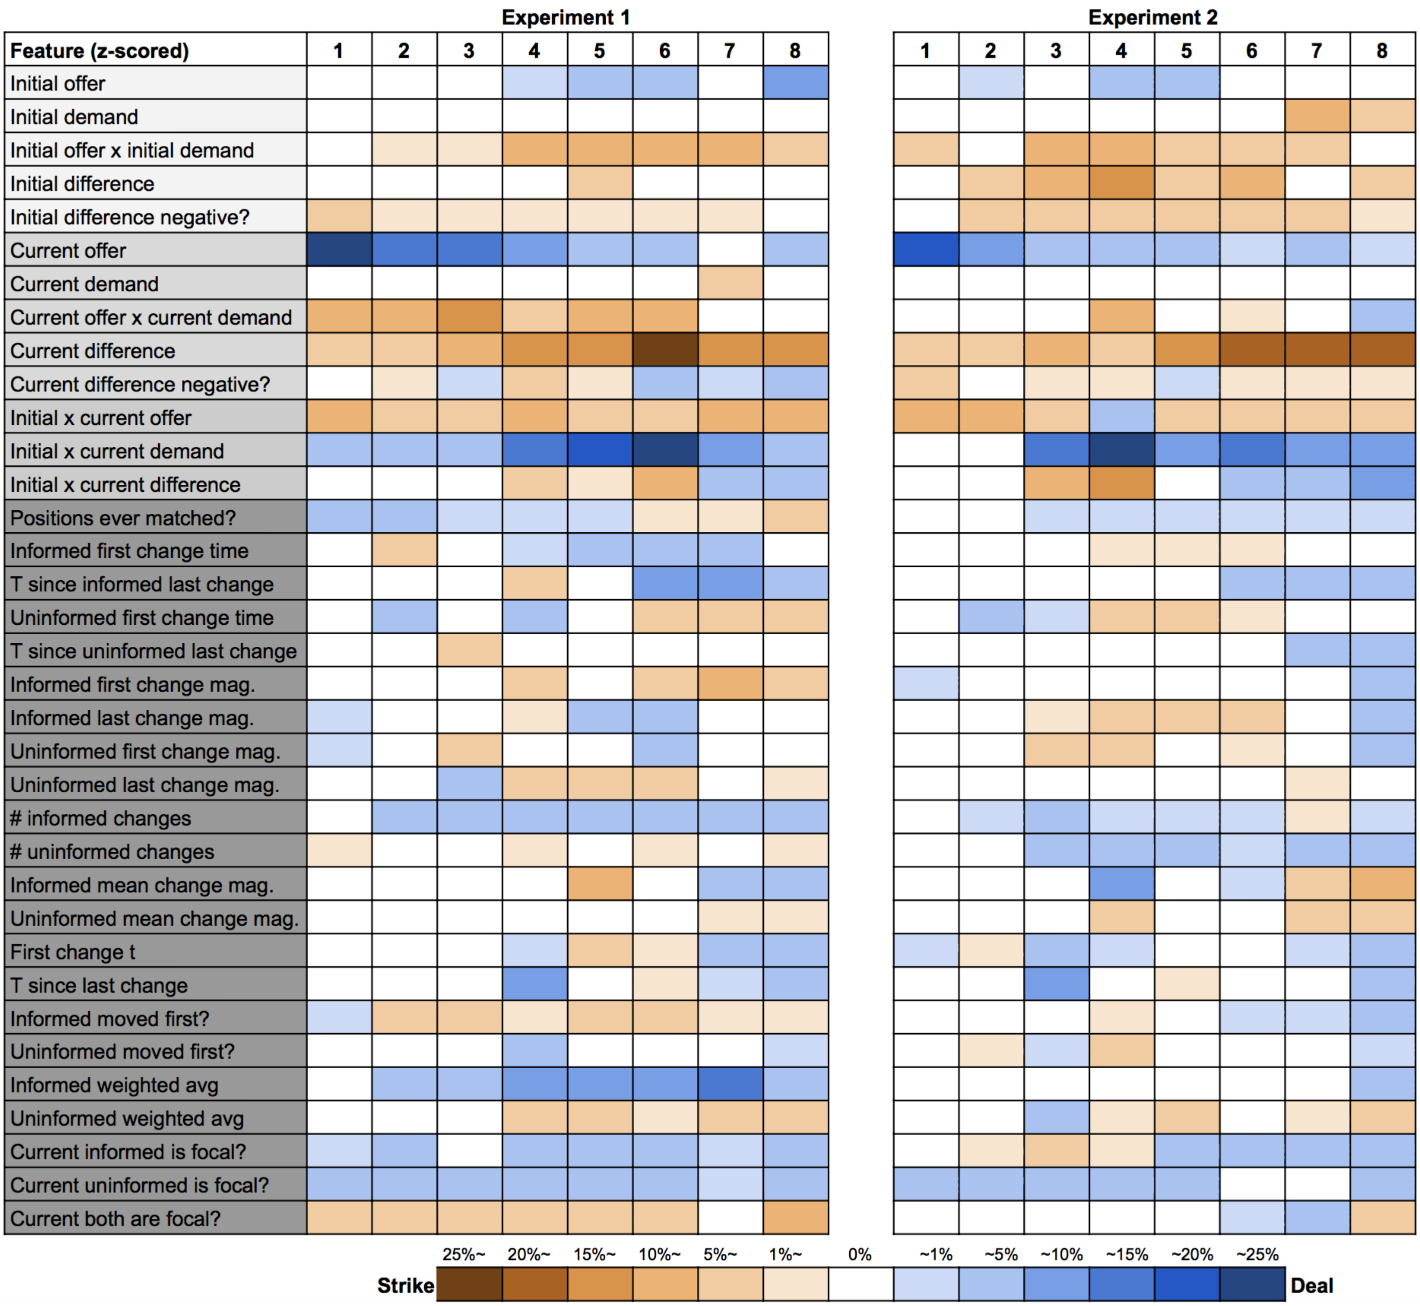

Figure 6: Bargaining Process Features Selected by the Classifier for Outcome Prediction (Deal= 1) and Their Estimated Marginal Effects. The left panel is the result from Camerer et al. [2019] (Experiment 1) and the right panel is the result from Experiment 2. The pie sizes are excluded.

A Appendix: Basics for Experiment 2

Contents

A.1 Summary Statistics for Different Pies and Treatments

In this section, we report the summary statistics for different pies and treatments in Table 2. From this table, we can see that no matter in which treatment the probability of disagreements decreases with pie size. Also, conditional on reaching a deal, players' payoffs increase with pie size. This pattern can be observed in Figure 7 and 8 as well.

In addition, surplus loss due to disagreements and the value of private information generally increase with pie size. Similar empirical patterns occur in all treatments. Therefore, we confirm that the first result in Camerer et al. [2019] is robust from different treatments.

