



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Dynamic Unstructured Bargaining with Private Information: Theory, Experiment, and Outcome Prediction via Machine Learning

Colin F. Camerer, Gideon Nave, Alec Smith

To cite this article:

Colin F. Camerer, Gideon Nave, Alec Smith (2019) Dynamic Unstructured Bargaining with Private Information: Theory, Experiment, and Outcome Prediction via Machine Learning. Management Science 65(4):1867–1890. <https://doi.org/10.1287/mnsc.2017.2965>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Management Science*. Copyright © 2018 The Author(s). <https://doi.org/10.1287/mnsc.2017.2965>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2018, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Dynamic Unstructured Bargaining with Private Information: Theory, Experiment, and Outcome Prediction via Machine Learning

Colin F. Camerer,^a Gideon Nave,^b Alec Smith^c
^aHumanities and Social Sciences and Computational and Neural Systems, California Institute of Technology, Pasadena, California 91125;

^bMarketing Department, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104; ^cDepartment of Economics, Virginia Tech, Blacksburg, Virginia 24061

Contact: camerer@caltech.edu,  <http://orcid.org/0000-0003-4049-1871> (CFC); gnave@wharton.upenn.edu,

 <http://orcid.org/0000-0001-6251-5630> (GN); alecsmith@vt.edu,  <http://orcid.org/0000-0002-1938-4660> (AS)

Received: June 15, 2016

Revised: June 9, 2017

Accepted: September 8, 2017

Published Online in Articles in Advance:
May 8, 2018

<https://doi.org/10.1287/mnsc.2017.2965>

Copyright: © 2018 The Author(s)

Abstract. We study dynamic unstructured bargaining with deadlines and one-sided private information about the amount available to share (the “pie size”). Using mechanism design theory, we show that given the players’ incentives, the equilibrium incidence of bargaining failures (“strikes”) should increase with the pie size, and we derive a condition under which strikes are efficient. In our setting, no equilibrium satisfies both equality and efficiency in all pie sizes. We derive two equilibria that resolve the trade-off between equality and efficiency by favoring either equality or efficiency. Using a novel experimental paradigm, we confirm that strike incidence is decreasing in the pie size. Subjects reach equal splits in small pie games (in which strikes are efficient), while most payoffs are close to either the efficient or the equal equilibrium prediction, when the pie is large. We employ a machine learning approach to show that bargaining process features recorded early in the game improve out-of-sample prediction of disagreements at the deadline. The process feature predictions are as accurate as predictions from pie sizes only, and adding process and pie data together improves predictions even more.

History: Accepted by Uri Gneezy, behavioral economics.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Management Science*.” Copyright © 2018 The Author(s). <https://doi.org/10.1287/mnsc.2017.2965>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.

Funding: Generous support was provided by the National Science Foundation [SES-0850840] and the Behavioral and Neuroeconomics Discovery Fund at Caltech. Open access was sponsored by C. Camerer.

Supplemental Material: Data are available at <https://doi.org/10.1287/mnsc.2017.2965>.

Keywords: bargaining • dynamic games • private information • mechanism design • machine learning

1. Introduction

Bargaining is everywhere in economic activity: from price haggling in flea markets, to wage negotiations between unions and firms, to high-stakes diplomacy. Even in competitive, large-scale markets, sequences of market trades often result from individual buyer–seller partners bargaining over a range of mutually agreeable contract terms, knowing their outside options from the market. Bargaining failures such as holdouts and strikes—due to disputes over what each side should get—are also common and reduce welfare.

Strikes are surprising because in almost every case, the bargain that was eventually struck after a costly strike could have been agreed to much earlier in the bargaining, which would have saved lost profits, legal bills, and many other collateral costs. Then why do strikes happen? The standard approach in the game theory of private-information bargaining is that the

willingness to endure a strike is the only way for one side to credibly convince the bargaining partner that their existing offer is inadequate. Although making a deal appears to be a better outcome for both sides, when players’ incentives and information are taken into account, strikes are not just efficient: they can also be unavoidable (Kennan and Wilson 1990).

Private information bargaining theories, and tests of these theories, have developed in two ways:

(1) The most popular way is bargaining theories based on highly structured settings (e.g., Ståhl 1972 or Rubinstein 1982; for a review, see Ausubel et al. 2002). “Structure” means that the rules of how bargaining proceeds are clearly specified in the theory. 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.¹ Following the burst of progress in game theory on structured private-information bargaining, a large experimental literature emerged testing these theories.²

The clear assumptions about structure in the theory made experimental design and theory-testing straightforward.

(2) The less popular way of theorizing and experimentation in economics is based on unstructured bargaining. 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 study unstructured bargaining.

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 theory. Establishing these regularities can *lead* theorizing, rather than test theory.

Second, 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. Of course, theories can gain precision by ignoring these process data. However, if process variables are systematically associated with outcomes, these empirical regularities both challenge simple equilibrium theories and invite new theory development. Indeed, we use process data in a new way: To predict which bargaining trials will result in deals and strikes. We use a penalized regression approach from machine learning, to select those features from a large number of process features and make out-of-sample, cross-validated predictions (guarding against overfitting). The 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 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.³

Advice like this can be easily tested by carefully controlled experimental designs that allow structure-free bargaining while keeping the process fully tractable, such as our paradigm.

Third, 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 “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 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 nonobvious 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 deal rates can be made. That is, even if the bargaining protocol lacks structure, predictions can have plenty of restricted “structure.”⁴

2. Background

2.1. Experimental Economic Literature of Unstructured Bargaining

The experimental literature on bargaining is vast, so below we only focus on studies closely related to ours.⁵ Before theoretical breakthroughs in the understanding of structured bargaining, most experiments used unstructured communication. The main focus of interest was process-free solution concepts such as the Nash bargaining solution (Nash 1950), and important extensions (e.g., Kalai and Smorodinsky 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, p. 259). Under information asymmetries, disagreements may arise due to coordination difficulties. Several papers by Roth and colleagues then explored what happens when players bargain over points that 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. Many of the disagreements could be traced to self-serving differences between which of two focal points should be adopted—whether to allocate points equally or the money, resulting from points, equally. Focal points have remained an important theme in more recent work.⁶ 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.”

Several experiments have observed what happens in unstructured bargaining with *two-sided* private information (Valley et al. 2002). The typical finding is that in face-to-face and unstructured communication via message passing, there are *fewer* disagreements than predicted by theory.⁷ However, when players bargaining can only make a single offer, 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).

The closest precursor to our design is that of 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. 1991).⁸ 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 six pie sizes with rapid bargaining (10 seconds per trial), where bargaining occurs only through visible offers and counteroffers, 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.

From the literature studying structured bargaining, Mitzkewitz and Nagel (1993; henceforth MN) is a closely related design. They study ultimatum bargaining with incomplete information. MN use the same distribution over pie sizes in ultimatum bargaining that we employ in unstructured bargaining. The pattern of payoffs and disagreements in our results is similar to that of MN’s “offer” game, in which the informed player makes an ultimatum proposal. Our results generalize their conclusion that fairness and equality concerns matter in asymmetric information ultimatum bargaining to a less structured environment.

2.2. The Equality vs. Efficiency Trade-off

Another branch of literature in economics that is related to our study is the experimental work investigating how humans resolve trade-offs between equality and efficiency. While this question is still under (heated) debate,⁹ it is largely accepted that people are heterogeneous with respect to how they prioritize these factors.¹⁰

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

2.3. Negotiation Research

Finally, our study closely relates to negotiation research (Pruitt 2013), a branch of social psychology 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 (in utility terms) at which the negotiator places a non-negligible likelihood that value would be accepted by the other party(ies)” (White and Neale 1994, p. 304). Aspirations played an important role in determining the bargainers’ initial offers and were shown to influence bargaining outcome variables such as deal rates and surplus division.¹¹

The reminder of this paper is organized as follows. In Section 3, 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 we propose two equilibria that solve this trade-off by either favoring the former or the latter. We present a novel experimental design in Section 4, and summarize its general results in Section 5. We use machine learning to examine how bargaining process data can be associated with bargaining outcome variables in Section 6, and conclude in Section 7.

3. Theory

In this section, we develop a theory that provides testable predictions of deal rates and surplus division. Our model combines two methods to analyze bargaining: mechanism design and focal points. We extend

the model of strikes developed in Kennan (1986) and Forsythe et al. (1991) to an arbitrary finite number of states. This extension yields nonobvious predictions of the frequency of disagreement (the strike rate) in each state, using only the game structure, rationality, and incentive-compatibility constraints. Assuming interim efficiency allows further predictions. We then suggest a focal point approach to the problem of equilibrium selection. Combining these two approaches yields testable predictions about both deal rates and payoffs in each state.

3.1. Game and Notation

Two players must agree on how to split a surplus (or “pie”), a random variable denoted by π . The informed player knows the actual size of the pie. The uninformed player knows that the informed player knows the pie size. The finite set of states of the world are indexed by $k \in \{1, 2, \dots, K\}$, and the pie size in state k is π_k . Without loss of generality, we assume $\pi_k > \pi_j$ when $k > j$. The probability distribution of pie sizes $\Pr(\pi_k) = p_k$ is commonly known. The players have a finite amount of time T , which is commonly known, to reach an agreement. They bargain over the payoff of the uninformed player, denoted by w , by continuously communicating their bids. Players cannot commit to a particular bargaining position. In the case of agreement on an uninformed player's payoff w , the informed player gets $\pi - w$. If no deal is made by time T , both players' payoffs are zero.

3.2. The Direct Bargaining Mechanism

By the revelation principle (Myerson 1979, 1984), for any Nash equilibrium in the bargaining game, there exists a payoff-equivalent equilibrium of a simplified game (“a direct mechanism”) in which the informed player truthfully reveals the pie size to a neutral “mediator” who determines the payoffs and the probability of a strike based on that report (Forsythe et al. 1991). Following FKS, we assume that bargainers negotiate inscrutably over the set of direct mechanisms of the following type.

In the direct mechanism, the informed player announces the true size of the pie, π_k . The pie is then decreased by a known fraction, $1 - \gamma_k$, which can be interpreted as the strike probability in state k , leaving an expected pie size of $\gamma_k \pi_k$. We refer to γ_k as the deal probability, and to $1 - \gamma_k$ as the strike probability. The uninformed bargainer receives x_k , and the informed player gets the rest of the pie, $\gamma_k \pi_k - x_k$. To make predictions regarding observed behavior, we rely on the fact that the payoff x_k in the direct mechanism is tantamount to the expected payoff of the uninformed player in state k of the bargaining game: $x_k = \gamma_k w_k$ such that w_k is the uninformed payoff conditional on a deal in state k . A mechanism therefore involves $2K$ parameters, $\{\gamma_k, x_k\}_{k=1}^K$.

3.2.1. Individual Rationality (IR). Individual rationality requires that both players prefer to participate in the mechanism. Therefore, the IR requirement is that for all k ,

$$\gamma_k \pi_k - x_k \geq 0, \quad (1)$$

$$x_k \geq 0. \quad (2)$$

3.2.2. Incentive Compatibility (IC). A mechanism is incentive compatible (IC) if it is optimal for the informed player to tell the truth—i.e., her expected payoff is (weakly) maximized when she announces the true size of the pie. This requires

$$\gamma_k \pi_k - x_k \geq \gamma_j \pi_k - x_j, \quad \text{for all } k, \text{ for all } j \neq k. \quad (3)$$

The IR and IC conditions together lead to the following result.

