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

October 15, 2020

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 is 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, Crémer and McLean, 1985, 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. We follow [Camerer et al., 2019] to provide a brief description of this history to help explain why we are enthusiastic about modern applications of machine learning.

1.1 A brief history of bargaining experiments

Prior to breakthroughs on theories of structured bargaining, most experiments were conducted using unstructured communication. Research mainly focused on process-free solution concepts lead by the Nash bargaining solution [Nash Jr, 1950], and their extensions [e.g. Kalai et al., 1975]. Numerous bargains [Nydegger and Owen, 1974, Roth and Malouf, 1979] led to an equal split of the amount of surplus available to share, which we refer to as "pie." When there are informational asymmetries, disagreements may occur because of coordination difficulties. Hence, Roth and colleagues have subsequent papers where players bargain over points which could be redeemed for different monetary values [Roth et al., 1981, Roth and Murnighan, 1982, Roth, 1985]. Theory

*The experimental data, codes and online appendix can be found on Open Science Framework https://osf.io/9j4cm/.

predicts 100% agreement in these games, but experimental results show that a modest percentage of trials (10-20%) end in disagreement, likely due to differences regarding which "focal points" are acceptable.[1] Roth et al. [1988] also drew attention to the "deadline effect," in which a large majority of agreements are made just before the (known) deadline.

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 [Ausubel et al., 2002]. In these theories and experiments, "structure" refers to theory clearly specifying the rules of how bargaining proceeds, which predictions of bargaining outcomes are sensitive to.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 [Camerer, 2003].

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. These experiments typical find *fewer* disagreements than predicted by theory in face-to-face and unstructured communication via message-passing (comparable to findings in "cheap talk" games where senders willingly reveal "too much" [Crawford, 2003, Cai and Wang, 2006, Wang et al., 2010].)

1.2 Returning to a Less Popular Route

Since the rise of structured bargaining theories, experimentation in economics on unstructured bargaining has all but disappeared. However, as shown by Karagözoğlu [2019], there are many good reasons to investigate the process of unstructured bargaining, such as the need to satisfy procedural justice. Hence, Camerer et al. [2019] return to this older venue, and explore unstructured bargaining with one-sided private information in laboratory experiment for the following reasons:

First, establishing more empirical regularities in naturally-occurring settings is a prerequisite for theorizing. Since most natural two-player bargaining settings have little penalties for deviating from structured conventions, studying unstructured bargaining is of particular importance. In particular, strategic behavior under continuous-time interaction [Friedman and Oprea, 2012] should be documented, as well as deadline effects [Roth et al., 1988, Gächter and Riedl, 2005] which are not predicted by most theories (though see Fuchs and Skrzypacz [2013]).

Second, theory can still be applied to make clear, interesting predictions even when bargaining is unstructured. For instance, clear predictions about unstructured bargaining can emerge, thanks to the wonderful "revelation principle" [Myerson, 1979, 1984]. This principle generates empirical predictions for non-cooperative equilibria based purely on the information structure, regardless of the bargaining protocol.

Third, unstructured bargaining generates very rich data during the bargaining process. Players are allowed to make offers at any time, retract them, etc. Natural language can be analyzed, perhaps including vocal properties in verbal communication [Capra, 2019]. 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," especially if one does not have a theory to say anything about them. Nevertheless, when outcomes are systematically influenced by process variables, these empirical regularities would challenge existing equilibrium theories and invite new developments in theory.

To this end, we focus on predicting which bargaining trials will results in deals and strikes, using a penalized regression approach from machine learning to select predictive 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 predicts strikes roughly as accurately as the pie sizes can, while combining both process and pie size makes even better predictions.

Since practical negotiation advice often consists of simple heuristics, process data could also be very useful to carefully test them experimentally [Pruitt, 2013]. In particular, initial offers have long been postulated to serve as bargaining anchors, and perspective taking, as well as various other

[1] See the literature starting at Schelling [1960] leading to Isoni et al. [2013, 2014], Hargreaves Heap et al. [2014].

psychological manipulations could potentially bias bargaining outcomes [Kristensen and Gärling, 1997, Galinsky and Mussweiler, 2001, Van Poucke and Buelens, 2002, Mason et al., 2013, Ames and Mason, 2015]. However, Jeong et al. [2019] show that making first offers in a "warm and friendly" communication style surprisingly leads to less favorable outcomes in buyer-seller bargaining, while Weir et al. [2020] find a null result priming distributive and integrative language in the context of dam maintenance and wildlife preservation.