Result 1. *Deal rates and payoffs are increasing with the pie size.*

Table 2: Summary Statistics for Different Pies and Treatments

Treatment	Pie Size	1	2	3	4	5	6	Mean
Baseline Treatment	Informed Payoff[a]	0.37	0.95	1.56	2.23	3.07	3.87	2.01
		(0.03)	(0.04)	(0.04)	(0.03)	(0.05)	(0.06)	
	Uninformed Payoff[a]	0.63	1.05	1.44	1.77	1.93	2.13	1.49
		(0.03)	(0.04)	(0.04)	(0.03)	(0.05)	(0.06)	
	Deal Rate	0.42	0.48	0.54	0.69	0.73	0.81	0.61
		(0.06)	(0.05)	(0.03)	(0.02)	(0.02)	(0.02)	
	Surplus Loss[b]	0.58	1.04	1.39	1.25	1.36	1.16	1.13
		(0.06)	(0.10)	(0.10)	(0.10)	(0.10)	(0.11)	
	Information Value[c]	-0.11	-0.05	0.05	0.31	0.83	1.39	0.40
		(0.03)	(0.03)	(0.04)	(0.04)	(0.07)	(0.10)	
Female Informed	Informed Payoff[a]	0.45	1.02	1.66	2.34	3.09	3.94	2.08
		(0.02)	(0.01)	(0.06)	(0.04)	(0.04)	(0.05)	
	Uninformed Payoff[a]	0.55	0.98	1.34	1.66	1.91	2.06	1.42
		(0.02)	(0.01)	(0.06)	(0.04)	(0.04)	(0.05)	
	Deal Rate	0.54	0.56	0.65	0.71	0.75	0.79	0.66
		(0.02)	(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	
	Surplus Loss[b]	0.46	0.89	1.06	1.18	1.26	1.27	1.02
		(0.02)	(0.06)	(0.09)	(0.14)	(0.18)	(0.18)	
	Information Value[c]	-0.05	0.03	0.21	0.48	0.87	1.49	0.51
		(0.02)	(0.01)	(0.08)	(0.07)	(0.06)	(0.13)	
Male Informed	Informed Payoff[a]	0.35	1.03	1.65	2.37	3.18	3.97	2.09
		(0.08)	(0.05)	(0.09)	(0.05)	(0.04)	(0.15)	
	Uninformed Payoff[a]	0.65	0.97	1.35	1.63	1.82	2.03	1.41
		(0.08)	(0.05)	(0.09)	(0.05)	(0.04)	(0.15)	
	Deal Rate	0.47	0.51	0.62	0.66	0.70	0.74	0.62
		(0.03)	(0.02)	(0.06)	(0.02)	(0.04)	(0.03)	
	Surplus Loss[b]	0.53	0.98	1.13	1.35	1.48	1.57	1.18
		(0.03)	(0.04)	(0.17)	(0.07)	(0.19)	(0.15)	
	Information Value[c]	-0.13	0.03	0.17	0.49	0.96	1.41	0.49
		(0.06)	(0.05)	(0.09)	(0.06)	(0.08)	(0.18)	
Experienced	Informed Payoff[a]	0.46	1.07	1.71	2.29	3.19	3.87	2.10
	Uninformed Payoff[a]	0.54	0.93	1.29	1.71	1.81	2.13	1.40
	Deal Rate	0.44	0.54	0.63	0.72	0.81	0.81	0.66
	Surplus Loss[b]	0.56	0.91	1.10	1.10	0.93	1.17	0.96

	Information Value[c]	-0.03	0.08	0.26	0.42	1.12	1.41	0.54
High Stake	Informed Payoff[a]	0.40	0.93	1.58	2.26	3.09	3.95	2.04
	Uninformed Payoff[a]	0.60	1.07	1.42	1.74	1.91	2.05	1.46
	Deal Rate	0.56	0.62	0.59	0.63	0.73	0.77	0.65
	Surplus Loss[b]	0.44	0.77	1.24	1.47	1.37	1.40	1.11
	Information Value[c]	-0.11	-0.08	0.09	0.33	0.86	1.46	0.42

Means and standard errors (which are shown in parentheses) are calculated by treating each session's mean as a single observation. Since there is only one session for experienced and high-stake treatment, the standard errors for these two treatments are not computable.
[a] Averages are calculated for deal games only.
[b] Surplus loss = the mean expected loss of pie due to strikes.
[c] Information value = the mean difference between the informed and uninformed payoffs.

INSERT FIGURE 7 AND 8 HERE

A.2 Focal Points

Figure 9 to Figure 12 show distributions of uninformed player's payoff conditional on reaching a deal in different treatments. See Camerer et al. [2019] for the distribution of the baseline. Focusing on the female-informed and male-informed sessions, as we pool the data from all pairs that reached a deal, we can observe 84.9% of the payoffs are 0.5, 1, 1.5, 2, 2.5 or 3, matching values that are exactly halves of possible pie sizes.

In addition, equal-split is the most common outcome. Among female-informed and male-informed treatments, 49.45% of the outcomes are equal-splitting when pie size is small or medium ($\pi \leq \$4$). This is even more common in female-informed sessions (55.98%) than male-informed sessions (42.94%). Notice that when pie size is small, the predicted payoffs of the efficient and equal-split equilibrium are the same (the top two rows in Figure 9 to Figure 12).

When pie size is large ($\pi \geq \$5$), the efficient equilibrium and equal-split equilibrium predict differently. The bottom row in Figure 9 to Figure 12 show that the mode is at \$2 and there are second modes at half the pie. Pooling the data from female-informed and male-informed treatments, 25.59% of payoffs are \$2. When pie size is \$5, 18.9% of the payoffs are at the equal-split equilibrium, while 13.0% of the payoffs are equal-splits when pie size is \$6. This empirical pattern is observed in all treatments and coincides with the finding in Camerer et al. [2019]. Thus, the second result in Camerer et al. [2019] is also replicated in Experiment 2.

INSERT FIGURE 9 TO 12 HERE

Result 2. *In all treatments, when pie size is small or medium ($\pi \leq \$4$), the modes of the uninformed player's payoff distributions equal to half the pie; when pie size is large ($\pi \geq \$5$), the modes are at \$2, though there are second modes at half the pie.*

A.3 Bargaining Dynamics

Figure 13 to Figure 16 show the dynamics of mean bargaining positions across different points of time. From these figures, we can see that no matter in which treatment, the informed player's offer increases with time, while the uninformed players' demand decreases. Therefore, we can first conclude that the empirical pattern found in Camerer et al. [2019] is robust to the gender difference, experience and stakes.

Result 3. *In all treatments, the informed player's offer increases, and the uninformed player's demand decreases with time.*

Although the general dynamic changes are similar, there are still some differences in information transmission among treatments. Comparing the female-informed and male-informed treatments, information transmission in male-informed sessions is better than that in female-informed sessions since in the latter the final positions are not clearly separated for pie size greater than \$4.

On the other hand, from Figure 15, we can see that the experienced treatment has the worst information transmission because the range of the average final position is only about \$0.6. By contrast, the high-stake session has the best information transmission when pie size is small, in which there is clear separation in the uninformed player's final position.

INSERT FIGURE 13 TO 16 HERE

A.4 Deadline Effect

Figure 17 and Figure 18 show the CDF of deals over time , which sharply increases as the deadline approached. This "endgame effect" is common to all pie sizes in both treatments. Basically, more than half of the deals are made in the last two seconds no matter in which session. Furthermore, deals are reached sooner when the pie is larger. This confirms the last empirical trend identified in Camerer et al. [2019]. One thing worth noticing here is that in the experienced session, players are more likely to reach a deal in the early seconds when pie size is large.