Lemma 1. *If the bargaining mechanism satisfies IR and IC:*

1. *Deal rates are monotonically increasing in the pie size π_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.*

Proof. Incentive compatibility requires

$$\gamma_k \pi_k - x_k \geq \gamma_{k+1} \pi_k - x_{k+1},$$

$$\gamma_{k+1} \pi_{k+1} - x_{k+1} \geq \gamma_k \pi_{k+1} - x_k.$$

These two equations imply that

$$(\gamma_{k+1} - \gamma_k) \pi_{k+1} \geq x_{k+1} - x_k \geq (\gamma_{k+1} - \gamma_k) \pi_k, \quad (4)$$

and therefore

$$(\gamma_{k+1} - \gamma_k)(\pi_{k+1} - \pi_k) \geq 0. \quad (5)$$

By definition, $\pi_{k+1} \geq \pi_k$, so then $\gamma_{k+1} \geq \gamma_k$, and therefore the deal rate γ_k is monotonically increasing in the pie size (Lemma 1.1). Then, since the difference in deal rates $\gamma_{k+1} - \gamma_k$ is weakly positive, by the right-hand side of Equation (4), the difference in the uninformed player's payoffs $x_{k+1} - x_k$ is also weakly positive, and therefore the uninformed player's payoffs are monotonically increasing in the pie size (Lemma 1.2). Finally, if the deal rate in state k is 1, by Lemma 1.1, it must also be 1 in states $j > k$. Replacing $\gamma_k = \gamma_{k+1} = 1$ in both the right- and left-hand sides of Equation (4), it follows that $x_k = x_{k+1}$. Therefore, the uninformed player's payoff is identical in states k and $k + 1$, and, by induction, in all states $j > k$. (Lemma 1.3). $\square$

3.2.3. Efficiency. In our setting, a mechanism is efficient (more precisely, is “interim-incentive efficient”; Holmström and Myerson 1983) if it is Pareto optimal

for the set of $K + 1$ agents: the K informed players in each of the different states k , and the uninformed player.

Lemma 2. *The strike condition: For IR and IC mechanisms, strikes in state k are ex ante efficient if*

$$\frac{\pi_k}{\pi_{k+1}} < \frac{(1 - \sum_{j=1}^k p_j)}{(1 - \sum_{j=1}^{k-1} p_j)} = \frac{\Pr(\pi \geq \pi_{k+1})}{\Pr(\pi \geq \pi_k)}. \quad (6)$$

Proof. See Section A.1 of Appendix A.

The relations between pie size ratios and conditional probabilities of pie size in Equation (6) are called “strike conditions.” For a given pie size k , strikes are interim efficient so long as the probability distribution over pie sizes does not place too much weight on state k relative to the weight on higher states.¹²

By Lemma 1.1, if there exists a cutoff state, π_c , in which $\gamma_c = 1$ (no strikes), then strikes are inefficient in all states π_k such that $k \geq c$. Furthermore, as the uninformed player’s payoff must be the same in all states where no disagreements occur (Lemma 1.3), this implies that if strikes are inefficient in more than a single state, there exists no equilibrium where both efficiency and payoff equality hold for all states. Thus, there is a built-in tension between efficiency and equality under some informational settings.

3.3. Equilibrium Selection Using Focal Points

In theory, the IR and IC constraints limit the scope of possible bargaining outcomes and predict when strikes are likely to occur. This is remarkable considering that the bargaining protocol is unstructured. However, these conditions do not precisely pin down the numerical strike rates $1 - \gamma_k$ and the equilibrium payoffs (conditional on a deal being reached) w_k for each state. There are many such sets of parameter values that will satisfy IR and IC, and that are equilibrium outcomes.

To make a more precise prediction, we incorporate an equilibrium selection approach that relies on the extensive literature emphasizing the importance of focal points in bargaining games (Schelling 1960; Roth 1985; Kristensen and Gärling 1997; Janssen 2001; Binmore and Samuelson 2006; Janssen 2006; Bardsley et al. 2010; Isoni et al. 2013, 2014).

Absent other salient features of bargaining, the natural focal point is an equal split (i.e., $w_k = \pi_k/2$). Indeed, equal splits often emerge in bargaining experiments (e.g., Roth and Malouf 1979, Roth and Murnighan 1982).¹³ Note that equal sharing is also common in sharecropping contracts (Young and Burke 2001), corporate budget allocations to divisions (Bardolet et al. 2011), bequests to heirs (Menchik 1980, Behrman and Rosenzweig 2004), and sharing of university invention royalties (Kotha et al. 2015). Regardless of the source

of equal sharing, here, we simply use this regularity as a basis for generating precise numerical predictions of the deal rates.

In practice, we propose that the equilibrium payoff of the uninformed player, conditional on a deal, will equal half of the pie size ($w_k = \pi_k/2$) as long as an equal split satisfies the IR and IC conditions (Lemma 1), and subject to efficiency conditions that we discuss below. We use this premise to calibrate our model and derive two competing predictions that resolve the tension between efficiency and equality (discussed in Section 3.2.3) by either prioritizing the former or the latter.

3.3.1. The Efficient Equilibrium. To prioritize efficiency over equality, we set the deal rate to 1 whenever the strike condition (Lemma 2) does not hold (i.e., whenever strikes are inefficient). Then, we split the pie equally given this constraint. Suppose that strikes are inefficient for all pies that are greater than π_c . As discussed above, this implies that the uninformed player’s payoff must be the same for all $\pi_k \geq \pi_c$ (Lemma 1.3). To yield a clear prediction about the equilibrium uninformed payoffs w_k^* , we divide the pie equally in lower-value pie states given this constraint:

$$w_k^* = \begin{cases} \frac{\pi_k}{2}, & \forall \pi_k \leq \pi_c, \\ \frac{\pi_c}{2}, & \forall \pi_k > \pi_c. \end{cases} \quad (7)$$

In our experiment, π takes on values that are the integer dollar amounts 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 $\pi \geq 4$, there should be no strikes. Combining this efficiency constraint with Lemma 1.3 and the focal principle of equal splitting implies that an equal split of $\pi = 4$ (i.e., the uninformed player’s payoff is 2) can be an equilibrium, but then the same amount (2) must also be the equilibrium payoff of the uninformed player for the larger pie sizes 5 and 6.

The assumption of efficiency, the strike condition (Equation (6)), and the use of focal payoffs (Equation (7)) enable us to pin down the exact deal rates for all pie sizes. We first set $\gamma_4 = \gamma_5 = \gamma_6 = 1$, as required by the strike condition when pie sizes are uniformly distributed over $\{1, 2, 3, 4, 5, 6\}$. The IC conditions and the assumption of efficiency require that the left-hand side of Equation (4) holds at equality, because raising the deal rate in any state weakly improves the payoffs of both players.¹⁴ Therefore, efficiency requires that

$$(\gamma_{k+1} - \gamma_k)\pi_{k+1} = x_{k+1} - x_k, \quad \forall k. \quad (8)$$

Noting again that the uninformed player’s payoff in each state x_k in the direct bargaining mechanism is equal to the payoff in case of a deal times the strike rate,

we fix $x_k = \gamma_k(\frac{1}{2}\pi_k)$ for all $k < 4$, and $x_k = 2$ for all $k \geq 4$ as required by the efficiency condition, and substitute these expressions into Equation (8) to obtain

$$\begin{cases} \gamma_k = \frac{(1/2)\pi_{k+1}}{\pi_{k+1} - (1/2)\pi_k} \gamma_{k+1}, & \forall k < 4, \\ \gamma_k = 1, & \forall k \geq 4. \end{cases} \quad (9)$$

Solving this set of equations yields the prediction that

$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = (\frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1, 1, 1). \quad (10)$$

We can also use Equation (8) to show that under certain conditions, which are met by both of our candidate equilibria, deal rates are linear in the pie size.

Remark 1. If the difference between pie sizes $\pi_{k+1} - \pi_k$ is constant, for all k , and the uninformed player's payoff, conditional on a deal, is equal to half the pie size (so that $x_k = \frac{1}{2}\gamma_k\pi_k$), then $\gamma_k - \gamma_{k-1} = \gamma_{k+1} - \gamma_k$; that is, the change in deal rates is constant.

The proof is in Section A.2 of the appendix.

3.3.2. The Equal Split Equilibrium. As discussed in Section 3.2.3, some efficiency must be sacrificed to achieve equality for every pie size. In the equal split equilibrium, we first impose equal splits and only then maximize efficiency given this constraint:

$$w_k^* = \frac{\pi_k}{2}. \quad (11)$$

As the deal rates are increasing with the pie size (Lemma 1.1), and as the uninformed payoff must be identical in all states where there are no strikes (Lemma 1.3), 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 equilibrium, we set $\gamma_6 = 1$. Then, we again make use of Equation (8) to calculate deal rates:

$$\begin{cases} \gamma_k = \frac{(1/2)\pi_{k+1}}{\pi_{k+1} - (1/2)\pi_k} \gamma_{k+1}, & \forall k < 6, \\ \gamma_k = 1, & k = 6. \end{cases} \quad (12)$$

Solving this set of equations yields the prediction that

$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = (\frac{2}{7}, \frac{3}{7}, \frac{4}{7}, \frac{5}{7}, \frac{6}{7}, 1). \quad (13)$$

4. Experiment

In this section, we present a novel experimental paradigm of dynamic bargaining, which allows both parties to communicate offers whenever they please, while keeping their behavior tractable.

4.1. Design

Our design is a continuous-time bargaining game with one-sided private information. At the start of each session, participants were randomly divided into two equally sized type groups: informed and uninformed. The types were fixed for the session's 120 bargaining round. Each round had the following steps:

Step 1. Each player was randomly matched with a partner from the other group in a stranger protocol (to prevent sequential effects such as reputation building).

Step 2. In each game, an integer pie size, $\pi \in \{\$1, 2, 3, 4, 5, 6\}$, was drawn from a commonly known discrete uniform distribution:

$$\Pr(\pi_k) = \frac{1}{6}, \quad \forall \pi \in \{\$1, 2, 3, 4, 5, 6\}.$$

Step 3. The informed player was told the true value of π for that round.

Step 4. Each pair bargained over the uninformed player's payoff, denoted by w . Players communicated their monetary offers, in multiples of \$0.2,¹⁵ using mouse clicks on a graphical interface that was designed for this purpose by z-Tree software (Fischbacher 2007)¹⁶ (see Figure 1). The offer values were between \$0 and \$6.

Step 5. During the first two seconds of bargaining, both players fixed their initial offers, without seeing the offers of their partner (see Figure 1(a)). The initial cursor location (i.e., before the first click) was randomized.

Step 6. Once the initial offers were set, players bargained continuously for 10 seconds using mouse clicks (see Figure 1(b)).¹⁷

Step 7. When players' positions matched each other, visual feedback was given to both of them in the form of a vertical stripe connecting their offer lines (see Figure 1(c)). If none of the players changed their position for the next 1.5 seconds following the offer-match feedback, a deal was made (i.e., the deal is closed 1.5 seconds after the initial match). Thus, to make a deal, the latest time in which players' bids could match was $t = 8.5$ seconds.

Step 8. If no deal had been made within 10 seconds of bargaining, both players' payoffs from that round were \$0.

Step 9. After each game, both players were told their payoffs and the actual pie size, for five seconds (see Figure 1(d)).

4.2. Methods

We conducted 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. There were a total of $N = 110$ subjects (mean age 21.3, SD 2.4; 47 females (see Appendix B for detailed session information)). In the beginning of each session, subjects were randomly assigned to isolated computer

Figure 1. (Color online) Bargaining Interface



Notes. (a) Initial offer screen: in the first two seconds of bargaining, players set their initial position, oblivious to the initial position of their partner. The pie size at the top left corner appears only for the informed type. (b) Players communicate their offers using mouse click on the interface. (c) When demands match, feedback in the form of a green vertical stripe appears on the screen. If no changes are made in the following 1.5 seconds, a deal is made. (d) Following the game, both players are notified regarding their payoffs and the pie size.

workstations and were handed printed versions of the instructions (see Appendix D). The instructions were also read aloud by the experimenter (who was the same person in all sessions). All of the participants completed a short quiz to check their understanding of the task. Subjects played 15 practice rounds to become familiar with the game and the interactive interface before the actual play of 120 rounds. Participants' payoffs were based on their profits in a randomly chosen 15% of the rounds, plus a show-up fee of \$5. Each session lasted between 70 and 90 minutes (including check-in, reading of instructions, experimental task, and payment). The data consist of each subject's bargaining positions in each game and the outcomes of 120 rounds.¹⁸

5. Experimental Results

5.1. Main Findings

Below, we report our main findings. Supporting statistical analyses and hypothesis tests are contained in the subsequent two sections. We observed the following empirical regularities:

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

The mean deal rates and payoffs, conditional on a deal being reached, for each pie amount are summarized in Table 1 and Figure 2(a). While the *probability* of disagreement decreased with the pie size, the mean amount of surplus lost because of strikes (Table 1) was positively correlated with the pie, as relatively small amounts of money are lost when strikes occur in small pie games. However, as a fraction of the total available surplus, the loss is greatest in the low-pie-size games, and smallest in the high-pie-size games. In addition, the difference between uninformed and informed payoffs is larger in high-pie-size games.