In this paper, we replicate Camerer et al. [2019] whose design has its closest precursor in Forsythe, Kennan, and Sopher (henceforth FKS), as both studied unstructured bargaining with one-sided private information about the sizes of several possible pies [Forsythe et al., 1991b]. With two possible pie sizes, FKS apply the revelation principle [Myerson, 1979, 1984] to identify a "strike condition" predicting when disagreements would be ex-ante efficient. They then experimentally verify (qualitatively) their theory with free-form communication. Camerer et al. [2019] generalize the FKS model to allow for any finite number of pie sizes, resulting in equilibria which maximize efficiency or equality that create different predictions. Thererfore, unlike FKS, their experimental design has 6 different pie sizes and record 10 seconds per trial of visible offers and counter-offers with little restrictions. This dynamic strategic environment with information asymmetry extends the recent literature on free form bargaining with full information [Herreiner and Puppe, 2004, Galeotti et al., 2018].

Our main finding, using National Taiwan University subjects and some small design changes, is a close replication of earlier results using US subjects in California. 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 also modest 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.

The reminder of this paper is organized as follows. In section 2.1, we summarize qualitative properties of bargaining in equilibrium derived from mechanism design theory. The experimental design we replicate is presented in section 2.2, and its general results summarized in section 3.1. We replicate the machine learning results in section 3.2. Finally, section 4 points to possible new directions of future research.

2 Theory and Experiments

In this paper, we adopt the theoretical framework from Camerer et al. [2019] to generate comparative statics predictions regarding the frequency of disagreements in each state with only the game structure, incentive compatibility (IC) and individual rationality (IR) constraints. Since the mechanism design approach only characterizes the class of possible equilibria rather than predicts specific outcomes, 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.

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 indexed by $k \in \{1, 2, \ldots, K\}$, 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 payoff of the uninformed player is w, and is bargained over by the players continuously communicating their bids within a certain amount of time T—which is also common knowledge. If the players agree on 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 private information regarding 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) 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

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

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

A direct mechanism is interim-efficient if the payoff profile is Pareto-optimal for the informed player in each of the K possible states and the uninformed player (in expectation) [Holmström and Myerson, 1983].

Based on the IC, IR-1, IR-2 and conditions, Camerer et al. [2019] prove the following two lemmas regarding bargaining outcomes and interim-efficient strikes.

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.*

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

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

Note that $x_k = \gamma_k w_k$, where w_k is the uninformed player's payoff conditional on a deal being made in state k.

The 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 which 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 $\{\$1, \dots, 6\}$.

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. [2020]). Based on players' tendency 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 subjects to efficiency conditions. By either prioritizing the former or the latter, Camerer et al. [2019] derive two competing equilibrium predictions, which are of the **the efficient equilibrium**:

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

and **the equal split equilibrium**: (See Online Appendix C for details of the derivation)

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

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

Our experiment 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 w. Both players communicate their offers, in multiples of \$0.1,[2] using a mouse click on a graphic interface which was programmed with z-Tree software (Fischbacher [2007]). Both players start with two seconds to 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 enter 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).[3] 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 a total of eight experimental sessions in the Social Science Experimental Laboratory (SSEL) at Caltech and the California Social Science Experimental Laboratory (CASSEL) at UCLA. At the beginning of each session, subjects are randomly seated at isolated computer workstations and given printed versions of the instructions, which are also read aloud by the experimenter. All participants complete a short quiz to check their understanding. Subjects play 15 practice rounds to become familiar with the game and interface, and then play 120 real rounds. They are paid a randomly chosen 15% of the rounds, plus a show-up fee of \$5. Each session lasts between 70 and 90 minutes, which includes check-in, instructions, experimental task, and payment.

2.2.3 Experiment 2

The follow-up experiment is conducted in the Taiwan Social Sciences 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. In addition, we conduct one experienced session and one high-stake session in order test whether our results are robust to experience and stakes. In the experienced session, we recruit subjects who have participated one of the six previous sessions.

[2] Camerer et al. [2019] (Experiment 1) set the resolution to be in multiples of \$0.2, since they thought \$0.1 was too fine a resolution for coordinating in a short game. However, the result in Experiment 1 shows that 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.

[3] 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.

In the high-stakes session, we multiply the stakes by 5. Notice that there is no gender constraint in the experienced and high-stake 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 pre-set exchange rate (1 ECU = NT$15) specified in the instructions. In the high-stake 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. [2018]. 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 realized. Before undergoing DOSE, all subjects evaluated their subjective willingness to take risk on a scale from 0 (not willing to take any risk at all) to 10 (willing to take any risk). The evaluation would not affect the payoff. Each session lasts around 2.5 hours.