Result 4. *Most deals are made close to the deadline.*

INSERT FIGURE 17 AND 18 HERE

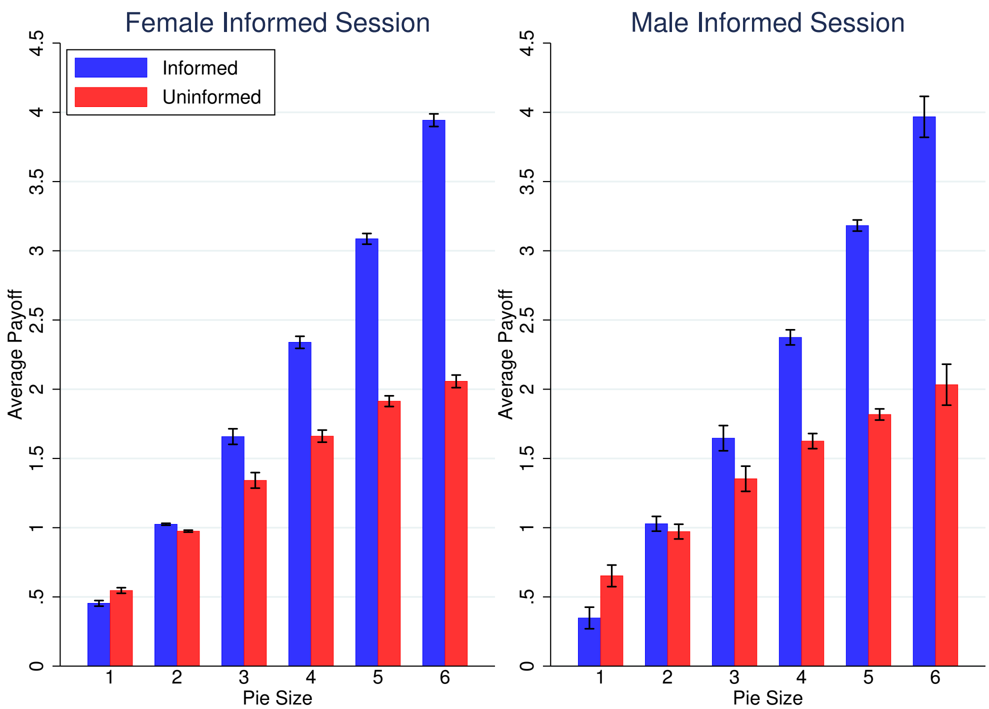

Figure 7: Mean payoffs by pie size and subject type, rounds ending in a deal. The blue bars represent the average payoff of informed players and the red bars are for uninformed players. The standard errors (overlaid on the bars) are calculated at the session level.

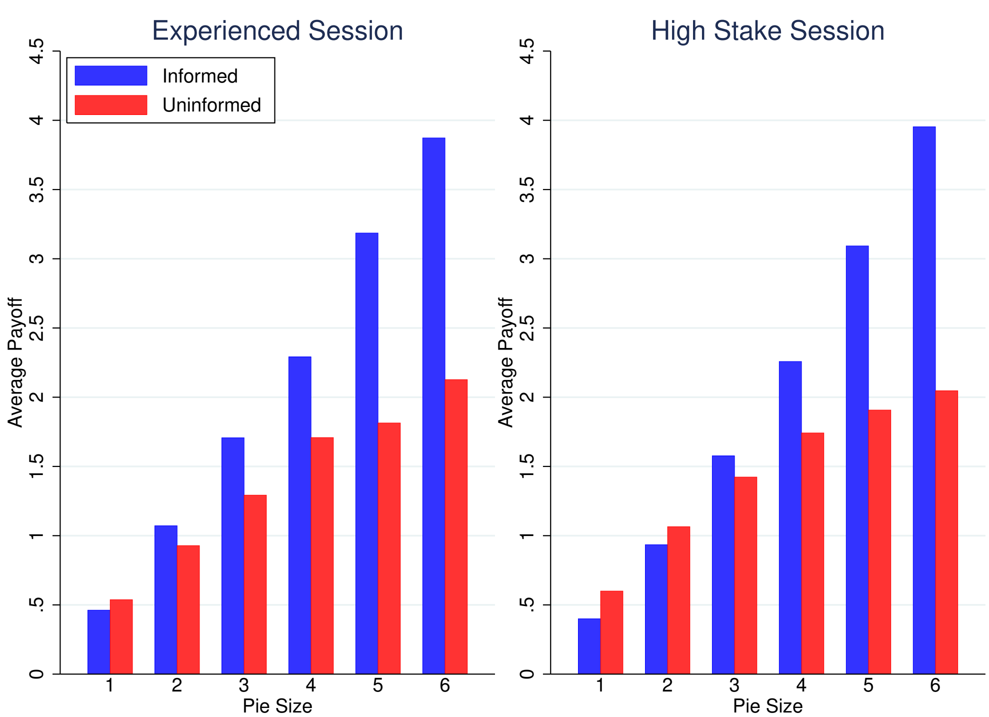

Figure 8: Mean payoffs by pie size and subject type, rounds ending in a deal. The blue bars represent the average payoff of informed players and the red bars are for uninformed players. Since there is only one session for each treatment, the standard errors can not be computed.

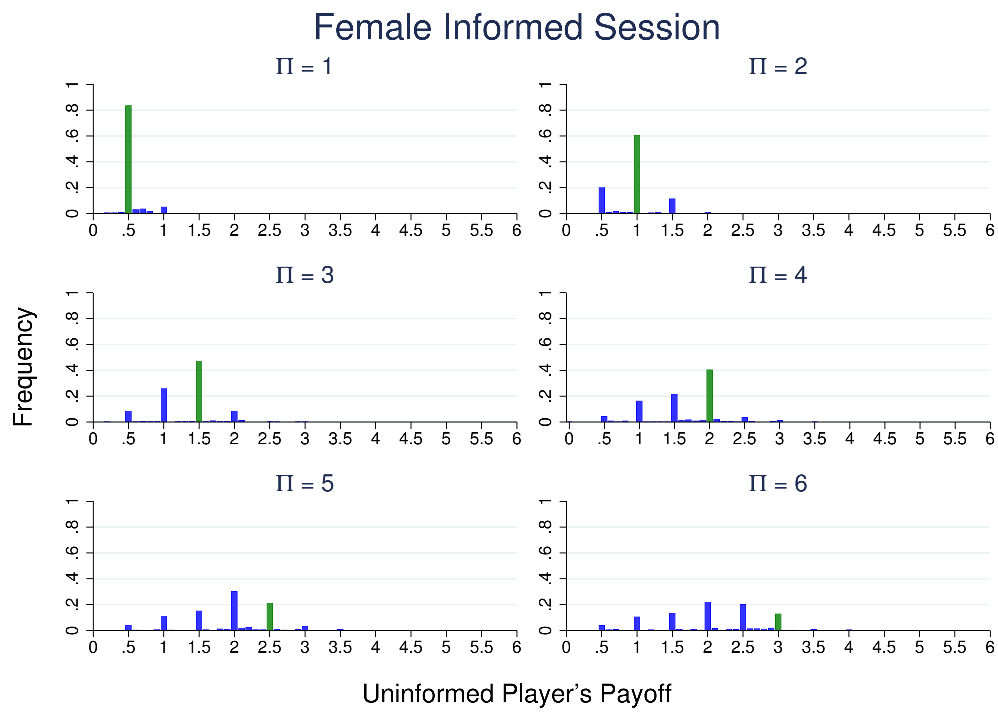

Figure 9: Uninformed player's payoff relative frequencies in female-informed sessions (conditional on reaching a deal). The bin size is 0.1 and the green bar locates half the pie in each distribution.