Result 2. *When the pie is small or medium ($\pi \leq \$4$), the modes of the uninformed players' payoffs distribution are half of the pie; in large-pie games ($\pi > \$4$) the modes are \$2 and there are local maxima at the half of the pie.*

The distributions of uninformed players' payoffs are in Figure 3. The mean payoffs (conditional on a deal


Table 1. Average Payoffs and Deal Rates by Pie Size

Pie size	1	2	3	4	5	6	Mean
Informed payoff	0.37 (0.03)	0.95 (0.04)	1.56 (0.04)	2.23 (0.03)	3.07 (0.05)	3.87 (0.06)	2.01
Uninformed payoff	0.63 (0.03)	1.05 (0.04)	1.44 (0.04)	1.77 (0.03)	1.93 (0.05)	2.13 (0.06)	1.49
Deal rate	0.42 (0.06)	0.48 (0.05)	0.54 (0.03)	0.69 (0.02)	0.73 (0.02)	0.81 (0.02)	0.61
Surplus loss	0.58 (0.06)	1.04 (0.10)	1.39 (0.10)	1.25 (0.10)	1.36 (0.10)	1.16 (0.11)	1.13
Information value ^a	-0.11 (0.03)	-0.05 (0.03)	0.05 (0.04)	0.31 (0.04)	0.83 (0.07)	1.39 (0.10)	0.40

Notes. Averages are calculated for deal games only. Means and standard errors are calculated by treating each session's mean as a single observation. Standard errors in parentheses.

^aInformation value = the mean difference between the informed and uninformed payoffs.

being reached) are in Figure 2(b). Overall, 82% of the payoffs, conditional on a deal being reached, match values that are halves of one of the six possible pies.¹⁹ Equal splits are the most prevalent outcomes (51.4%) of small- and medium-pie games ($\pi \leq 4$), where the predicted payoffs of the efficient and equal equilibria coincide (Figure 3, top two rows). In large-pie

Figure 2. (Color online) Deal Rates and Mean Payoffs Across Pie Sizes

Note. Standard errors are calculated at the session level.

games ($\pi \geq 5$), equality and efficiency are in discord. The payoff distributions of these games (Figure 3, bottom row) have modes at the efficient (but unequal) uninformed payoff of 2 (31% of payoffs), and local maxima (19% of payoffs) at the equal-split (but inefficient) payoffs of 2.5 (when $\pi = 5$, 17% of payoffs) and 3 (when $\pi = 6$, 20% of payoffs). Thus, about half of the bargaining payoffs match one of the two equilibria. These results confirm that equality concerns did influence bargaining outcomes, generalizing the experimental literature studying complete information bargaining (Nydegger and Owen 1974, Roth and Malouf 1979) and ultimatum bargaining with private information (Mitzkewitz and Nagel 1993) to an unstructured environment with informational asymmetry.

Result 3. *The informed players' offers increase, and the uninformed players' demands decrease with time (within a trial).*

Result 3 is illustrated by the plots of mean bargaining positions shown in Figure 4. In Section 6, we explore how the process of bargaining affects outcomes.

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

More than half of the deals were made in the last two seconds of bargaining. Figure 5 shows the cumulative distribution function (CDF) of deals over time, which sharply increased as the deadline approached for all pies.²⁰ Generally, deals were reached sooner when the pie was larger. This result is in line with the “deadline effect” reported in previous studies of unstructured bargaining with full information (Roth et al. 1988, Gächter and Riedl 2005).


5.2. Comparison with Theoretical Predictions

We now turn to testing the predictions derived from the bargaining theory. For convenience, we refer to the informed and uninformed players' bargaining positions as “offers” and “demands,” respectively. We first use nonparametric statistics on the pooled data to test our theoretical predictions for both deal rates and payoffs. We then use linear regressions that include controls for the location of the experiment, the session, and the experience of the subjects, to compare our results with the focal equilibrium predictions from Section 3.3.

5.2.1. Nonparametric Hypothesis Tests on the Pooled Data.

Our analysis of the direct bargaining mechanism predicts that deal rates will be monotonically increasing in the pie size. Table 1 and Figure 2(a) suggest that deal rates rise smoothly with increasing pie size, in line with the qualitative prediction. Hypothesis tests on the pooled experimental data (treating each game between two players as an observation) are also consistent with the model's predictions. A nonparametric Wilcoxon-test for trend (Cuzick 1985) in deal rates over pie

Figure 3. (Color online) Uninformed Player’s Payoff Relative Frequencies (Deal Games, Binned in a \$0.25 Resolution)



Note. The green bar locates the half of the pie in each distribution.

sizes strongly rejects the null hypothesis of no trend ($z = 23.24, p < 0.001$, two-sided).

The next prediction resulting from Lemma 1 is that the uninformed players’ payoffs will be monotonically increasing in the pie size, conditional on a deal being

reached. A nonparametric Kruskal–Wallis test (with adjustments for ties) rejects the null hypothesis of equality of the distributions of the uninformed players’ payoffs in each pie size ($\chi^2(5) = 1,800.162, p < 0.001$, two-sided). Consistent with Lemma 1, a nonparametric

Figure 4. (Color online) Mean Bargaining Position for All Pie Sizes (All Rounds Pooled)





Figure 5. (Color online) Cumulative Distribution of Deal Times by Pie Size

Note. Median deal times are marked by an asterisk.

Wilcoxon-type test rejects the null hypothesis of no trend in payoffs conditional on the pie size ($z = 41.11$, $p < 0.001$, two-sided).

Lemma 1 also predicts that the uninformed player's payoff will be identical for all states in which the deal probability is 1. The efficient equilibrium prediction is that the deal probability will be 1 in pie sizes 4 and greater, while the equal-split equilibrium gives deal rates equal to 1 in pie size 6 only. In contrast to the predictions of both equilibria, strikes are common in all pie sizes (see Table 1 and Figure 2(a)) and occur even at the largest pie size of 6 (about 19%, averaging over eight sessions; see Table 1). A nonparametric Kruskal–Wallis test for equality of payoff distributions (with corrections for ties) rejects the null hypothesis that the mean payoffs to the uninformed player, conditional on a deal, are equal in pie sizes 4, 5, and 6 ($\chi^2(2) = 111.503$, $p < 0.001$, two-sided).

5.2.2. Regression Analyses. In Section 3.3, we relied on the assumption that players coordinate on focal payoffs to derive precise equilibrium predictions about deal rates and payoffs. The “efficient” equilibrium predicts a deal rate of $\frac{2}{5}$ in pie size 1, increasing by $\frac{1}{5}$ per pie unit up to pie size 4, and then a deal rate of 1 when the pie is greater than or equal to 4. The “equal split” equilibrium predicts a deal rate of $\frac{2}{7}$ when the pie equals 1, increasing by $\frac{1}{7}$ per unit, up to 1 in pie size 6.

We test these predictions using linear regressions, where the dependent variable is whether a deal was reached in a given bargaining game, as follows:

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

Here, y_{iust} equals 1 if a deal was reached in the bargaining round between informed player i and uninformed player u in session s and period t , and 0 otherwise.²¹ The predictions of the two models are

nested via the inclusion of the spline term $d_{iust}(\pi_{iust} - 4)$, where d_{iust} is a dummy variable that takes value 1 if $\pi_{iust} \geq 4$, and 0 otherwise; and the term $(\pi_{iust} - 4)$ generates a knot at pie size 4 (see Greene 2003, pp. 121–122). The vector X_{iust} contains control variables. With this specification, the equal split equilibrium prediction is $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{7}, \frac{1}{7}, 0)$, where the efficient equilibrium prediction is $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{5}, \frac{1}{5}, -\frac{1}{5})$.

Table 2 reports regression results. All models include standard errors clustered at the session level, to account for dependence in residuals within a particular session.²² Model A gives the base results from the regression of deal on pie size and the spline term, using pooled data from all sessions. Models B includes session controls, and Model C uses controls at the level of individual subject pairs (the smallest grouping available). Model D drops these controls and adds an indicator term controlling for the session location (Caltech=1 or UCLA=0), 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 the location and experience terms and the pie and spline terms. Model F drops the Caltech indicator and adds session-level controls. Model G (which is discussed in detail in the next section) adds the initial bargaining positions of the players as covariates.

In all models, the coefficient on pie size is significantly positive, indicating an 8.0% to 9.5% increase in deal probability for a 1 unit increase in the pie size. This estimate is robust to controlling for location and experience, and to the inclusion of session or pair-level controls (see Models B–E of Table 2). Thus, consistent with the predictions derived from the IC conditions in Lemma 1, we find that deal rates increase in the pie size. However, the slope coefficient on the pie size is smaller than predicted by either equilibrium model, with the upper bound for the 95% confidence interval of the coefficient estimate on pie size (about 0.11) lower than the minimum prediction of either model ($\frac{1}{7}$ or approximately 0.14 for the equal-split model, or $\frac{1}{5}$ for the efficient equilibrium). The constant term from the regression in Model A of Table 2 (approximately 0.28) is larger than predicted by either model, but is closer to the prediction of the equal split model ($\frac{1}{5}$). The coefficient on the spline term for pie sizes 4 and above is not significantly different from 0 in any of the models, and F -tests of the full model specifications consistently reject both models.²³ We do not find a statistically significant difference in deal rates between the subject populations at UCLA and Caltech.

In a similar fashion, we used linear regressions to test the focal-split predictions about payoffs conditional on the players reaching a deal. The efficient equilibrium predicts equal splits when the pie is small, and that the uninformed player's conditional payoff will be 2 when the pie is 4 or greater. The equal-split equilibrium

Table 2. 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
<i>Pie</i>	0.0940*** (0.0154)	0.0935*** (0.0156)	0.0900*** (0.0159)	0.0938*** (0.0155)	0.0951*** (0.0170)	0.0868*** (0.0129)	0.0802*** (0.0108)
<i>Spline at pie = 4</i>	−0.0284 (0.0263)	−0.0262 (0.0270)	−0.0200 (0.0278)	−0.0276 (0.0270)	−0.0425 (0.0454)	−0.0246 (0.0298)	−0.0171 (0.0263)
<i>Caltech</i>				0.0710 (0.0425)	0.1067 (0.1211)		
<i>Rounds 61–120</i>				0.0347* (0.0154)	−0.0049 (0.0495)	−0.0084 (0.0471)	−0.0041 (0.0477)
<i>Caltech × Pie</i>					−0.0138 (0.0309)		
<i>Caltech × Spline</i>					0.0260 (0.0505)		
<i>Rd.61–120 × Pie</i>					0.0106 (0.0151)	0.0126 (0.0140)	0.0145 (0.0132)
<i>Rd.61–120 × Spline</i>					0.0048 (0.0203)	−0.0025 (0.0174)	−0.0101 (0.0149)
<i>Initial demand</i>							−0.0546*** (0.0103)
<i>Initial offer</i>							0.0257 (0.0305)
<i>Constant</i>	0.2805*** (0.0612)	0.2812*** (0.0434)	0.2902*** (0.0439)	0.2309*** (0.0497)	0.2332*** (0.0439)	0.2861*** (0.0279)	0.4402*** (0.0531)
Observations	6,432	6,432	6,432	6,432	6,432	6,432	6,432
AIC	8,543.859	8,400.798	7,106.364	8,502.368	8,501.844	8,395.001	8,278.694
BIC	8,564.166	8,414.336	7,119.902	8,536.213	8,549.228	8,428.847	8,326.077
Session controls	No	Yes	No	No	No	Yes	Yes
Pair controls	No	No	Yes	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

assumes a 50/50 split for all pie sizes. We test these predictions via linear regressions that include a spline term, of the form

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

where w_{iust} represents the payoff, conditional on a deal being reached, to uninformed player u , paired with informed player i , in session s and period t . The equal-split equilibrium prediction is $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, 0)$, as the coefficient on the spline term is 0. The efficient equilibrium prediction is $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$, with the $-\frac{1}{2}$ coefficient on the spline term capturing the prediction that uninformed payoffs will be constant at pie sizes 4 and higher.