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 Online Appendix A and B for further analysis on the payoffs and the bargaining dynamics.

Table 1 provides the summary statistics of average bargaining outcomes in different treatments. The average bargaining outcomes are similar across treatments. Differences in the average payoffs across treatments are less than $0.1 and differences of average deal rates are within 5%. We highlight some of our findings in the following: The average surplus loss is the lowest in the experienced treatment and the highest in the male-informed treatment. Turning to the information value, which can be interpreted as the advantage of knowing the pie size, we observe that it is the largest in the experienced treatment and lowest in the baseline treatment. Bargaining outcomes are generally robust across different treatments and stakes on the aggregated level.

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.

Next, we break down deal rates according to different pie sizes for different treatments. Figure 2 and Figure 3 show that in all treatments, deal rates increase with the pie size. This confirms our theoretical prediction in Lemma 1. Moreover, 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, 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

We defer further results from Experiment 2 to the Online Appendix. These results include analyses of the bargaining dynamics (see Online Appendix A) and testing predictions in Lemma 1 (see Online Appendix B). In general, 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 observe that 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 a large amount of bargaining process data beyond initial demands and offers to predict disagreements before the deadline. Hence, we search for a small set of such features that is predictive, employing cross-validation (Stone [1974]) to control for over-fitting.

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 choose from the 35 behavioral features introduced by Camerer et al. [2019]. Among them are the current difference between the offer and demand, the time since the last position change, and which player had changed his or her position in the game first. We compared three outcome prediction models at eight time points in the bargaining process (i.e., $1, 2, \ldots, 8$ seconds after bargaining starts). The first model relied only on the pie size, the second used only process features and the third combined both. At each time point, we carried out the following nested cross-validation procedure: For each of the eight sessions in Experiment 2, we used the data of the remaining seven sessions to train our model. The model classifies trials into disagreements or deals by estimating a logistic regression with a least absolute shrinkage and selection operator (LASSO) penalty (Tibshirani [1996]). The tuning parameter, λ, is optimized via ten-fold cross-validation, performed within each training set. We then made out-of-sample outcome predictions (disagreement or deal) for the hold-out session.

To compare the three models, we use the "receiver operating characteristic" (ROC) curves [Hanley and McNeil, 1982, Bradley, 1997], standard in signal detection theory to quantifies the performance of binary classifiers under various trade-offs between type I and type II errors. The 45-degree line in Figure 4 indicates a random classifier whose true positive and false positive rates are identical. A better classifier has higher true positive rates (moving up on the y axis) and lower false positive rates (moving left on the x axis). The "area under the curve" (AUC), or difference between the ROC and the 45-degree line in the upper-left direction, is an index of how well the classifier does.

INSERT FIGURE 4 AND 5 HERE

Figure 4 shows the ROC curve at $t = 2, 5, 7$ seconds for both Experiment 1 and 2. The ROC analysis indicates that process data do better than random at every time point in both experiments. Moreover, the fitness of models with process data increase with time, but the same is not true for the model with pie size only.

While patterns of AUC are similar in Experiment 1 and 2, there are still some subtle differences. In Experiment 1, the model with pie and process features 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 the three, its predictive power is not significantly stronger than the model with process features only.

To further investigate which behavioral process features predict strikes, we follow Camerer et al. [2019] and 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. The general feature patterns in Experiment 2 are consistent with those in Experiment 1. 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. One surprising finding in Camerer et al. [2019] is that initial bargaining positions contain predictive information regarding the possibility of reaching a deal, even as we approach the deadline, and even after controlling for current offers. In Experiment 2, such effect of initial positions is even stronger. We also find a negative interaction between initial offer and initial demand and a negative interaction between initial and current offer, which again confirms the arguments in Camerer et al. [2019].[4]

INSERT FIGURE 6 HERE

4 Open Questions, Challenges, and Future Directions

Unstructured bargaining seems to hand over the reins of endogenous "treatments" to the experimental subjects. But 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. Instead, the challenge lies in interpreting the results, which can be ad hoc and sometimes opaque. Besides, our replication results indicate that even a lot of process data in a highly controlled setting results in only modest AUC values, so there is still much room to improve.

Furthermore, 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 therefore predicts comparably with one using many process features. One interesting question future studies should explore is whether alternative process models could predict significantly better, since our combined model improves only modestly beyond using only pie size.

Thirdly, Experiment 2 shows a modest effects of gender. While gender effects in bargaining are interesting, a lot more statistical power is probably needed. In fact, 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.