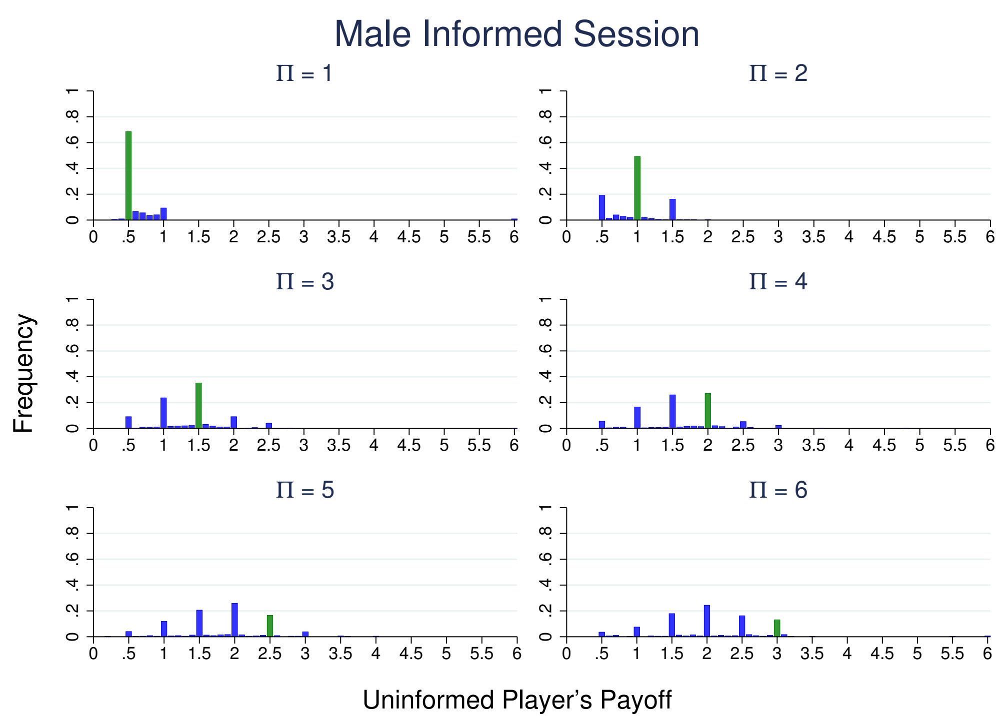

Figure 10: Uninformed player's payoff relative frequencies in male-informed sessions (conditional on reaching a deal). The bin size is 0.1 and the green bar locates half the pie in each distribution.

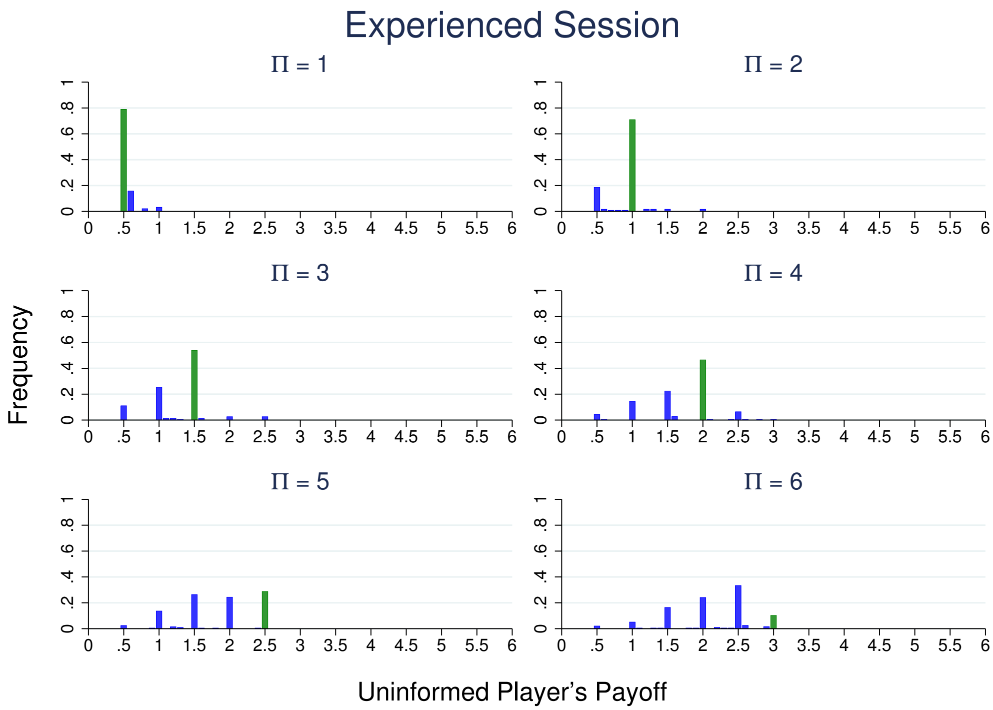

Figure 11: Uninformed player's payoff relative frequencies in the experienced session (conditional on reaching a deal). The bin size is 0.1 and the green bar locates the half of the pie in each distribution.