The regression results are shown in Table 3, with the same model specifications as in Table 2. As above, the results are robust to the inclusion of controls for either the session or for subject pairs (Models A–C). The coefficient on the pie size is positive and statistically significant, consistent with the prediction derived from the IC conditions that uninformed payoffs increase with the pie size. In line with the efficient equilibrium prediction (of a constant uninformed payoff when the pie is larger than 4), the coefficient of the spline term is negative and significant. However, these analyses only partially support the efficient equilibrium hypothesis: both

the slope term (between 0.34 and about 0.37 in Models A–F) and the intercept term (which ranges from 0.23 to 0.34 in Models A–F) are lower than the efficient equilibrium predictions that both terms will equal 0.5. The spline coefficient does not fully offset the pie coefficient when the pie is large. Formally, both equilibrium models are rejected by an F -test, as is the equal-split equilibrium prediction that the marginal effect of pie size is zero in high pie sizes (e.g., in Model A, $F(1, 7) = 598.48, p < 0.0001$). We also find a small effect of location when interacted with pie size (Model C).

With 120 rounds of play, our experiment design allows us to investigate whether behavior converges to the predicted equilibria as participants gain experience. Models B–E in Tables 2 and 3 include an indicator for the second half of the experiment (rounds 61–120) as a proxy for experience. We find a small effect of experience on deal rates (the coefficient on rounds 61–120 was positive and marginally significant (Table 2, Model B)) as well as uninformed payoffs (the coefficient on rounds 61–120 is positive and significant in Model B of Table 3, indicating that the uninformed player earns about \$0.10 more per game in the second half of the experiment). The interaction term of rounds 61–120 and pie size is also marginally significant in Model E, suggesting some convergence

Table 3. 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
<i>Pie</i>	0.3708*** (0.0103)	0.3718*** (0.0098)	0.3603*** (0.0152)	0.3697*** (0.0106)	0.3428*** (0.0208)	0.3595*** (0.0153)	0.2589*** (0.0232)
<i>Spline at pie = 4</i>	−0.2252*** (0.0234)	−0.2230*** (0.0215)	−0.2196*** (0.0209)	−0.2237*** (0.0232)	−0.2540*** (0.0333)	−0.2200*** (0.0228)	−0.1527*** (0.0199)
<i>Caltech</i>				0.0462 (0.0714)	−0.1286 (0.1103)		
<i>Rounds 61–120</i>				0.1048** (0.0397)	0.0422 (0.0237)	0.0258 (0.0231)	−0.0089 (0.0311)
<i>Caltech × Pie</i>					0.0355* (0.0187)		
<i>Caltech × Spline</i>					0.0583 (0.0467)		
<i>Rd.61–120 × Pie</i>					0.0144 (0.0132)	0.0212 (0.0130)	0.0359* (0.0158)
<i>Rd.61–120 × Spline</i>					0.0126 (0.0366)	−0.0035 (0.0389)	−0.0208 (0.0368)
<i>Initial demand</i>							0.0345* (0.0147)
<i>Initial offer</i>							0.3968*** (0.0750)
<i>Constant</i>	0.3042*** (0.0572)	0.2990*** (0.0440)	0.3416*** (0.0577)	0.2304* (0.0977)	0.3549** (0.1037)	0.2894*** (0.0483)	0.1100** (0.0349)
Observations	3,819	3,819	3,819	3,819	3,819	3,819	3,819
AIC	6,254.813	6,156.995	4,561.273	6,217.291	6,184.654	6,122.826	4,869.375
BIC	6,273.557	6,169.491	4,573.769	6,248.530	6,228.389	6,154.065	4,913.109
Session controls	No	Yes	No	No	No	Yes	Yes
Pair controls	No	No	Yes	No	No	No	No

Notes. Coef., coefficient; SE, standard errors. Standard errors (in parentheses) are clustered at the session level.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

toward the focal equilibria predictions. However, even in the final rounds, strikes are common in all pie sizes.²⁴ We interpret this as evidence that additional factors drive bargaining outcomes beyond the pie size.

In summary, we find strong support for theoretical predictions derived from incentive and rationality (or participation) constraints from mechanism design, while predictions derived from efficiency considerations fare less well. Both deal rates and payoffs increase with the pie size, as predicted by Lemma 1. Disagreement rates are qualitatively more consistent with the equal-split equilibrium, though deal rates are generally higher in low pie sizes than our equilibrium predictions, and both deal rates and payoffs are less responsive to changes in the pie sizes than predicted by our models.²⁵ Interestingly, there is evidence of a lessened sensitivity of the uninformed player's payoffs to the pie size in high pie sizes ($\pi \geq 4$), which is qualitatively consistent with our efficient equilibrium model. However, the magnitude of this change is lower than predicted by the model.

Our results demonstrate that theoretical predictions, derived from mechanism design models that assume risk-neutral, selfish players, and focal-point-based equilibrium selection, can take us a long way,

even in unstructured settings, but also reveal the limitations of this approach. Bargaining outcomes qualitatively match a mix of two equilibrium patterns, and some game outcomes match neither equilibria. Furthermore, theoretical predictions critically depend on the pie size—which is private information that would typically be unobservable to econometricians in field data. In the next section, we use bargaining process data to overcome some of these limitations.


6. Using Process Data to Understand Disagreements

In counterpoint to the process-free approach of mechanism design, a large experimental literature in psychology and economics studies the role of procedural factors in determining bargaining outcomes. In this section, we explore how bargaining process affects outcomes.

6.1. The Influence of Bargaining Positions on Bargaining Outcomes

Many experimental studies have found that initial demands and offers can influence bargaining outcomes (Chertkoff and Conley 1967, Yukt 1974b, Gächter and

Figure 6. Uninformed Player's Initial Demands (Pooled Across All Games, Binned in a \$0.25 Resolution)



Riedl 2005, Karagözoğlu and Riedl 2015). Some negotiation researchers view these initial demands as reflecting players' aspirational payoffs—i.e., the most desirable payoffs that they can achieve, according to their beliefs (Yukl 1974a, White and Neale 1994, Kristensen and Gärling 1997, Galinsky and Mussweiler 2001, Van Poucke and Buelens 2002). Motivated by this research, we investigated the relationship between initial bargaining positions and outcomes in our task.

Figure 6 shows the distribution of the uninformed players' initial bargaining positions. The mode of the distribution (pooled across all pie sizes) was three (31%)—matching the highest possible equal equilibrium payoff.²⁶ An additional local maxima at two (19%) matched the highest possible payoff in the efficient equilibrium. Thus, the majority of the uninformed players' initial demands *exactly* match their maximal payoffs in either the efficient or equal equilibria, with a greater proportion matching the equal-split equilibrium.

In Model G of both Tables 2 and 3, we report results from regressions that add the initial offer of the informed player and the initial demand of the uninformed player to the regression models reported in Section 5.2. We examine the effects of initial offers on deals (Table 2) and on payoffs conditional on a deal being reached (Table 3). We find a significant and economically meaningful effect of initial demands by the uninformed player on the probability of a deal, such that a \$1 increase in initial demands reduced the probability of a deal by about 4.6%. However, it also led to an increase in the uninformed player's conditional payoff of \$0.035, though this increase is only marginally statistically significant. The initial offer of the informed player is not a statistically significant predictor of whether deals are reached, but it strongly predicts the uninformed player's payoff conditional on a deal being reached, above and beyond the pie size. This effect is economically large: a \$1 increase in the initial offer is predicted to result in a \$0.397 increase in

the payoff of the uninformed player when controlling for the size of the pie.²⁷

These results show that the process of bargaining plays an important role in determining whether a deal is reached, beyond the actual realization of the pie size. As private information might be unobservable to an econometrician in more natural settings, this finding has important practical implications, which we explore further in the following sections.

6.2. Predicting Disagreements Using Bargaining Process Data

Our unstructured paradigm 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 example, 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 (Krajbich et al. 2009, Belloni et al. 2012, Einav and Levin 2014, Varian 2014, Smith et al. 2014, Mullainathan and Spiess 2017, Bajari et al. 2015) and other social sciences (Dzyabura and Hauser 2011, Youyou et al. 2015, Yarkoni and Westfall 2017, Nave et al. 2018).

One possibility is that there is little predictive information in such features, after controlling for overfitting. Indeed, if players know what the predictive features are, they should alter their behavior to avoid costly disagreements, erasing the features' predictive power.²⁸ Another possibility is that there are numerous small influences on disagreement that the players simply do not notice and that may be picked up by our modeling.

We chose 34 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 Appendix C). For each of the eight experimental sessions, 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).²⁹ By applying these trained models, we then made out-of-sample predictions of the binary bargaining outcomes for each of the experimental sessions.

As noted earlier, the pie size is a strong predictor of disagreements. The challenges for our machine learning approach are twofold. First, we investigate whether

process features have predictive power similar to the pie size when studied alone. In other words, we test whether process data allow for predicting bargaining outcomes when the pie size, which is private information, is treated as if it were unobservable. Second, we investigate whether process features add predictive power when used *together* with the pie size.

To assess the predictive power of process data, we estimated three strike prediction models at eight different points in the bargaining process, separated by one-second intervals (i.e., 1, 2, ..., 8 seconds after bargaining started). One model relies only on the pie size,³⁰ the second uses only process features, and the third uses both pie size and process features.³¹ For each time stamp, predictions were carried out using the following nested cross-validation procedure: For each of the eight sessions, we trained 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 , was optimized via ten-fold cross-validation,³² performed within each training set.³³ Finally, using that trained model, we conducted out-of-sample predictions for the holdout sessions.

We evaluate our results using “receiver operating characteristic” (ROC) curves (Hanley and McNeil 1982, Bradley 1997). ROC 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. A familiar example is a household smoke alarm: the alarm can be tuned to be very sensitive, indicating a fire when a burnt toast creates too much smoke, or it can be tuned to be insensitive, ignoring the smoke from burnt toast but also possibly ignoring smoke from a genuine fire caused by a half-lit cigar accidentally knocked onto a copy of the *Daily Prophet* newspaper.

The use of an ROC curve reflects the fact that one can always create more true positives (in our example predicting more strikes), but doing so comes at the cost of then predicting more false positives (predicting strikes that do not happen). When using these methods, one would often like to know the trade-off between correctly detecting true positives more accurately and also reducing the probability of false positives. A curve mapping all pairs of true and false positive levels therefore allows for choosing an optimal policy for every given relative cost of the two types of errors.

To calculate the ROC, we subjected the out-of-sample predicted deal probabilities (calculated by applying the estimated logistic LASSO regression weights to the out-of-sample process data) to different decision thresholds—i.e., for a decision threshold $\tau \in [0, 1]$, all predicted values less than τ were classified as “strike,” whereas predicted values greater than or equal to τ were classified as “deal.”³⁴ Every point on the ROC, therefore, represents a decision threshold, such that its

coordinates represent the empirical false positive and true positive rates, calculated using the threshold.

For a random classifier, the true positive and false positive rates are identical (the 45-degree line in Figure 7). 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; Bradley 1997) is an index of how well the classifier does.³⁵


The ROC analysis shows that process data do better than random for every time stamp (for illustration, see Figures 7(a)–(c)). Furthermore, the AUC of the classifier that uses solely process features is similar to the AUC of a classifier using solely the pie size, for times greater than two seconds into the bargaining process (see Figure 7(d)). Combining of pie size and process features improves accuracy further: a classifier using both pie size and process data outperforms the classifier using the pie size alone for every time stamp in the bargaining process (Figure 7(d)). These results show that the process of bargaining itself can lead to bargaining failures, above and beyond the pie size alone.³⁶

6.3. Which Bargaining Process Features Predict Disagreements?

To further investigate which behavioral process features predict disagreements, we used a “post-LASSO” procedure (Belloni and Chernozhukov 2013, Belloni et al. 2012).³⁷ Figure 8 summarizes the marginal effects of the most predictive process features (z-scored for every time point), such that an “interaction” represents a multiplication of two variables. The marginal effects of all process features investigated are reported in Appendix C.

Not surprisingly, the most predictive process features are the current informed player’s offer (positively associated with a deal) and the current difference between the players’ bargaining positions (positively associated with a strike). More surprisingly, the players’ *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. The informed player’s initial offer is positively associated with a chance of a deal, and the effect is moderated by the uninformed player’s initial demand, as implied by a negative interaction between the two factors. Thus, initial offers are mostly associated with deals when the initial demands are low. There was also an intriguing negative interaction between the initial and current offers: the *current* offer becomes particularly associated with a deal when the *initial* offer is low. This result is consistent with an idea from negotiation research, that initial offers serve as reference points in bargaining. When initial offers are

Figure 7. (Color online) Strike Prediction Using Bargaining Process Data



Notes. (a–c) Receiver operating characteristic (ROC) for predicting disagreements, two and seven seconds into the bargaining game. The dashed lines represent the false and true positive rates of a random classifier. (d) 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.

low, they make later, more generous offers seem more attractive, and increase the chances of a deal (Galinsky and Mussweiler 2001).