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] allowed subjects to transmit verbal messages during bargaining. 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. While they allowed messages and recorded them, they did not analyze them at all because they deemed the resulting game—treating messages as strategy choices—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. For example, Jeong et al. [2019] use NLP to analyze first offer messages to identify "warm and friendly" communications.

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.

[4] However, not all effects are transparent; some even reverse across time (initial x current offer in Experiment 2).

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

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.

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.

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. *Behavioral Game Theory: Experiments in Strategic Interaction*. The Roundtable Series in Behavioral Economics. Princeton University Press, 2003. ISBN 9780691090399.

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.

C. M. Capra. Understanding decision processes in guessing games: a protocol analysis approach. *Journal of the Economic Science Association*, 5(1):123–135, Aug 2019. ISSN 2199-6784.

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

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

J. Crémer and R. P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

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. *American Economic Review*, 81(1):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.

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*, 51(6):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.

M. Jeong, J. Minson, M. Yeomans, and F. Gino. Communicating with warmth in distributive negotiations is surprisingly counterproductive. *Management Science*, 65(12):5813–5837, 2019.

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

E. Karagözoğlu. On "going unstructured" in bargaining experiments. In J.-F. Laslier, H. Moulin, M. R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers*, pages 295–304. Springer International Publishing, Cham, 2019. ISBN 978-3-030-18050-8.

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.

P.-H. Lin, A. L. Brown, T. Imai, J. T.-y. Wang, S. W. Wang, and C. F. Camerer. Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments. *Nature Human Behaviour*, 4(9):917–927, 2020.

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*, 60(2): 395–421, 1992.

R. B. Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–73, 1979.

R. B. Myerson. Two-person bargaining problems with incomplete information. *Econometrica*, 52 (2):461–487, 1984.

J. F. Nash Jr. The bargaining problem. *Econometrica*, 18(2):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.

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

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*, 50(5):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. *American Economic Review*, 78(4):806–823, 1988.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

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

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.

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, 2010.

S. W. Wang, J. Chapman, E. Snowberg, and C. F. Camerer. Loss attitudes in the u.s. population: Evidence from dynamically optimized sequential experimentation (dose). *mimeo*, 2018.

M. J. Weir, C. M. Ashcraft, N. Leuchanka Diessner, B. McGreavy, E. Vogler, and T. Guilfoos. Language effects on bargaining. *PLOS ONE*, 15(3):1–20, 03 2020.