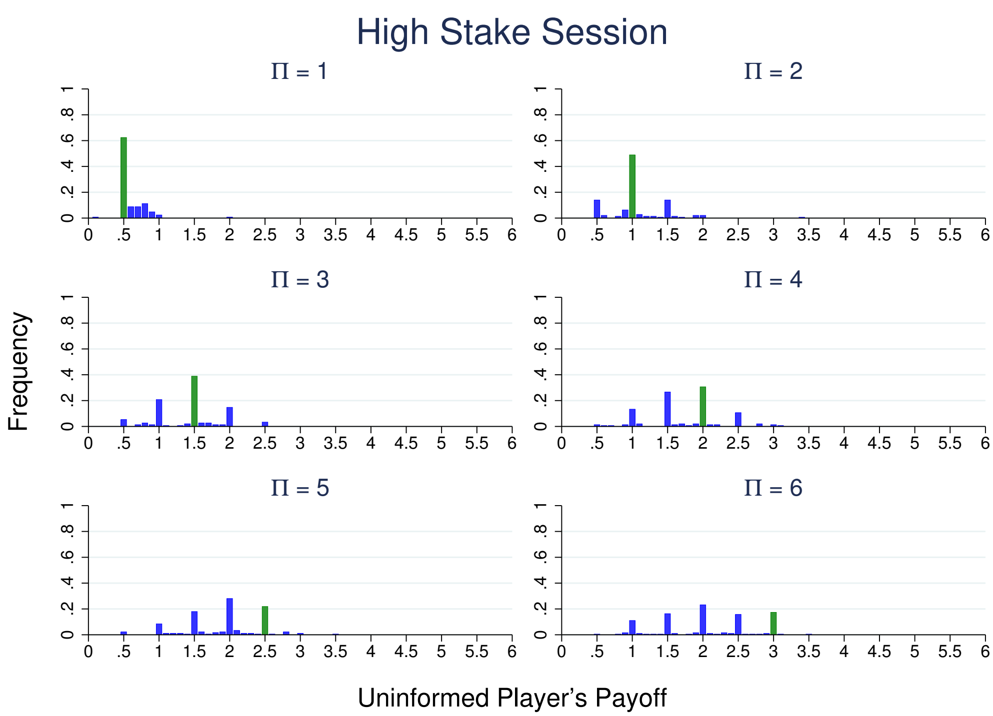

Figure 12: Uninformed player's payoff relative frequencies in the high stake session (conditional on reaching a deal). The bin size is 0.1 and the green bar locates the half of the pie in each distribution.

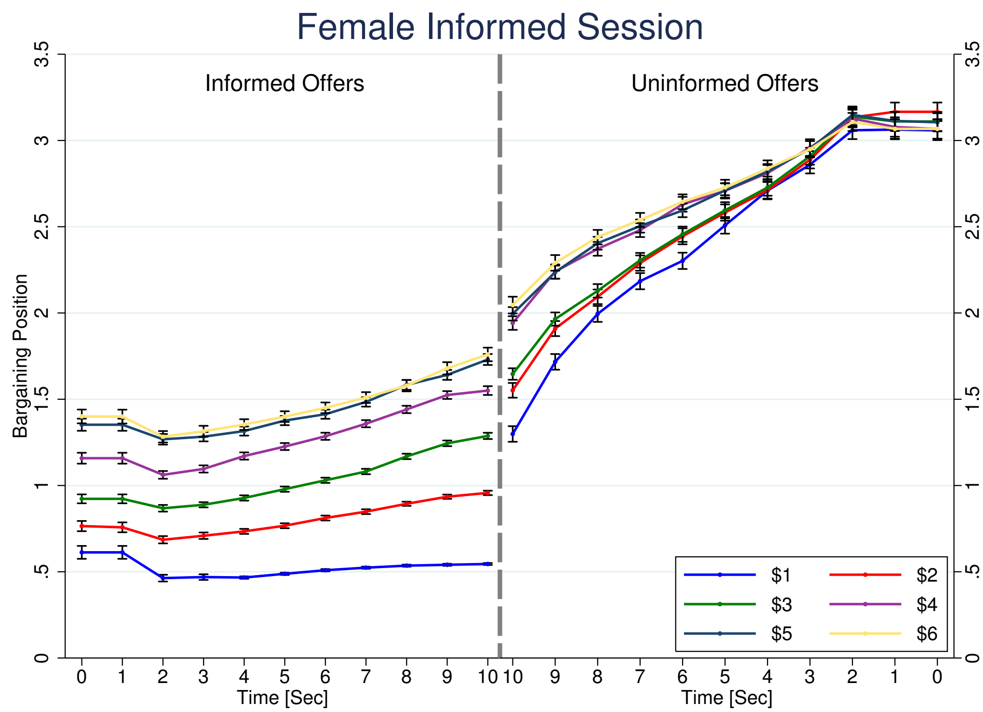

Figure 13: Mean Bargaining Position for All Pie Sizes in Female Informed Session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

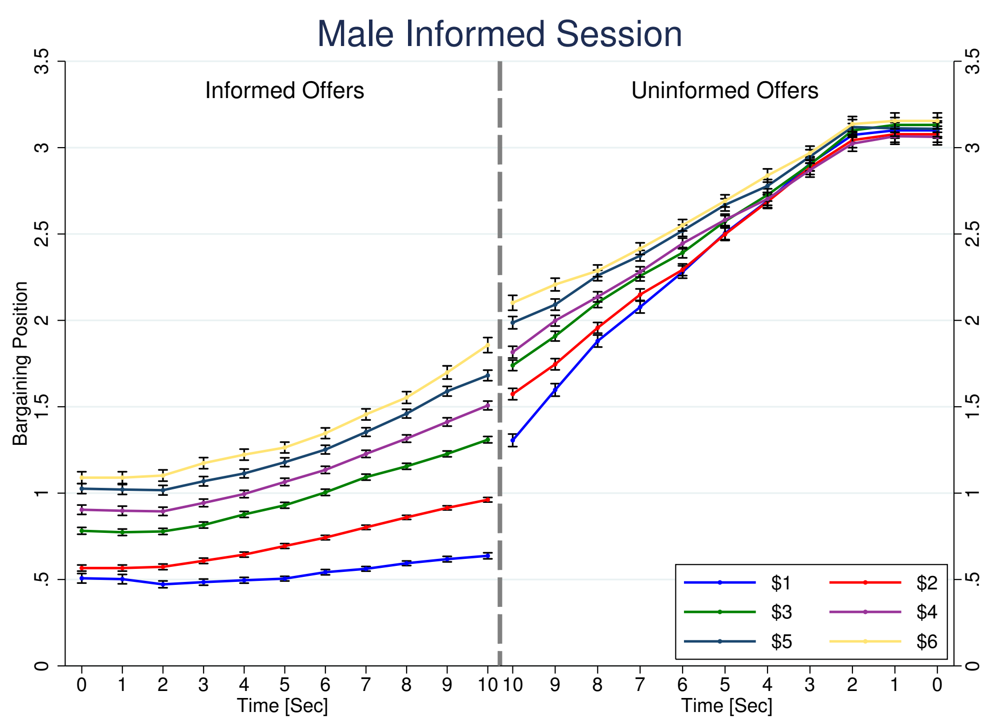

Figure 14: Mean Bargaining Position for All Pie Sizes in Male Informed Session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