Our analyses further revealed a rich set of behavioral features that reliably predicted disagreements throughout the bargaining process, even after controlling for the current bargaining positions (see Appendix C for all marginal effects). While these findings do not allow causal inference, and therefore should be interpreted with caution, they provide an avenue for further investigations of how bargaining outcomes are related to bargaining processes characteristics. For example, an increased activity on the informed player side (i.e., many position changes) is a precursor of an upcoming deal, as early as two seconds into the bargaining process. The use of focal points (i.e., offers and demands that match halves of the integer pies) was positively associated with an upcoming deal, unless both players’ positions match *different* focal points, as implied by a negative marginal interaction effect. This finding suggests that disagreements may arise as a result of a coordination failure when players use different focal points


to communicate their claims, in line with Roth’s focal theory of bargaining (Roth 1985).³⁸

7. Conclusion

Much of the recent literature on bargaining has studied structured bargaining. We reiterate here our motivations for studying unstructured bargaining in dynamic and uncertain environments. First, much real-world bargaining is unstructured and involves private information; unstructured bargaining generates process data that can be used to predict strikes ahead of time; and theory can be used to make precise predictions even with minimal structure.

In this paper, we study dynamic unstructured bargaining in a game with one-sided private information. We combine mechanism design theory with an equilibrium selection approach that builds on a well-documented empirical regularity: the appeal of an equal split as a bargaining focal point. Our approach is agnostic regarding the driving force behind equal splits. A large theoretical literature attempts to address the question of why equal splits are focal; equal splits might result, for example, from inequality aversion,

Figure 8. (Color online) Bargaining Process Features Selected by the Classifier for Outcome Prediction (Deal = 1) and Their Estimated Marginal Effects



Note. Pie sizes are excluded.

concerns about fairness, or social norms. Another explanation might be lying aversion (Gneezy et al. 2013), and our experiment’s design, which incorporates feedback after each round of bargaining, may encourage truthful revelation. However, our design also involves random, anonymous rematching of bargaining partners after each game, which might be expected to act in the opposite direction.

Our theoretical model predicts that the rate of bargaining failures will be decreasing in the pie size. The additional assumption of interim incentive efficiency implies that the distribution of surplus will favor the informed player when the pie size crosses a threshold. We find support for both of these hypotheses in our data. In addition, we also observe an interesting departure from the “efficient” benchmark: bargaining failures arise even at the highest pie levels and even after many rounds of play, and the surplus is divided equally in many high-stakes games, in contrast to the efficient equilibrium prediction.

In theory, the uninformed players’ payoffs must be identical in all pies where no disagreements occur, generating an inherent trade-off between efficiency and equality. We propose two ways to resolve this tension, by either favoring efficiency and dividing the pie equally given the efficiency constraint (“efficient” equilibrium) or by imposing equal splits and only then maximizing efficiency (“equal split” equilibrium). While the modes of the distributions of the informed

players’ final offers more closely match the efficient equilibrium, deal rates qualitatively match the prediction of the equal split equilibrium. Further, the uninformed players’ initial offers reflect aspirations of equal splits in the largest pie, suggesting that some uninformed players might use disagreements as a means to impose equal splits despite the loss of efficiency.

Although our results show that theoretical predictions based on the assumption of self-interest go a long way even in an unstructured setting, they also highlight their limitations. The data qualitatively match the mix of the two equilibria patterns, but some games do not match either. Further, the theoretical prediction depends on the realization of the pie size, which is private information and therefore might not be observable in many realistic circumstances. We propose overcoming this obstacle by analyzing bargaining process data.

Our machine learning approach shows that process data is incrementally informative for predicting strikes when the pie size is included in the model, and is as informative as knowing the pie size when the latter is unobservable (before the deadline has arrived). These results suggest that some bargaining failures may result from process “mistakes” that could have been avoided if players had behaved differently. Process data may be used to avert strikes and other inefficient disagreements by offering “course corrections” in the bargaining process. Bargaining process data could potentially be much richer, and therefore substantially

more informative, than the series of cursor locations on which our exploratory investigation has focused. An increase in the time of each bargaining period beyond 10 seconds and the incorporation of bargaining features such as verbal communication, nonverbal gestures (e.g., facial expressions, body language), and physiological responses (e.g., skin conductance, pupil dilation, brain activity) are likely to improve predictive performance. Our results should therefore be considered as a lower bound regarding the predictive power of process data, and an invitation to bargaining researchers to expand the scope of the data that are collected.

Finally, we acknowledge that our laboratory bargaining institution deliberately omits many features of natural bargaining. Lifelike bargaining is often face-to-face, has little anonymity, uses natural language, includes repetition and resulting reputations, and typically has two-sided private information. Moreover, individual differences in the bargainers' economic preferences (e.g., risk taking) and psychological traits (e.g., impulsivity, personality measures) may also play a role and provide additional explanatory power. The addition of more lifelike features can also be easily done step-by-step, as part of a research program reviving interest in unstructured bargaining. Typically, the addition of natural institutional properties and accounting for individual differences make it harder to figure out theoretically what behavior will result. The opposite is true when machine learning is used: the addition of more natural properties simply adds more "features" that can be used for prediction.

Acknowledgments

The author order is alphabetical, and all authors contributed equally to this project. The authors thank Eskil Forsell for research assistance, and Taisuke Imai, Boris van Leeuwen, and Joseph Wang for useful comments on this work.

Appendix A. Mathematical Appendix

A.1. Proof of Lemma 2: The Strike Condition

A mechanism is interim efficient if it is Pareto optimal for the set of $K + 1$ agents: the informed player for each pie size, and the uninformed player.

Following FKS, we first show that strikes in the "best" pie size π_K are never efficient for the class of direct mechanisms that we consider. That is, if the mechanism $\mu = \{\gamma_k, x_k\}_{k=1}^K$ is efficient, then it must be the case that $\gamma_K = 1$.

If μ is an efficient mechanism, then the incentive compatibility conditions must hold, and so by Lemma 1, $\gamma_K \geq \gamma_k$ for all $k \leq K$. If $\gamma_K = 1 - \delta < 1$, we can define a new mechanism μ^* with $\gamma_K^* = 1$, $\gamma_k^* = \gamma_k + \delta$, for all $k < K$, and $x_k^* = x_k$, for all k . The mechanism μ^* does not affect the uninformed player's expected payoff and does not affect the IC constraints (as both sides of the inequality increase by the same amount, $\delta\pi_K$), but it increases the informed player's payoff by $\delta\pi_K$ in state K and by $\delta\pi_k$ in states $1, \dots, K-1$, so the original mechanism cannot be efficient.

Next, if $\gamma_k, k < K$, can be increased without violating the IC constraint, the informed bargainer's payoff in states $j \neq k$ are unaffected, as is the uninformed bargainer's, while player I_k , the informed bargainer in state k , is made better off. Therefore, efficiency requires that the left-hand side of Equation (4) holds at equality:

$$(\gamma_{k+1} - \gamma_k)\pi_{k+1} = x_{k+1} - x_k, \quad \forall k. \quad (\text{A.1})$$

We make use of Equation (A.1) to derive Lemma 2, the strike condition, below.

Lemma 2. *The strike condition: For IR and IC mechanisms, strikes in state k are ex ante efficient if*

$$\frac{\pi_k}{\pi_{k+1}} < \frac{(1 - \sum_{j=1}^k p_j)}{(1 - \sum_{j=1}^{k-1} p_j)} = \frac{\Pr(\pi \geq \pi_{k+1})}{\Pr(\pi \geq \pi_k)}. \quad (\text{A.2})$$

Proof. To derive the strike condition, consider mechanisms $\mu = \{\gamma_k, x_k\}_{k=1}^K$ and $\mu^* = \{\gamma_k + \delta_k, x_k + d_k\}_{k=1}^K$, which satisfy IR and IC, and assume that both satisfy (A.1). Since μ^* satisfies (A.1), we have

$$(x_{k+1} + d_{k+1}) - (x_k + d_k) = ((\gamma_{k+1} + \delta_{k+1}) - (\gamma_k + \delta_k))\pi_{k+1}. \quad (\text{A.3})$$

By subtracting (A.1) from (A.3), we find a useful condition that

$$d_{k+1} - d_k = (\delta_{k+1} - \delta_k)\pi_{k+1}. \quad (\text{A.4})$$

Next, assume that strikes are not efficient in states $k + 1, \dots, K$, so that $\gamma_j = 1$ if $j > k$, but assume that $\gamma_k < 1$. This implies that $d_{k+1} = \dots = d_K$.

Let ΔV_k and ΔU represent the difference in payoffs between μ^* and μ for the informed player in state k and the uninformed player, respectively. If μ^* dominates μ , then $\Delta V_k \geq 0$ for all k , and $\Delta U \geq 0$, and at least one of these inequalities is strict.

First, consider the K conditions for the informed player.

$$\begin{aligned} \Delta V_1 &= \delta_1 \pi_1 - d_1 \geq 0; \\ &\vdots \\ \Delta V_j &= \delta_j \pi_j - d_j \geq 0, \quad j < k; \\ \Delta V_k &= \delta_k \pi_k - d_k \geq 0; \\ \Delta V_j &= \delta_k \pi_{k+1} - d_k \geq 0, \quad j > k. \end{aligned}$$

Multiplying the conditions for players $I_1, \dots, I_k$ by p_j and summing them up gives

$$\sum_{j=1}^k p_j \pi_j \delta_j \geq \sum_{j=1}^k p_j d_j.$$

Multiplying the equation for player k by $(1 - \sum_{j=1}^{k-1} p_j)$ gives

$$\left(1 - \sum_{j=1}^{k-1} p_j\right) \delta_k \pi_k \geq \left(1 - \sum_{j=1}^{k-1} p_j\right) d_k.$$

Adding up these two conditions gives

$$\sum_{j=1}^k p_j \pi_j \delta_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) \delta_k \pi_k \geq \sum_{j=1}^k p_j d_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) d_k. \quad (\text{A.5})$$

Next, we consider the uninformed player. If μ^* dominates μ , it must be the case that the uninformed player's payoff from μ^* is at least as large as in μ .

$$\begin{aligned}\Delta U &= \sum_{j=1}^K p_j d_j = \sum_{j=1}^k p_j d_j + \left(1 - \sum_{j=1}^k p_j\right) d_{k+1} \geq 0 \\ &\quad \sum_{j=1}^k p_j d_j + \left(1 - \sum_{j=1}^k p_j\right) (d_k - \delta_k \pi_{k+1}) \geq 0 \\ &\quad \sum_{j=1}^{k-1} p_j d_j + p_k d_k + \left(1 - \sum_{j=1}^k p_j\right) d_k - \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1} \geq 0 \\ &\quad \sum_{j=1}^{k-1} p_j d_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) d_k - \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1} \geq 0 \\ &\quad \sum_{j=1}^{k-1} p_j d_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) d_k \geq \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1} \quad (\text{A.6})\end{aligned}$$

Combining Equations (A.5) and (A.6) gives

$$\begin{aligned}\sum_{j=1}^{k-1} p_j \pi_j \delta_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) \delta_k \pi_k &\geq \sum_{j=1}^{k-1} p_j d_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) (d_k) \\ &\geq \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1}. \quad (\text{A.7})\end{aligned}$$

And this implies that

$$\sum_{j=1}^{k-1} p_j \pi_j \delta_j + \left(1 - \sum_{j=1}^{k-1} p_j\right) \delta_k \pi_k \geq \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1}. \quad (\text{A.8})$$

To examine whether strikes are efficient in state k , suppose μ^* and μ have identical strike rates in all states $j < k$. Then, δ_j equals 0 for all $j < k$, implying that

$$\left(1 - \sum_{j=1}^{k-1} p_j\right) \delta_k \pi_k \geq \left(1 - \sum_{j=1}^k p_j\right) \delta_k \pi_{k+1}. \quad (\text{A.9})$$