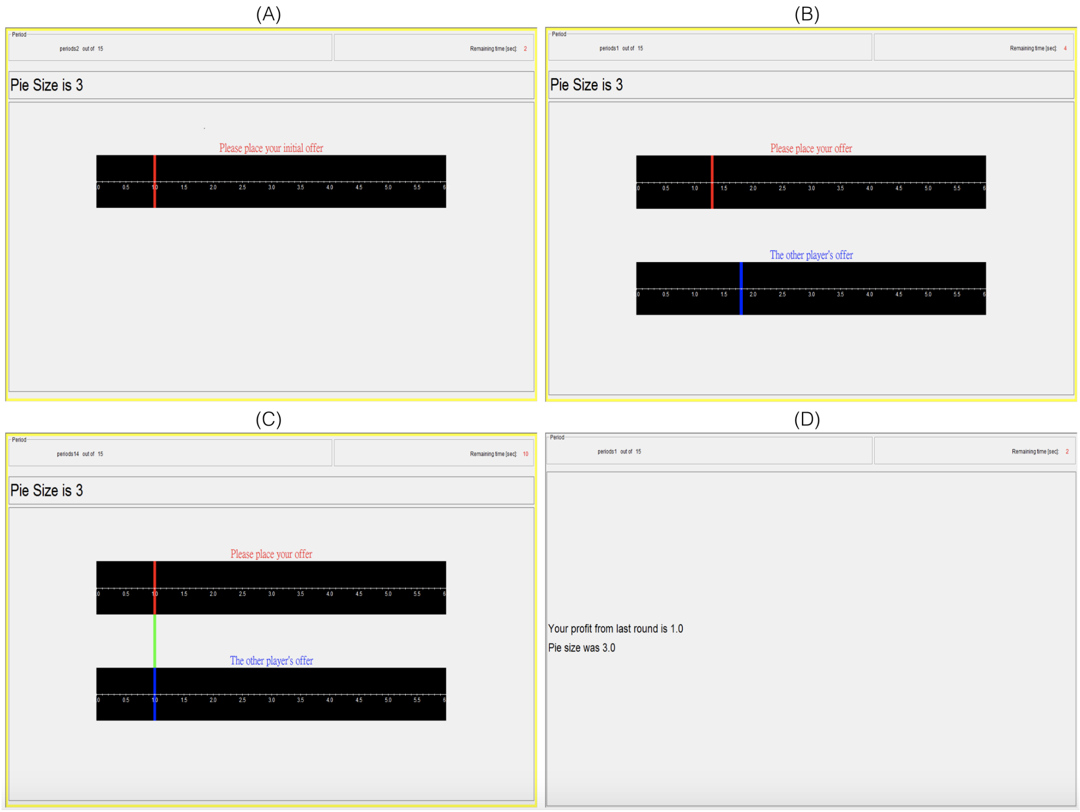

Figure 1: (A) Initial offer screen: in the first two seconds of bargaining, both players can set their initial positions without revealing to the opponent. The pie size is on the top left corner and it only appears on informed player's screen. (B) Players communicate their offers using mouse click on 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.

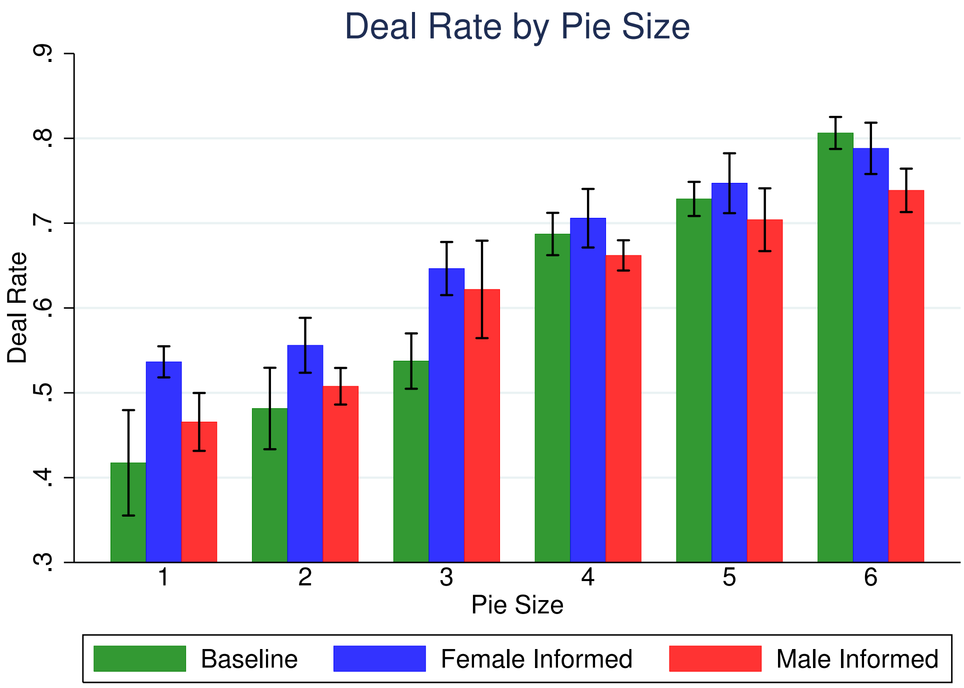

Figure 2: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for female-informed sessions and male-informed sessions, respectively. The standard errors (overlaid on the bars) are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