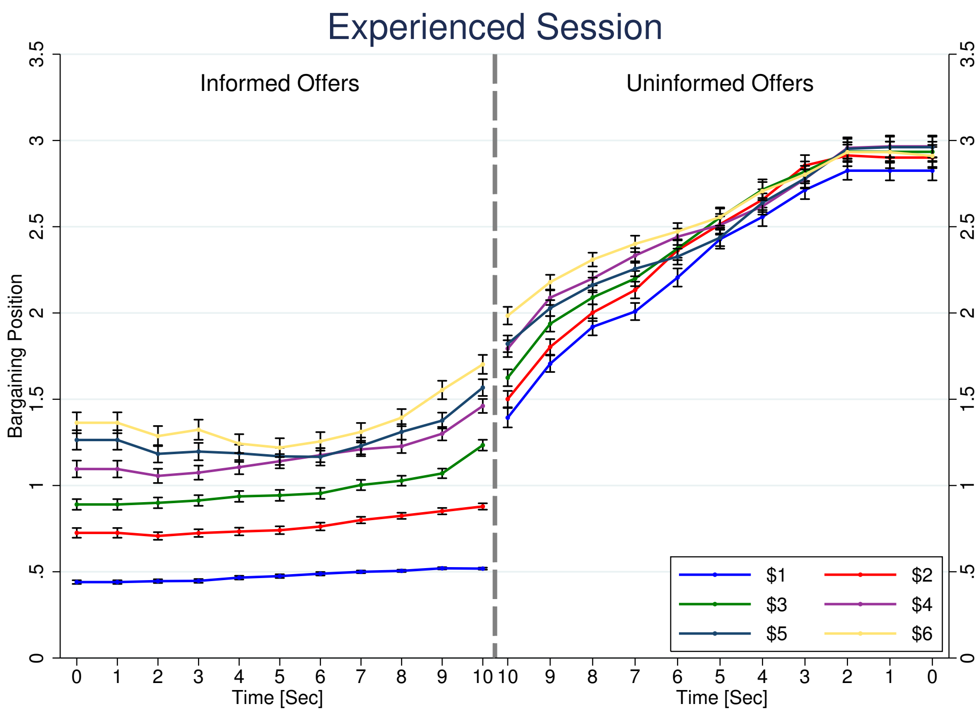

Figure 15: Mean Bargaining Position for All Pie Sizes in the experienced session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

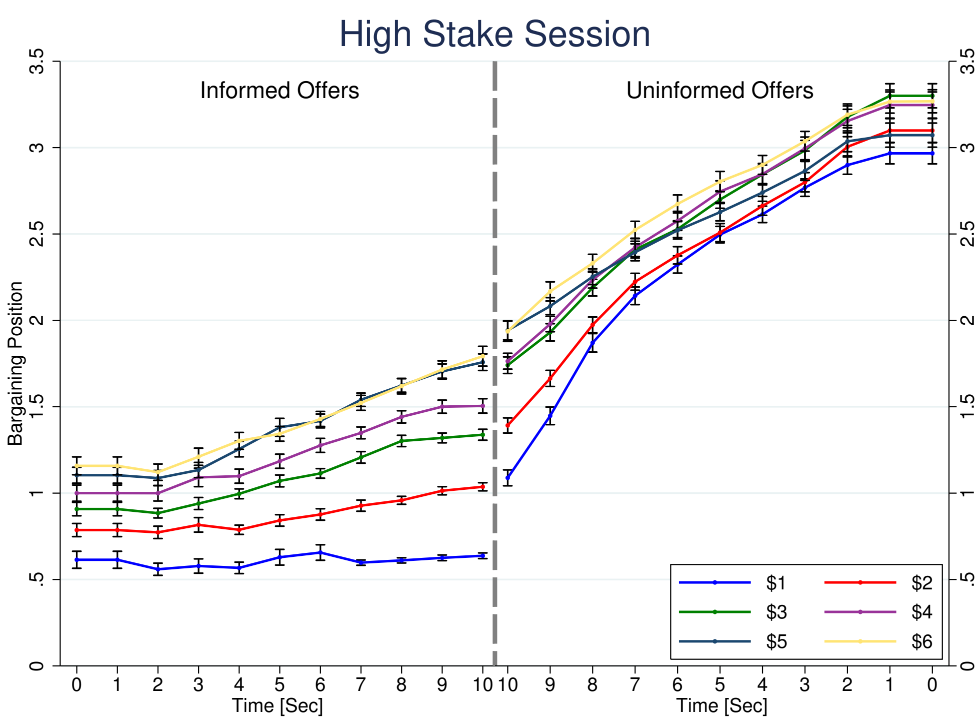

Figure 16: Mean Bargaining Position for All Pie Sizes in high stake session (All Rounds Pooled). The mean position is sampled at every second and standard errors are overlaid.

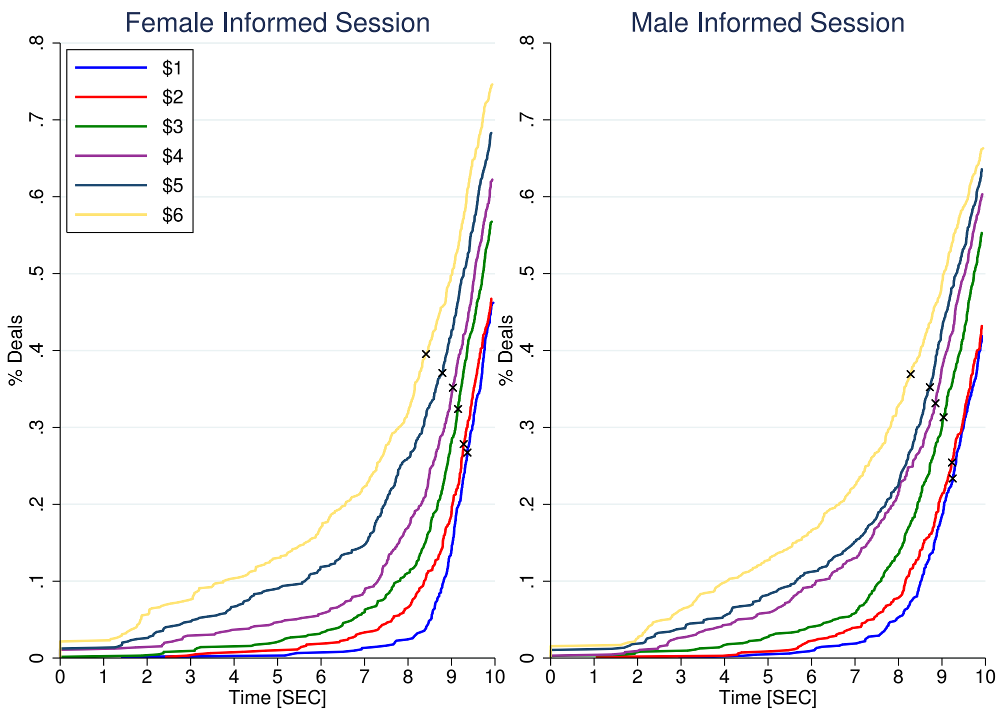

Figure 17: Cumulative Distribution of Deal Times by Pie Size in Female and Male Informed Sessions. Median deal times in different pie sizes are marked by a cross.