Then, $\delta_k > 0$ implies that strikes are inefficient in state k if

$$\frac{\pi_k}{\pi_{k+1}} \geq \frac{(1 - \sum_{j=1}^k p_j)}{(1 - \sum_{j=1}^{k-1} p_j)},$$

implying that strikes are efficient in state k if

$$\frac{\pi_k}{\pi_{k+1}} < \frac{(1 - \sum_{j=1}^k p_j)}{(1 - \sum_{j=1}^{k-1} p_j)},$$

or alternatively

$$\frac{\pi_k}{\pi_{k+1}} < \frac{\Pr(\pi \geq \pi_{k+1})}{\Pr(\pi \geq \pi_k)}. \quad \square \quad (\text{A.10})$$

A.2. Proof of Remark 1: Deal Rates Are Linear in the Pie Size for Equal Splits

Proof. If the uninformed player's payoff is equal to half of the informed player's payoff in the direct bargaining mechanism, then $x_k = \frac{1}{2} \gamma_k \pi_k$. Substituting this into (A.1), we have

$$\frac{1}{2} (\gamma_{k+1} \pi_{k+1} - \gamma_k \pi_k) = (\gamma_{k+1} - \gamma_k) \pi_{k+1}. \quad (\text{A.11})$$

If we assume that the difference between successive pie sizes is constant, then $\pi_k - \pi_{k-1} = \Delta$ (note that this condition holds

in our experiment). Adding and subtracting Δ to the left-hand side of Equation (A.11), we have

$$\begin{aligned}\frac{1}{2} (\gamma_{k+1} \pi_{k+1} - \gamma_k (\pi_k + \Delta - \Delta)) &= (\gamma_{k+1} - \gamma_k) \pi_{k+1}, \\ \frac{1}{2} (\gamma_{k+1} \pi_{k+1} - \gamma_k (\pi_{k+1} - \Delta)) &= (\gamma_{k+1} - \gamma_k) \pi_{k+1}, \\ \frac{1}{2} (\gamma_{k+1} - \gamma_k) \pi_{k+1} + \frac{1}{2} \gamma_k \Delta &= (\gamma_{k+1} - \gamma_k) \pi_{k+1},\end{aligned}$$

which implies that

$$\gamma_k \Delta = (\gamma_{k+1} - \gamma_k) \pi_{k+1}, \quad (\text{A.12})$$

and therefore also that

$$\gamma_{k-1} \Delta = (\gamma_k - \gamma_{k-1}) \pi_k. \quad (\text{A.13})$$

Subtracting (A.13) from (A.12), we have

$$\begin{aligned}(\gamma_k - \gamma_{k-1}) \Delta &= (\gamma_{k+1} - \gamma_k) \pi_{k+1} - (\gamma_k - \gamma_{k-1}) \pi_k, \\ (\gamma_k - \gamma_{k-1}) \Delta + (\gamma_k - \gamma_{k-1}) \pi_k &= (\gamma_{k+1} - \gamma_k) \pi_{k+1}, \\ (\gamma_k - \gamma_{k-1}) (\pi_k + \Delta) &= (\gamma_{k+1} - \gamma_k) \pi_{k+1}.\end{aligned}$$

Since $\pi_k + \Delta = \pi_{k+1}$, this implies that

$$\gamma_k - \gamma_{k-1} = \gamma_{k+1} - \gamma_k. \quad \square \quad (\text{A.14})$$

Equation (A.14) shows that the rate of change of the deal rate as a function of the pie size is constant whenever the efficiency condition (Equation (A.1)) holds and the incremental change in pie sizes is constant.

We can also use the results in Equations (A.12) and (A.14) to derive deal rates.

For the efficient equilibrium, we set γ_4, γ_5 , and γ_6 equal to 1. Equation (A.12) is then trivially satisfied for pie sizes 4 and higher. Then, Equation (A.12) implies that

$$\gamma_3 = (1 - \gamma_4)4 \quad \gamma_3 = \frac{4}{5},$$

and therefore $\gamma_k - \gamma_{k-1} = \frac{1}{5}$, so that

$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = (\frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1, 1, 1). \quad (\text{A.15})$$

For the equal-split equilibrium, we set $\gamma_6 = 1$; then, we have

$$\gamma_5 = (1 - \gamma_6)6 \quad \gamma_5 = \frac{6}{7},$$

and therefore $\gamma_k - \gamma_{k-1} = \frac{1}{7}$, so that

$$(\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6) = (\frac{2}{7}, \frac{3}{7}, \frac{4}{7}, \frac{5}{7}, \frac{6}{7}, 1), \quad (\text{A.16})$$

which is consistent with our results for the equal-split equilibrium.

Appendix B. Session Information

Summary information of all of the experimental sessions (location, number of subjects, and gender by role) is recapitulated in Table B.1.

Table B.1. Session Information, I-Informed, U-Uninformed

Session no.	Location	Date	N	I Male	I Female	U Male	U Female
1	Caltech	12/1/2011	10	3	2	3	2
2	Caltech	12/8/2011	10	2	3	2	3
3	Caltech	1/9/2012	8	3	1	2	2
4	Caltech	1/11/2012	16	5	3	5	3
5	Caltech	2/28/2012	8	3	1	1	3
6	UCLA	5/11/2012	18	6	3	6	3
7	UCLA	5/11/2012	20	4	6	6	4
8	UCLA	5/11/2012	20	6	4	6	4
Total			110	32	23	31	24

Appendix C. List of Process Features and Associated Marginal Effects

Figure C.1 summarizes all of the process features used to predict bargaining outcomes. We provide further details of calculating some of the features below.


- *Initial difference negative?* A binary indicator that equals 1 if the initial offer of the informed player is greater than the initial uninformed player's demand and zero otherwise.

- *Positions ever matched?* A binary indicator that equals 1 if the players' bargaining positions had previously matched and they later changed their minds.

- *Informed/uninformed first change T.* The first time in the game in which the informed/uninformed player has updated his or her initial bargaining position.

- *T since informed/uninformed last change T.* The time since the last time in which the informed/uninformed player has updated his or her bargaining position.

Figure C.1. (Color online) Bargaining Process Features Used for Outcome Prediction (Deal = 1) and Their Estimated Marginal Effects



- *Informed/uninformed first/last change mag.* The magnitude of the last informed/uninformed position change.
- *# informed/uninformed changes.* The number of times that the informed/uninformed player has changed his or her bargaining position since the start of the game.
- *Informed/mean change mag.* The mean magnitude of change in the informed/uninformed player when he or she changed bargaining positions.
- *First change T.* The first time in the game in which either player has updated his or her initial bargaining position.
- *T since last change.* The time since the last time in which either player has updated his or her bargaining position.
- *Informed/uninformed moved first?* A binary indicator that equals 1 if the informed/uninformed player was the first to change his or her bargaining position in the game.
- *Informed/uninformed weighted avg.* A weighted sum of the informed/uninformed bargaining positions across time.

$$\sum_{t=0}^T w_t x_t, \quad (\text{C.1})$$

such that t denotes time (between 0 and the current time T , sampled in a 0.1 sec resolution) and x_t is bargaining position in time t . The weight w_t equals

$$w_t = \frac{t^2}{\sum_{q=0}^T q^2}, \quad (\text{C.2})$$

where q is an aggregation index. This results a linear combination where later bargaining positions are weighted more heavily than earlier ones.

- *Current informed/uninformed/both are focal?* A binary indicator that equals 1 if the informed/uninformed/both players bargaining positions match the half of either possible pie size (i.e., 0.4, 0.6, 1, 1.4, 1.6, 2, 2.4, 2.6, 3).

Appendix D. Instructions

This is an experiment about bargaining. You will play 120 rounds of a bargaining game.

In the game, one participant (the informed player) is told the total amount of money (pie size) in each round. This amount will be \$1, 2, 3, 4, 5, or 6, chosen randomly in each trial. The amount will appear at the top-left corner of the screen.

The other player is not informed of the pie size.

During each round, participants bargain over the uninformed player's payoff.

The roles are randomly selected and fixed for the duration of the experiment. Before each round, informed and uninformed players are randomly matched.

Participants negotiate by clicking on a scale from \$0 to \$6 (see Figure 1). Amounts on the scale represent the uninformed player's payoff.

During the first two seconds, participants select their initial offers. Note that the initial location of the cursors is random. In the following 10 seconds, the participants bargain, using the mouse to select payoffs for the uninformed player. Clicking on a different part of the scale moves the cursor.

A deal occurs when the cursors are in the same place for 1.5 seconds. When both cursors are in the same place on the scale, a green rectangle will appear (see Figure 2).

If a deal is made, the informed player's payoff is equal to the pie size minus the negotiated uninformed player's payoff. If the agreement exceeds the total amount of money, the payoff will be negative.

If no deal has been made after 10 seconds of bargaining, both participants get \$0.

Following each trial, the uninformed player will be shown of the pie size.

The game has a total of 120 trials.

Before the experiment begins, there will be 15 training trials, to allow you to practice.

At the end of the game, you will receive payment based on randomly selected 15% of your trials.

You will receive a \$5 participation fee in addition to whatever you earn from playing the game.

Quiz

Total amount is \$3. Cursors were matched in \$1. How much money does the informed participant get? How much does the uninformed participant get?

Total amount is \$2. Cursors were matched in \$4.1. How much money does the informed participant get? How much does the uninformed participant get?

One second before the end of the trial, both participants have agreed on payoff of \$2 and the green rectangle appears. What is going to happen when the trial ends?

Both participants have agreed on payoff of \$2 and the green rectangle appears. After one second, the uninformed player changed his offer to \$2.5. What is going to happen?

Endnotes

¹ See Cramton (1984), Fudenberg et al. (1985), Rubinstein (1985), Grossman and Perry (1986), Gul and Sonnenschein (1988), and Ausubel and Deneckere (1993).

² See Ochs and Roth (1989), Camerer et al. (1993), Mitzkewitz and Nagel (1993), Güth 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), and Kriss et al. (2013).

³ See Kristensen and Gärling (1997), Galinsky and Mussweiler (2001), Van Poucke and Buelens (2002), Mason et al. (2013), and Ames and Mason (2015).

⁴ For other recent "protocol-free" bargaining theories, see Fuchs and Skrzypacz (2013), Fanning (2016), and Simsek and Yildiz (2016).

⁵ For reviews, see Kennan and Wilson (1993), Ausubel et al. (2002), and Thompson et al. (2010).

⁶ See Schelling (1960), Roth (1985), Kristensen and Gärling (1997), Janssen (2001), Binmore and Samuelson (2006), Janssen (2006), Bardsley et al. (2010), Isoni et al. (2013, 2014), Hargreaves Heap et al. (2014), Luhan et al. (2017), and Sontuoso and Bhatia (2017).

⁷ 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. 2010).

⁸ Anbarci and Feltoich (2013, 2016), Karagözoğlu and Riedl (2015), and Luhan et al. (2017) are other recent experimental papers that investigated unstructured bargaining under various information conditions.

⁹ See Kritikos and Bolle (2001), Charness and Rabin (2002), Engelmann and Strobel (2004), Engelmann and Strobel (2006), Fehr et al. (2006), Bolton and Ockenfels (2006), Durante et al. (2014), and El Harbi et al. (2015).

¹⁰For 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).

¹¹See Yükl (1974b), 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), and Ames and Mason (2015).

¹²See also the discussion in Kennan and Wilson (1993, section 6.1).

¹³Many possible explanations have been proposed for the prevalence of equal splitting, including social norms (Andreoni and Bernheim 2009), pure dislike of unequal distributions (Fehr and Schmidt 1999), beliefs about the preferences of one's partner (Chmura et al. 2005), and evolutionary stability (Bolton 1997, Young 1993).

¹⁴This result is used in the derivation of the strike condition for Lemma 2. See Equation (A.1) in Appendix A. See also equations (1) and (2) of Forsythe et al. (1991).

¹⁵The multiple of \$0.2 was chosen as 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). All of the experimental results are consistent when considering only the even pie games, in which subjects can exactly agree on half of the pie.

¹⁶A video demonstration of the task is available on <https://www.youtube.com/watch?v=y7pKh1EJsvM&> (published November 27, 2014).

¹⁷We chose 10 seconds as a long enough time period to observe the bargaining process, and short enough to collect a large number of observations per subject in each one of the six pie sizes. Given likely subject fatigue, there is an obvious trade-off between having a lot of learning across many sessions (especially with six pie sizes, which permits nuanced prediction about deal rates) and then having to make the sessions short. We have piloted longer time periods, and the results were not substantially different. Also, the participants have 15 practice rounds, making them familiar with the fast bargaining process.

¹⁸A small fraction (less than 2.5%) of the games were excluded from analysis because of a software bug in the first sessions.