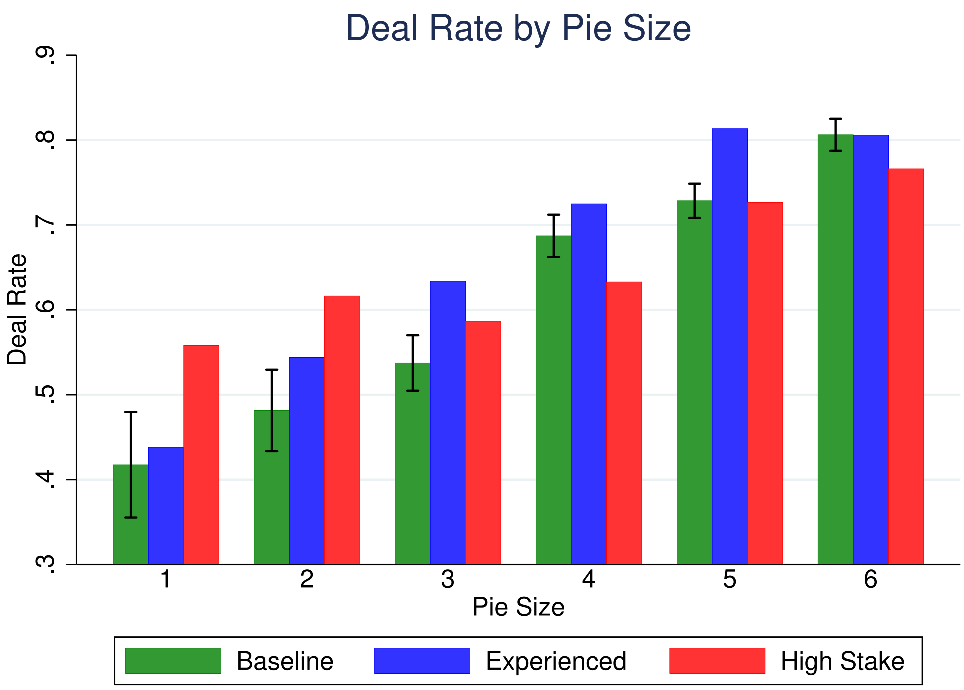

Figure 3: The deal rates under different pie sizes and treatments. The green bars stand for the average deal rates of baseline sessions at different pie size. The blue and red bars are for the experienced sessions and high-stake sessions, respectively. The standard errors (overlaid on the bars) are calculated 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.

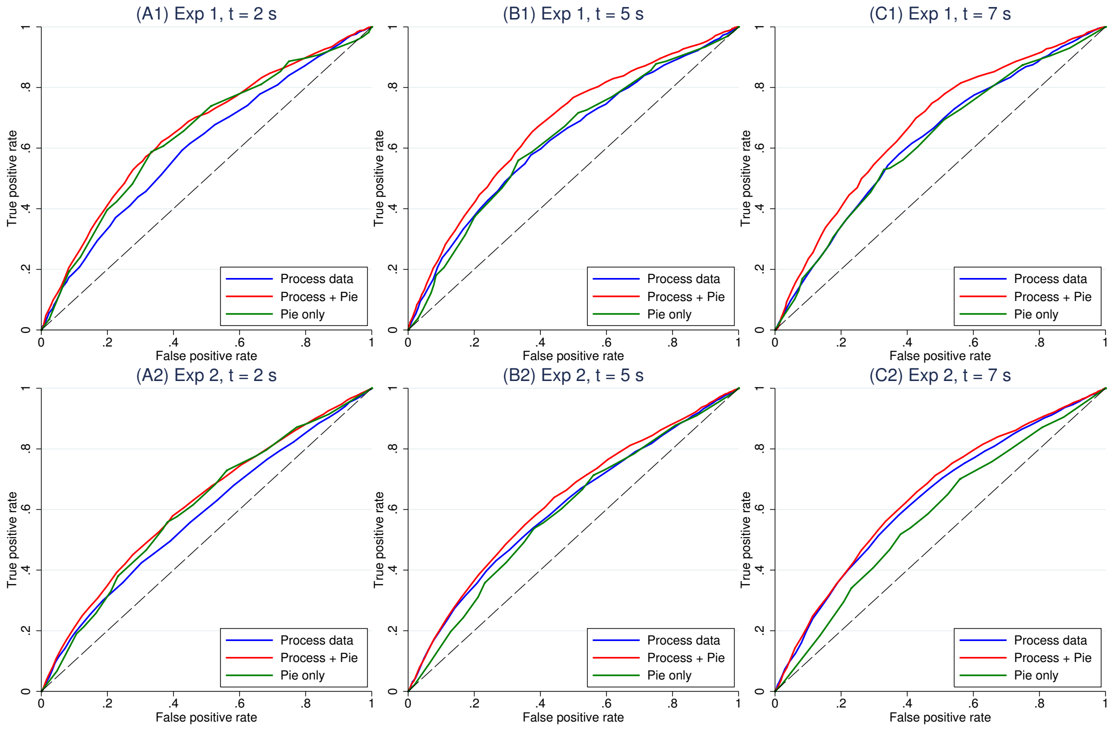

Figure 4: Receiver operating characteristic (ROC) for predicting disagreements, two, five and seven seconds into the bargaining game. The dashed lines represent the false and true positive rates of a random classifier. (A1–C1) show the data from Camerer et al. [2019] (Experiment 1) and (A2–C2) plot the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