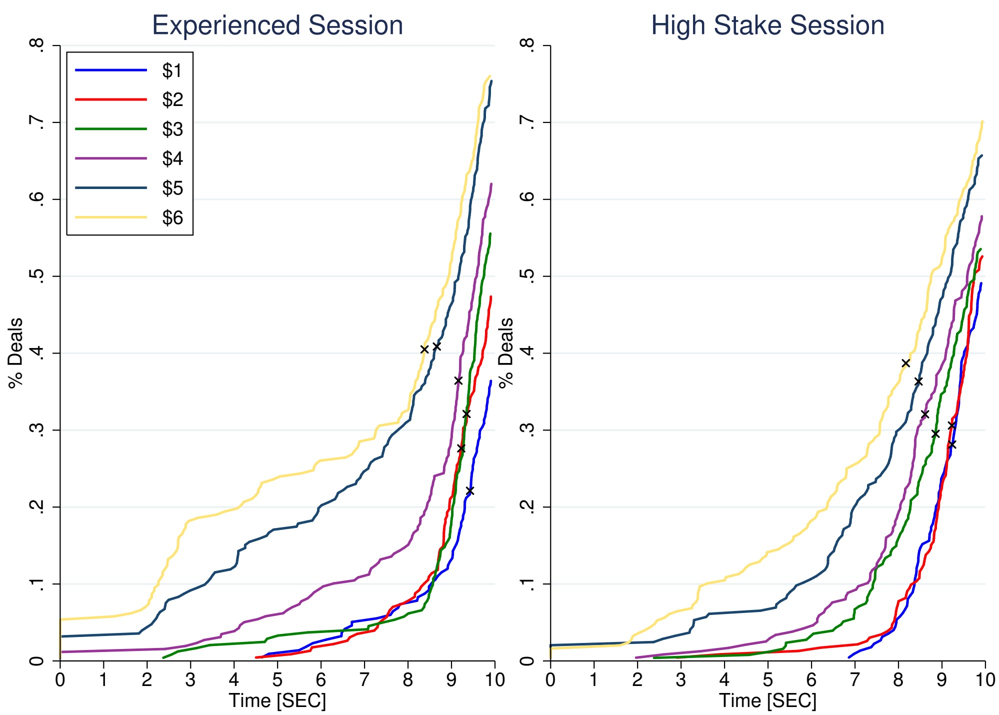

Figure 18: Cumulative Distribution of Deal Times by Pie Size in the experienced and high stake. Median deal times in different pie sizes are marked by a cross.

B Appendix: Comparison with Theoretical Predictions

Contents

B.1 Non-parametric Hypothesis Tests

In this section, we conduct non-parametric hypothesis tests on the predictions of Lemma 1. First of all, Lemma 1 predicts that deal rates increase with pie sizes. Table 2 and Figure 2 show this monotonic pattern in both female-informed and male-informed sessions. A non-parametric Wilcoxon-type test for trend in deal rates over pie sizes reject the null hypothesis of no trend (female-informed sessions: $z = 11.98, p < 0.001$, two-sided; male-informed sessions: $z = 12.63, p < 0.001$, two-sided).

The second prediction of Lemma 1 is that the uninformed player's payoff monotonically increases with pie size, conditional on reaching a deal. A non-parametric Kruskal-Wallis test (with adjustments for ties) rejects the null hypothesis that distributions of uninformed player's payoff are the same for each pie size in both treatments (female-informed sessions: $\chi^2(5) = 1284.828, p < 0.001$, two-sided; male-informed sessions: $\chi^2(5) = 1140.427, p < 0.001$, two-sided). In addition, a non-parametric Wilcoxon-type test rejects the null hypothesis of no trend in the payoffs across pie size conditional on reaching a deal (female-informed sessions: $z = 34.92, p < 0.001$, two-sided; male-informed sessions: $z = 32.93, p < 0.001$, two-sided).

Lastly, Lemma 1 also predicts that the uninformed player's payoff is identical for all pie sizes where the deal rate is 1. In particular, the efficient equilibrium predicts that the deal rates would be 1 for all $\pi \geq \$4$. However, the equal-split equilibrium predicts the deal rate is 1 only when pie size is 6. Table 1 and Figure 2 show that strikes are common even when pie size is 6 in both female-informed and male-informed sessions. A non-parametric Kruskal–Wallis test for equality of payoff distributions (with corrections for ties) rejects the null hypothesis that mean payoffs of the uninformed player, conditional on reaching a deal, are the same for pie sizes 4, 5, and 6 (female-informed Session: $\chi^2(2) = 95.623, p < 0.001$, two-sided; male-informed Sessions: $\chi^2(2) = 79.255, p < 0.001$, two-sided).

B.2 Regression Analyses

In this section, we perform two sets of linear regressions to test the theoretical predictions on deal rates and the uninformed player's payoffs conditional on reaching a deal.

B.2.1 Deal Rate Regression

First of all, we predict whether a deal is reached or not using the following specification:

$$y_{iust} = \alpha_0 + \alpha_1 \pi_{iust} + \alpha_2 d_{iust} (\pi_{iust} - 4) + \mathbb{X}_{iust}\beta + \epsilon_{iust}.$$

Here, y_{iust} is the dummy variable for reaching a deal between informed player i and uninformed player u in period t of session s. The spline term $d_{iust}(\pi_{iust} - 4)$ consists of two parts: the dummy variable d_{iust} for pie sizes greater or equal to 4 and $(\pi_{iust} - 4)$ as the increment of pie size beyond 4. The efficient equilibrium predicts that the deal rate is $\frac{2}{5}$ when pie size is 1, increasing by $\frac{1}{5}$ per unit in pie size. Hence, the deal rate would be 1 when pie size is greater or equal to 4. In contrast, the equal-split equilibrium predicts the deal rate to be $\frac{2}{7}$ when pie size is 1, increasing by $\frac{1}{7}$ per unit in pie size. Therefore, the efficient equilibrium predicts $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{5}, \frac{1}{5}, -\frac{1}{5})$, while the equal-split equilibrium predicts that $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{7}, \frac{1}{7}, 0)$.

INSERT TABLE 3 HERE

Table 3 reports the regression results. All models include standard errors clustered at the session level, to account for dependence in residuals within a particular session. Model A provides the baseline results predicting deal rate with pie size and the spline term, using pooled data from both female-informed and male-informed sessions. Models B includes session controls, and Model

C employs controls at the level of individual subject pairs (the smallest grouping available). Model D drops these controls and adds an indicator term controlling for female-informed sessions (Female = 1), as well as an indicator term for the last 60 rounds of the experiment (rounds 61–120) to capture the effect of experience. Model E adds interactions between female-informed/experience and pie size/spline. In model F, we further control for informed and uninformed player's risk preference ($ln(\rho_i), ln(\rho_u)$) and degree of loss aversion (λ_i, λ_u), as well as the initial offer and initial demand. Model G drops the female-informed session indicator and other controls but adds session-level controls. Model H adds back the controls for risk preference, loss aversion and the initial bargaining positions.

In all models, the coefficient on pie size is always significantly positive, ranging from 6.3% to 8.1%. This estimate is robust to different specification, which supports the prediction of the IC condition in Lemma 1. Yet, the slope coefficient on pie size is smaller than predicted from either equilibrium. The constant term from Model A is about 0.42 which is larger than predicted from either equilibrium. The spline term is negatively significant only when we control for interactions with experience, indicating that the players tend to coordinate on the efficient equilibrium at the beginning, but later move toward the equal-split equilibrium. In fact, experience plays an important role in reaching a deal. Players are roughly 8.5% more likely to reach a deal in later rounds. Also, the slope coefficient on the pie size is also significantly smaller in later rounds. On the other hand, it seems like whether the informed player is female is not a significant factor to predict deal rate. The dummy variable for female-informed sessions is only barely significant in Model E. Lastly, the initial position of the uninformed player would also affect deal rate.