¹⁹In our experimental interface, players communicated their bids in integer multiples of 0.2 and therefore could not make offers of exactly 0.5, 1.5, or 2.5. We consider offers that are within 0.1 of these values (that are as close as one could get to them) as matching half of integer pies.

²⁰The deal is closed 1.5 seconds after the bargainers' positions have matched.

²¹We employ linear probability models for the ease of testing and interpreting coefficient values and interpreting the effects of control variables; logistic regressions give similar results.

²²Hansen (2007) shows that the cluster-robust variance estimator is appropriate when the number of groups is small, as long as the number of within-group observations is large. See the discussions in Angrist and Pischke (2008, chap. 8) and Cameron and Miller (2015).

²³With session controls and standard errors clustered at the session level, a test of the full model would require 10 constraints (3 for the model coefficients, and 1 for each of the session controls), which is more than our degrees of freedom (7). We therefore report results from tests of the joint hypotheses on $(\alpha_0, \alpha_1, \alpha_2)$ on the pooled regression model (Model A), using standard errors clustered at the session level. An F -test of the null hypothesis $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{5}, \frac{1}{5}, -\frac{1}{5})$ rejects the null at $p < 0.001$, with $F(3, 7) = 134.78$. An F -test of the null hypothesis $(\alpha_0, \alpha_1, \alpha_2) = (\frac{1}{7}, \frac{1}{7}, 0)$ rejects the null at $p < 0.001$, with $F(3, 7) = 78.19$.

²⁴For instance, the deal rate when $\pi = 6$ is 0.83 in the last 20 rounds, pooling the data from all sessions.

²⁵In some interesting models of bargaining with self-interested preferences, and under certain experimental conditions, strikes can occur even with complete information (e.g., Roth and Malouf 1979, Roth et al. 1981, Roth and Murnighan 1982, Roth 1985, Haller and Holden 1990, Herreiner and Puppe 2004, Gächter and Riedl 2005, Gächter and Riedl 2006, Embrey et al. 2014). If the forces operating in such models and environments also apply in our private-information settings, the strike rates could be larger than those predicted by the mechanism design approach.

²⁶Pooling the demands makes sense because the uninformed players have no information regarding the realization of the pie that might be deduced from the behavior of the informed player at the initial-offer stage.

²⁷A \$1 increase corresponds to the distance between the 25th percentile (\$0.4) and the 75th percentile (\$1.4) of the distribution of initial offers.

²⁸By the revelation principle, every equilibrium in our setting corresponds to a payoff-equivalent equilibrium of the direct mechanism. As the direct mechanism is "process free," process features should not have predictive power in equilibrium after controlling for pie size.

²⁹A LASSO-penalized logistic regression maximizes the standard logistic regression log-likelihood function minus a penalty term equal to a weighted sum of their absolute values of the regression coefficients (their L_1 norm) to overcome potential overfitting of the training data.

³⁰To capture a nonlinear dependency of the deal rate on the pie size, the pie size is represented using five dummy variables, denoting the different pies (such that the \$1 pie is the baseline category).

³¹We included only trials that were still in progress (when a deal has not yet been achieved), and excluded trials in which the offer and demand were equal at the relevant time stamp.

³²The optimization procedure was performed using the "lassoglm" function implemented in MATLAB under its default setting. Thus, we first estimate $\lambda^{\max}$, the largest value of the penalty parameter λ that gives a nonnull model, and perform optimization by exploring a geometric sequence of 100 values between $0.0001\lambda^{\max}$ and $\lambda^{\max}$.

³³As the LASSO procedure is sensitive to the scale of the inputs, all independent variables were standardized (z-scored) prior to model training.

³⁴We used decision thresholds between 0 and 1 on a grid with a resolution of 0.01.

³⁵The AUC is closely related to the Mann–Whitney–Wilcoxon U -statistic (Hanley and McNeil 1982).

³⁶We formally tested the predictive accuracy of process data above and beyond the pie size by comparing the mean square error of the model that was trained using both process data and pie to the model that was trained using only the pie. Paired t -tests of the squared errors from both models (with adjustments for clustering at the session level) showed that adding process data significantly decreased out-of-sample prediction error for all times greater than or equal to three seconds (at the 5% level, Bonferroni corrected).

³⁷The "post-LASSO" procedure consisted of three steps. First, we optimized the LASSO tuning parameter λ using 10-fold cross-validation on the entire data set. Second, we conducted model selection by fitting a logistic LASSO regression using the optimized tuning parameter to the data. Finally, we fitted an ordinary logistic regression to the data, using the features with nonzero LASSO coefficients from the second stage.

³⁸When the design matrix includes several highly correlated variables that are related to the response variable (as in the case at hand), LASSO tends to pick only one of them and shrinks the influence of the rest to zero (Tibshirani 2011). Thus, LASSO might discard

additional weak (but important) predictors that are highly correlated with stronger predictors, which may also prove to be useful for understanding how bargaining process influences outcome. Crucially, previous investigations of LASSO's performance when multicollinearity is present (van de Geer et al. 2013, Hebiri and Lederer 2013, Oyeyemi et al. 2015) have shown that predictive accuracy is not harmed by multicollinearity, and even if this had been an issue, it would have biased the predictive accuracy of the process features downward. Thus, the predictive accuracy of process features that we report can be seen as a lower bound to what can potentially be achieved.

References

- Ames DR, Mason MF (2015) Tandem anchoring: Informational and politeness effects of range offers in social exchange. *J. Personality Soc. Psych.* 108(2):254.
- Anbarci N, Feltovich N (2013) How sensitive are bargaining outcomes to changes in disagreement payoffs? *Experiment. Econom.* 16(4):560–596.
- Anbarci N, Feltovich N (2016) Market institutions, prices and distribution of surplus: A theoretical and experimental investigation. Working paper, Deakin Business School, Geelong VIC, Australia.
- Andreoni J, Bernheim BD (2009) Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica* 77(5):1607–1636.
- Angrist JD, Pischke J-S (2008) *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, Princeton, NJ).
- Ausubel LM, Deneckere RJ (1993) Efficient sequential bargaining. *Rev. Econom. Stud.* 60(2):435–461.
- Ausubel LM, Cramton P, Deneckere RJ (2002) Bargaining with incomplete information. Aumann R, Hart S, eds. *Handbook of Game Theory with Economic Applications*, Vol. 3 (North-Holland, Amsterdam), 1897–1945.
- Bajari P, Nekipelov D, Ryan SP, Yang M (2015) Machine learning methods for demand estimation. *Amer. Econom. Rev.* 105(5):481–485.
- Bardolet D, Fox CR, Lovallo D (2011) Corporate capital allocation: A behavioral perspective. *Strategic Management J.* 32(13):1465–1483.
- Bardsley N, Mehta J, Starmer C, Sugden R (2010) Explaining focal points: Cognitive hierarchy theory versus team reasoning. *Econom. J.* 120(543):40–79.
- Behrman JR, Rosenzweig MR (2004) Parental allocations to children: New evidence on bequest differences among siblings. *Rev. Econom. Statist.* 86(2):637–640.
- Belloni A, Chernozhukov V (2013) Least squares after model selection in high-dimensional sparse models. *Bernoulli* 19(2):521–547.
- Belloni A, Chen D, Chernozhukov V, Hansen C (2012) Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica* 80(6):2369–2429.
- Binmore K, Samuelson L (2006) The evolution of focal points. *Games Econom. Behav.* 55(1):21–42.
- Blanco M, Engelmann D, Normann HT (2011) A within-subject analysis of other-regarding preferences. *Games Econom. Behav.* 72(2):321–338.
- Bolton GE (1997) The rationality of splitting equally. *J. Econom. Behav. Organ.* 32(3):365–381.
- Bolton GE, Ockenfels A (2006) Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *Amer. Econom. Rev.* 96(5):1906–1911.
- Bradley AP (1997) The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition* 30(7):1145–1159.
- Buelens M, Van Poucke D (2004) Determinants of a negotiator's initial opening offer. *J. Bus. Psych.* 19(1):23–35.
- Cai H, Wang JT-Y (2006) Overcommunication in strategic information transmission games. *Games Econom. Behav.* 56(1):7–36.
- Camerer CF, Johnson E, Rymon T, Sen S (1993) Cognition and framing in sequential bargaining for gains and losses. Binmore K, Kirman A, Tani P, eds. *Frontiers of Game Theory* (MIT Press, Cambridge, MA), 27–47.
- Cameron AC, Miller DL (2015) A practitioner's guide to cluster-robust inference. *J. Human Resources* 50(2):317–372.
- Charness G, Rabin M (2002) Understanding social preferences with simple tests. *Quart. J. Econom.* 117(3):817–869.
- Chertkoff JM, Conley M (1967) Opening offer and frequency of concession as bargaining strategies. *J. Personality Soc. Psych.* 7(2, Part 1):181.
- Chmura T, Kube S, Pitz T, Puppe C (2005) Testing (beliefs about) social preferences: Evidence from an experimental coordination game. *Econom. Lett.* 88(2):214–220.
- Cramton PC (1984) Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *Rev. Econom. Stud.* 51(4):579–593.
- Crawford VP (2003) Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. *Amer. Econom. Rev.* 93(1):133–149.
- Crosen R, Boles T, Murnighan JK (2003) Cheap talk in bargaining experiments: Lying and threats in ultimatum games. *J. Econom. Behav. Organ.* 51(2):143–159.
- Cuzick J (1985) A Wilcoxon-type test for trend. *Statist. Medicine* 4(4):543–547.
- Daniel TE, Seale DA, Rapoport A (1998) Strategic play and adaptive learning in the sealed-bid bargaining mechanism. *J. Math. Psych.* 42(2):133–166.
- De Bruyn A, Bolton GE (2008) Estimating the influence of fairness on bargaining behavior. *Management Sci.* 54(10):1774–1791.
- Durante R, Putterman L, Weele J (2014) Preferences for redistribution and perception of fairness: An experimental study. *J. Eur. Econom. Assoc.* 12(4):1059–1086.
- Dzyabura D, Hauser JR (2011) Active machine learning for consideration heuristics. *Marketing Sci.* 30(5):801–819.
- Einav L, Levin J (2014) Economics in the age of big data. *Science* 346(6210):1243089.
- El Harbi S, Bekir I, Grolleau G, Sutan A (2015) Efficiency, equality, positionality: What do people maximize? Experimental vs. hypothetical evidence from Tunisia. *J. Econom. Psych.* 47(April):77–84.
- Embrey M, Hyndman KB, Riedl A (2014) Bargaining with a residual claimant: An experimental study. Working paper, University of Sussex, Brighton, UK.
- Engelmann D, Strobel M (2004) Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *Amer. Econom. Rev.* 94(4):857–869.
- Engelmann D, Strobel M (2006) Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply. *Amer. Econom. Rev.* 96(5):1918–1923.
- Fanning J (2016) Reputational bargaining and deadlines. *Econometrica* 84(3):1131–1179.
- Fehr E, Schmidt KM (1999) A theory of fairness, competition, and cooperation. *Quart. J. Econom.* 114(3):817–868.
- Fehr E, Naef M, Schmidt KM (2006) Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *Amer. Econom. Rev.* 96(5):1912–1917.
- Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. *Experiment. Econom.* 10(2):171–178.
- Forsythe R, Kennan J, Sopher B (1991) An experimental analysis of strikes in bargaining games with one-sided private information. *Amer. Econom. Rev.* 81(1):253–278.
- Friedman D, Oprea R (2012) A continuous dilemma. *Amer. Econom. Rev.* 102(1):337–363.
- Fuchs W, Skrzypacz A (2013) Bargaining with deadlines and private information. *Amer. Econom. J. Microeconom.* 5(4):219–243.
- Fudenberg D, Levine D, Tirole J (1985) Infinite-horizon models of bargaining with one-sided incomplete information. Roth AE, ed. *Game-Theoretic Models of Bargaining*, Chap. 5 (Cambridge University Press, Cambridge, UK), 73–98.
- Gächter S, Riedl A (2005) Moral property rights in bargaining with infeasible claims. *Management Sci.* 51(2):249–263.
- Gächter S, Riedl A (2006) Dividing justly in bargaining problems with claims: Normative judgments and actual negotiations. *Soc. Choice Welfare* 27(3):571–594.
- Galeotti F, Montero M, Poulsen A (2015) Efficiency versus equality in bargaining. SSRN 2688599.