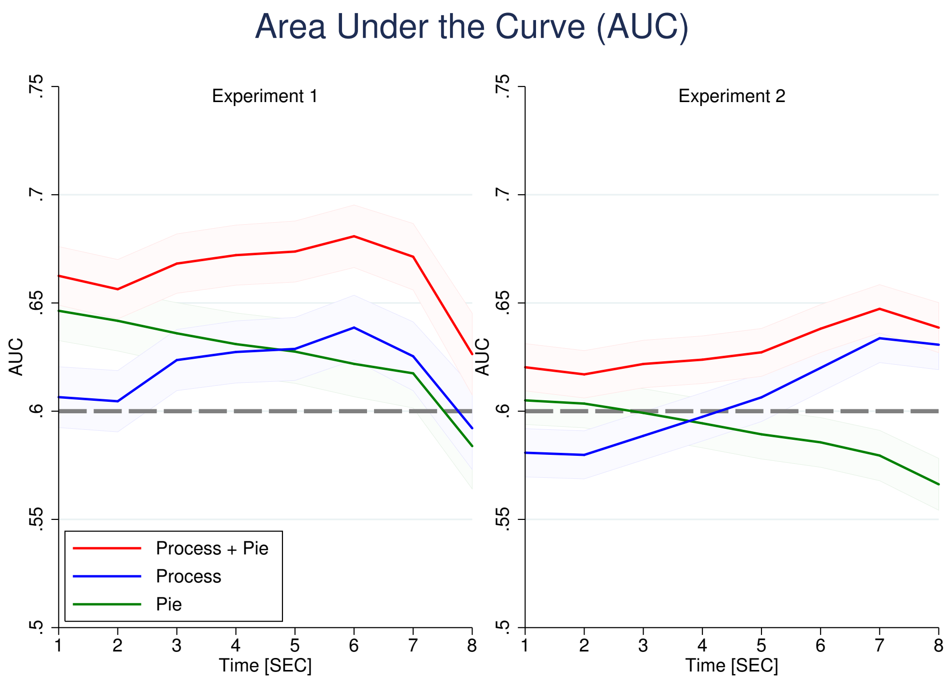

Figure 5: Area under the curve (AUC) of disagreements classifiers using process data, pie size, and the two combined. Note that the classifier's input included only trials that were still in progress (when a deal has not yet been achieved), and excluded trials in which the offers and demand were equal 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.