B.2.2 Wage Regression

Similarly, we perform linear regressions to test the focal-split predictions regarding payoffs conditional on players reaching a deal. The efficient equilibrium predicts equal-splits when the pie is small, and the uninformed player's conditional payoff would be 2 when the pie size is 4 or greater. The equal-split equilibrium predicts a 50/50 split for all pie sizes. We test these predictions with the following specification:

$$w_{iust} = \alpha_0 + \alpha_1 \pi_{iust} + \alpha_2 d_{iust} (\pi_{iust} - 4) + \mathbb{X}_{iust}\beta + \epsilon_{iust},$$

where w_{iust} is the uninformed payoff (conditional on reaching a deal) agreed upon by informed player i and uninformed player u in period t of session s. The efficient equilibrium predicts that $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$, yet the equal-split equilibrium predicts $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, 0)$.

INSERT TABLE 4 HERE

In all specifications, the coefficient of pie size is positively significant, ranging from 0.28 to 0.35, but these point estimates are lower than the theoretical prediction. The spline term is roughly -0.16 (Model A to D), which is negatively significant. However, as we control for interactions between whether the informed player is female and experience, the spline term is no longer significant (Model E and F). This indicates that players in female-informed sessions are slightly more likely to coordinate on the efficient equilibrium than those in male-informed sessions. On the other hand, the constant term ranges from 0.24 to 0.34 (Model A to D), which is also lower than the theoretical prediction. Model F and H includes risk preferences, degrees of loss aversion and initial positions and find that initial offers play a significant role in determining final payoffs of the uninformed.

Table 3: Linear Regressions—Predictors of Deals

	Model A Coef./SE	Model B Coef./SE	Model C Coef./SE	Model D Coef./SE	Model E Coef./SE	Model F Coef./SE	Model G Coef./SE	Model H Coef./SE
Pie	0.0649***	0.0649***	0.0634***	0.0649***	0.0809***	0.0786***	0.0756***	0.0729***
	(0.0106)	(0.0106)	(0.0084)	(0.0106)	(0.0133)	(0.0131)	(0.0088)	(0.0075)
Spline at $\pi = \$4$	-0.0253	-0.0260	-0.0183	-0.0258	-0.0509**	-0.0491**	-0.0412**	-0.0394**
	(0.0169)	(0.0169)	(0.0139)	(0.0168)	(0.0148)	(0.0154)	(0.0137)	(0.0119)
Female				0.0451	0.0722*	0.0638		
				(0.0275)	(0.0343)	(0.0369)		
Rounds $61 - 120$				0.0270	0.0872***	0.0893***	0.0846***	0.0864***
				(0.0185)	(0.0180)	(0.0186)	(0.0184)	(0.0193)
Female × Pie					-0.0107	-0.0118		
					(0.0206)	(0.0201)		
Female × Spline					0.0211	0.0206		
					(0.0322)	(0.0324)		
Rd.61-120 × Pie					-0.0215**	-0.0213**	-0.0207**	-0.0204**
					(0.0065)	(0.0067)	(0.0062)	(0.0066)
Rd.61-120 × Spline					0.0298	0.0296	0.0297	0.0295
					(0.0208)	(0.0204)	(0.0208)	(0.0206)
$ln(\rho_i)$						-0.0585		-0.0579
						(0.0458)		(0.0422)
λ_i						0.0039		0.0017
						(0.0047)		(0.0061)
$ln(\rho_u)$						0.0231		-0.0000
						(0.0254)		(0.0150)
λ_u						0.0081		0.0077
						(0.0062)		(0.0066)
Initial demand						-0.0302*		-0.0327*
						(0.0135)		(0.0130)
Initial offer						0.0160		0.0144
						(0.0107)		(0.0112)
Constant	0.4245***	0.4246***	0.4261***	0.3893***	0.3457***	0.4122***	0.3812***	0.4549***
	(0.0238)	(0.0290)	(0.0233)	(0.0210)	(0.0260)	(0.0573)	(0.0259)	(0.0511)
Observations	7,920	7,920	7,920	7,920	7,920	7,920	7,920	7,920
AIC	10,553.52	10,499.13	8,623.804	10,532.82	10,527.62	10,455.81	10,494.74	10,416.78
BIC	10,574.45	10,513.09	8,637.758	10,567.71	10,562.51	10,490.70	10,529.63	10,451.67
Session Controls	No	Yes	No	No	No	No	Yes	Yes
Pair Controls	No	No	Yes	No	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.
$*p < 0.1$; $**p < 0.05$; $***p < 0.01$.

Table 4: Linear Regressions—Predictors of Uninformed Payoffs Conditional on Deal

	Model A Coef./SE	Model B Coef./SE	Model C Coef./SE	Model D Coef./SE	Model E Coef./SE	Model F Coef./SE	Model G Coef./SE	Model H Coef./SE
Pie	0.3527***	0.3518***	0.3299***	0.3534***	0.3208***	0.2831***	0.3402***	0.2938***
	(0.0180)	(0.0186)	(0.0194)	(0.0181)	(0.0218)	(0.0293)	(0.0173)	(0.0233)
Spline at $\pi = \$4$	-0.1641**	-0.1619**	-0.1534**	-0.1653**	-0.1192	-0.1246	-0.1390**	-0.1390**
	(0.0447)	(0.0455)	(0.0452)	(0.0455)	(0.0725)	(0.0785)	(0.0358)	(0.0445)
Female				0.0060	-0.1249	-0.1591		
				(0.0599)	(0.0868)	(0.1006)		
Rounds $61 - 120$				0.0338	-0.0217	0.0116	-0.0292	0.0031
				(0.0217)	(0.0289)	(0.0273)	(0.0335)	(0.0333)
Female × Pie					0.0424	0.0260		
					(0.0322)	(0.0379)		
Female × Spline					-0.0475	-0.0390		
					(0.0918)	(0.0931)		
Rd.61-120 × Pie					0.0215	0.0158	0.0240	0.0191
					(0.0178)	(0.0143)	(0.0196)	(0.0160)
Rd.61-120 × Spline					-0.0433	-0.0127	-0.0474	-0.0202
					(0.0483)	(0.0374)	(0.0496)	(0.0378)
$ln(\rho_i)$						-0.2145		-0.2226
						(0.1611)		(0.1592)
λ_i						-0.0134		-0.0307
						(0.0703)		(0.0771)
$ln(\rho_u)$						-0.0156		0.0150
						(0.0483)		(0.0337)
λ_u						0.0178**		0.0149
						(0.0067)		(0.0099)
Initial demand						0.0365*		0.0344*
						(0.0154)		(0.0139)
Initial offer						0.2857***		0.2887***
						(0.0482)		(0.0455)
Constant	0.2617***	0.2636***	0.3408***	0.2394**	0.3350***	0.0890	0.2773***	0.0539
	(0.0503)	(0.0489)	(0.0566)	(0.0832)	(0.0645)	(0.1541)	(0.0445)	(0.1280)
Observations	5,058	5,058	5,058	5,058	5,058	5,058	5,058	5,058
AIC	8,647.621	8,572.667	5,678.255	8,647.008	8,637.756	7,743.683	8,572.392	7,644.838
BIC	8,667.207	8,585.725	5,691.312	8,678.651	8,670.400	7,776.327	8,605.036	7,677.481
Session Controls	No	Yes	No	No	No	No	Yes	Yes
Pair Controls	No	No	Yes	No	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.
$*p < 0.1$; $**p < 0.05$; $***p < 0.01$.