- Galinsky AD, Mussweiler T (2001) First offers as anchors: The role of perspective-taking and negotiator focus. *J. Personality Soc. Psych.* 81(4):657.
- Gneezy U, Rockenbach B, Serra-Garcia M (2013) Measuring lying aversion. *J. Econom. Behav. Organ.* 93:293–300.
- Greene WH (2003) *Econometric Analysis* (Pearson Education, Upper Saddle River, NJ).
- Grossman SJ, Perry M (1986) Sequential bargaining under asymmetric information. *J. Econom. Theory* 39(1):120–154.
- Gul F, Sonnenschein H (1988) On delay in bargaining with one-sided uncertainty. *Econometrica J. Econometric Soc.* 56(3):601–611.
- Güth W, Van Damme E (1998) Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *J. Math. Psych.* 42(2):227–247.
- Güth W, Huck S, Ockenfels P (1996) Two-level ultimatum bargaining with incomplete information: An experimental study. *Econom. J.* 106(436):593–604.
- Haller H, Holden S (1990) A letter to the editor on wage bargaining. *J. Econom. Theory* 52(1):232–236.
- Hanley JA, McNeil BJ (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* 143(1):29–36.
- Hansen CB (2007) Asymptotic properties of a robust variance matrix estimator for panel data when t is large. *J. Econometrics* 141(2):597–620.
- Hargreaves Heap S, Rojo Arjona D, Sugden R (2014) How portable is level-0 behavior? A test of level- k theory in games with non-neutral frames. *Econometrica* 82(3):1133–1151.
- Hebiri M, Lederer J (2013) How correlations influence lasso prediction. *IEEE Trans. Inform. Theory* 59(3):1846–1854.
- Herreiner DK, Puppe C (2004) Equitable allocations in experimental bargaining games: Inequality a version versus efficiency. Technical report, University of Bonn, Bonn, Germany.
- Holmström B, Myerson RB (1983) Efficient and durable decision rules with incomplete information. *Econometrica J. Econometric Soc.* 51(6):1799–1819.
- Isoni A, Poulsen A, Sugden R, Tsutsui K (2013) Focal points in tacit bargaining problems: Experimental evidence. *Eur. Econom. Rev.* 59:167–188.
- Isoni A, Poulsen A, Sugden R, Tsutsui K (2014) Efficiency, equality and labelling: An experimental investigation of focal points in explicit bargaining. *Amer. Econom. Rev.* 104(10):3256–3287.
- Jacquemet N, Zylbersztein A (2014) What drives failure to maximize payoffs in the lab? A test of the inequality aversion hypothesis. *Rev. Econom. Design* 18(4):243–264.
- Janssen MCW (2001) Rationalizing focal points. *Theory Decision* 50(2):119–148.
- Janssen MCW (2006) On the strategic use of focal points in bargaining situations. *J. Econom. Psych.* 27(5):622–634.
- Johnson EJ, Camerer C, Sen S, Rymon T (2002) Detecting failures of backward induction: Monitoring information search in sequential bargaining. *J. Econom. Theory* 104(1):16–47.
- Kagel JH, Wolfe KW (2001) Tests of fairness models based on equity considerations in a three-person ultimatum game. *Experiment. Econom.* 4(3):203–219.
- Kagel JH, Kim C, Moser D (1996) Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games Econom. Behav.* 13(1):100–110.
- Kalai E, Smorodinsky M (1975) Other solutions to Nash's bargaining problem. *Econometrica J. Econometric Soc.* 43(3):513–518.
- Karagözoğlu E, Riedl A (2015) Performance information, production uncertainty, and subjective entitlements in bargaining. *Management Sci.* 61(11):2611–2626.
- Kennan J (1986) The economics of strikes. Ashenfelter OC, Layard R, eds. *Handbook of Labor Economics*, Vol. 2, Chap. 19 (Elsevier B.V., Amsterdam), 1091–1137.
- Kennan J, Wilson R (1990) Can strategic bargaining models explain collective bargaining data? *Amer. Econom. Rev.* 80(2):405–409.
- Kennan J, Wilson RB (1993) Bargaining with private information. *J. Econom. Literature* 31(1):45–104.
- Kotha R, Kim Y-C, Camerer C, Lovo D (2015) Equality norms, incentives and reciprocity in scientific royalty contracts. Working paper, California Institute of Technology, Pasadena.
- Krajbich I, Camerer C, Ledyard J, Rangel A (2009) Using neural measures of economic value to solve the public goods free-rider problem. *Science* 326(5952):596–599.
- Kriss PH, Nagel R, Weber RA (2013) Implicit vs. explicit deception in ultimatum games with incomplete information. *J. Econom. Behav. Organ.* 93:337–346.
- Kristensen H, Gärling T (1997) The effects of anchor points and reference points on negotiation process and outcome. *Organ. Behav. Human Decision Processes* 71(1):85–94.
- Kritikos A, Bolle F (2001) Distributional concerns: Equity-or efficiency-oriented? *Econom. Lett.* 73(3):333–338.
- López-Pérez R, Pintér Á, Kiss HJ (2015) Does payoff equity facilitate coordination? A test of Schelling's conjecture. *J. Econom. Behav. Organ.* 117(September):209–222.
- Luhan W, Poulsen A, Roos M (2017) Real-time tacit bargaining, payoff focality, and coordination complexity: Experimental evidence. *Games Econom. Behav.* 102(March):687–699.
- Mason MF, Lee AJ, Wiley EA, Ames DR (2013) Precise offers are potent anchors: Conciliatory counteroffers and attributions of knowledge in negotiations. *J. Experiment. Soc. Psych.* 49(4):759–763.
- Menchik PL (1980) Primogeniture, equal sharing, and the U.S. distribution of wealth. *Quart. J. Econom.* 94(2):299–316.
- Mitzkewitz M, Nagel R (1993) Experimental results on ultimatum games with incomplete information. *Internat. J. Game Theory* 22(2):171–198.
- Mullainathan S, Spiess J (2017) Machine learning: An applied econometric approach. *J. Econom. Perspect.* 31(2):87–106.
- Myerson R (1979) Incentive compatibility and the bargaining problem. *Econometrica* 47(1):61–73.
- Myerson RB (1984) Two-person bargaining problems with incomplete information. *Econometrica J. Econometric Soc.* 52(2):461–488.
- Nash JF Jr (1950) The bargaining problem. *Econometrica J. Econometric Soc.* 18(2):155–162.
- Nave G, Minxha J, Greenberg DM, Kosinski M, Stillwell D, Rentfrow J (2018) Musical preferences predict personality: Evidence from active listening and Facebook likes. *Psych. Sci.*, DOI: 10.1177/0956797618761659.
- Nydegger RV, Owen G (1974) Two-person bargaining: An experimental test of the Nash axioms. *Internat. J. Game Theory* 3(4):239–249.
- Ochs J, Roth AE (1989) An experimental study of sequential bargaining. *Amer. Econom. Rev.* 79(3):355–384.
- Oyeyemi GM, Ogunjobi EO, Folorunsho AI (2015) On performance of shrinkage methods—A Monte Carlo study. *Internat. J. Statist. Appl.* 5(2):72–76.
- Pruitt DG (2013) *Negotiation Behavior* (Academic Press, New York).
- Radner R, Schotter A (1989) The sealed-bid mechanism: An experimental study. *J. Econom. Theory* 48(1):179–220.
- Rapoport A, Fuller MA (1995) Bidding strategies in a bilateral monopoly with two-sided incomplete information. *J. Math. Psych.* 39(2):179–196.
- Rapoport A, Daniel TE, Seale DA (1998) Reinforcement-based adaptive learning in asymmetric two-person bargaining with incomplete information. *Experiment. Econom.* 1(3):221–253.
- Rapoport A, Erev I, Zwick R (1995) An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Sci.* 41(3):377–394.
- Roth AE (1985) Toward a focal point theory of bargaining. Roth AE, ed. *Game-Theoretic Models of Bargaining* (Cambridge University Press, Cambridge, UK), 259–268.
- Roth AE, Malouf MW (1979) Game-theoretic models and the role of information in bargaining. *Psych. Rev.* 86(6):574–594.
- Roth AE, Murnighan JK (1982) The role of information in bargaining: An experimental study. *Econometrica J. Econometric Soc.* 50(5):1123–1142.
- Roth AE, Malouf MW, Murnighan JK (1981) Sociological versus strategic factors in bargaining. *J. Econom. Behav. Organ.* 2(2):153–177.

- Roth AE, Murnighan JK, Schoumaker F (1988) The deadline effect in bargaining: Some experimental evidence. *Amer. Econom. Rev.* 78(4):806–823.
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica: J. Econometric Soc.* 50(1):97–109.
- Rubinstein A (1985) A bargaining model with incomplete information about time preferences. *Econometrica: J. Econometric Soc.* 53(5):1151–1172.
- Schelling TC (1960) *The Strategy of Conflict* (Harvard University Press, Cambridge, MA).
- Simsek A, Yildiz M (2016) Durability, deadline, and election effects in bargaining. Technical report, National Bureau of Economic Research, Cambridge, MA.
- Smith A, Bernheim BD, Camerer CF, Rangel A (2014) Neural activity reveals preferences without choices. *Amer. Econom. J. Microeconom.* 6(2):1–36.
- Sontuoso A, Bhatia S (2017) A notion of prominence for strategic games. Technical report, Philosophy, Politics and Economics, University of Pennsylvania.
- Srivastava J (2001) The role of inferences in sequential bargaining with one-sided incomplete information: Some experimental evidence. *Organ. Behav. Human Decision Processes* 85(1):166–187.
- Ståhl I (1972) Bargaining theory. Ph.D. thesis, Economic Research Institute at the Stockholm School of Economics (EFI).
- Stone M (1974) Cross-validatory choice and assessment of statistical predictions. *J. Roy. Statist. Soc. Ser. B (Methodological)* 36(2):111–147.
- Thompson LL, Wang J, Gunia BC (2010) Negotiation. *Annual Rev. Psych.* 61:491–515.
- Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B (Methodological)* 267–288.
- Tibshirani R (2011) Regression shrinkage and selection via the lasso: A retrospective. *J. Roy. Statist. Soc. Ser. B (Statist. Methodology)* 73(3):273–282.
- Valley K, Thompson L, Gibbons R, Bazerman MH (2002) How communication improves efficiency in bargaining games. *Games Econom. Behav.* 38(1):127–155.
- van de Geer S, Lederer J, et al. (2013) The lasso, correlated design, and improved oracle inequalities. *From Probability to Statistics and Back: High-Dimensional Models and Processes—A Festschrift in Honor of Jon A. Wellner* (Institute of Mathematical Statistics, Beachwood, OH), 303–316.
- Van Poucke D, Buelens M (2002) Predicting the outcome of a two-party price negotiation: Contribution of reservation price, aspiration price and opening offer. *J. Econom. Psych.* 23(1):67–76.
- Varian HR (2014) Big data: New tricks for econometrics. *J. Econom. Perspect.* 28(2):3–27.
- Wang JT-Y, Spezio M, Camerer CF (2010) Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *Amer. Econom. Rev.* 100(3):984–1007.
- White SB, Neale MA (1994) The role of negotiator aspirations and settlement expectancies in bargaining outcomes. *Organ. Behav. Human Decision Processes* 57(2):303–317.
- White SB, Valley KL, Bazerman MH, Neale MA, Peck SR (1994) Alternative models of price behavior in dyadic negotiations: Market prices, reservation prices, and negotiator aspirations. *Organ. Behav. Human Decision Processes* 57(3):430–447.
- Yarkoni T, Westfall J (2017) Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspect. Psych. Sci.* 12(6):1100–1122.
- Young HP (1993) An evolutionary model of bargaining. *J. Econom. Theory* 59(1):145–168.
- Young HP, Burke MA (2001) Competition and custom in economic contracts: A case study of Illinois agriculture. *Amer. Econom. Rev.* 91(3):559–573.
- Youyou W, Kosinski M, Stillwell D (2015) Computer-based personality judgments are more accurate than those made by humans. *Proc. Natl. Acad. Sci. USA* 112(4):1036–1040.
- Yukl G (1974a) Effects of the opponent's initial offer, concession magnitude and concession frequency on bargaining behavior. *J. Personality Soc. Psych.* 30(3):323.
- Yukl GA (1974b) Effects of situational variables and opponent concessions on a bargainer's perception, aspirations, and concessions. *J. Personality Soc. Psych.* 29(2):227–236.