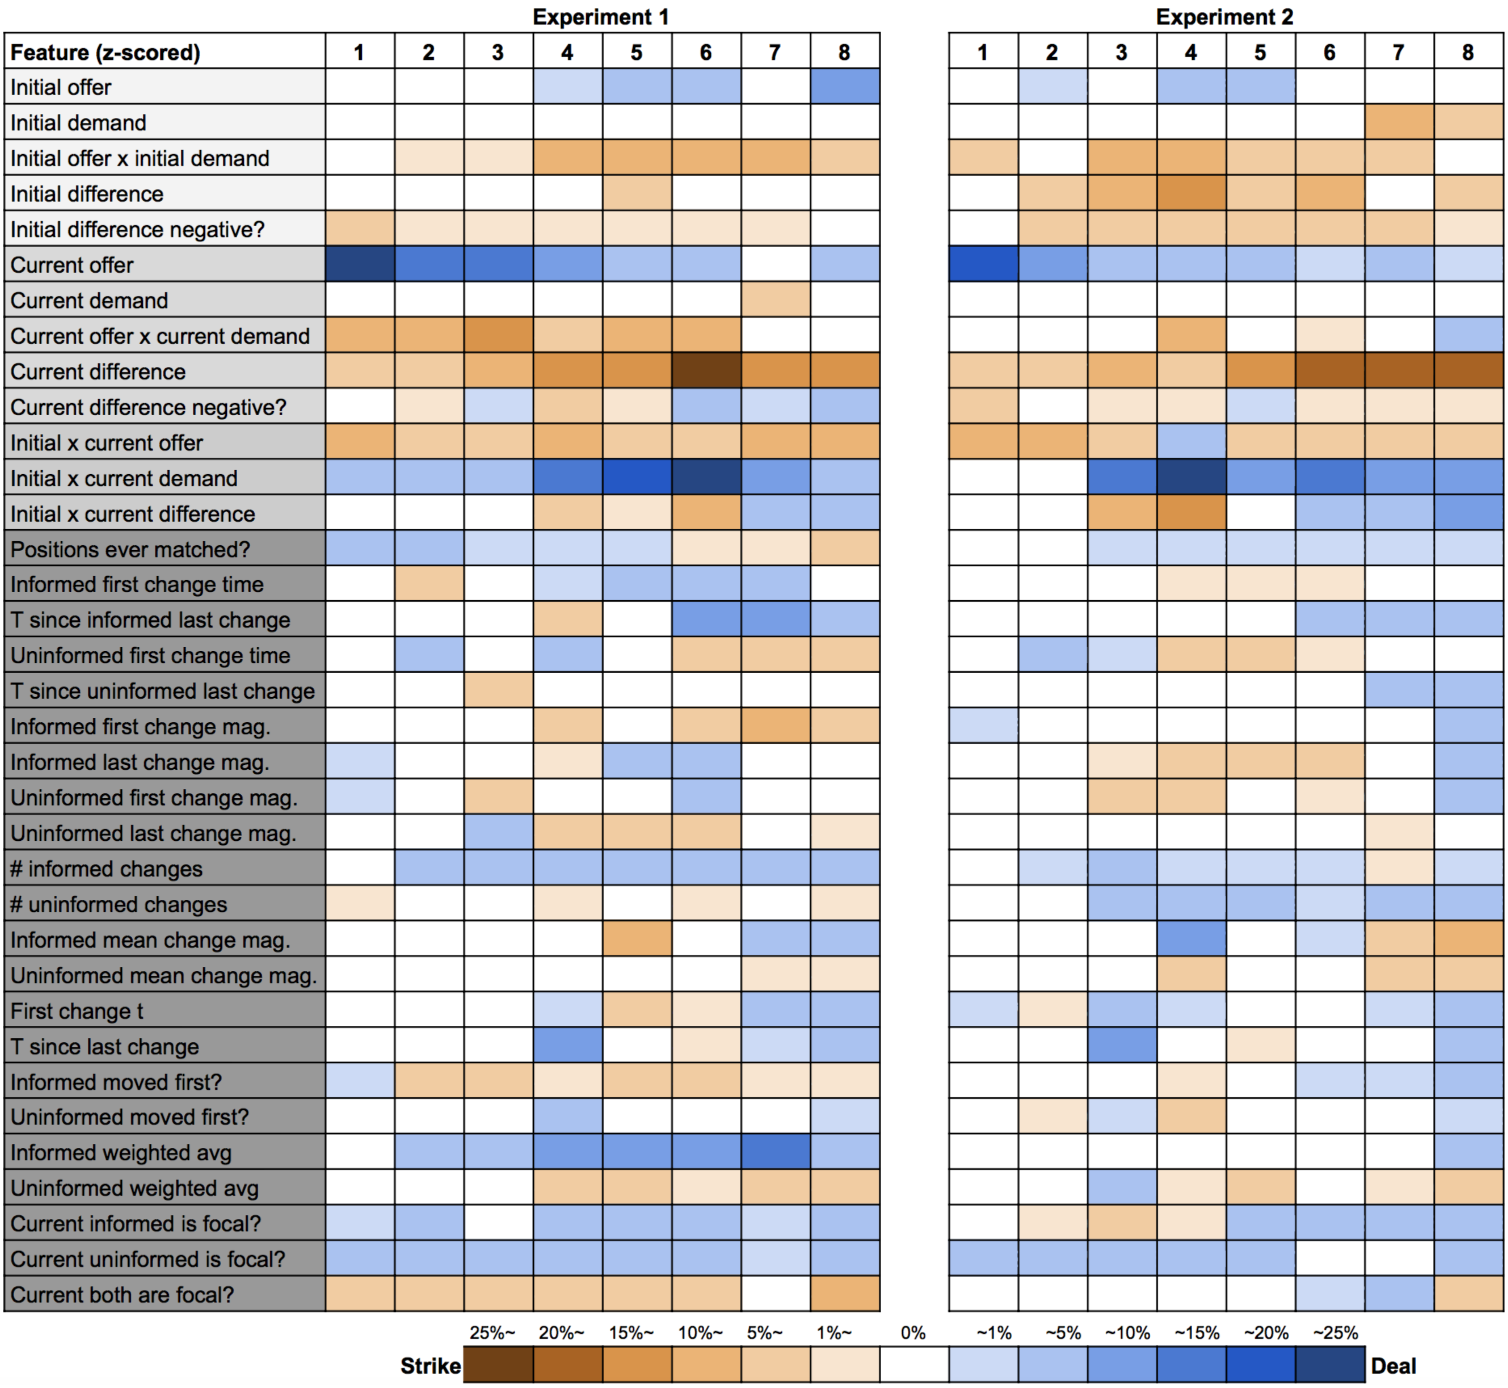

Figure 6: Bargaining Process Features Selected by the Classifier for Outcome Prediction (Deal= 1) and Their Estimated Marginal Effects. The left panel is the result from Camerer et al. [2019] (Experiment 1) and the right panel is the result from Experiment 2. The pie sizes are excluded.