Fairness in Incomplete Information Bargaining: Theory and Widespread Evidence from the Field*

Daniel Keniston, Bradley J. Larsen, Shengwu Li, J.J. Prescott,
Bernardo S. Silveira, and Chuan Yu†

September 21, 2023

Abstract

This paper documents a robust pattern from diverse sequential bargaining settings: agents favor offers that split the difference between the previous two offers. Our empirical settings include used-cars, insurance claims, home sale, trade tariffs, a TV game show, eBay, and auto-rickshaws. These even-split offers are more likely to be accepted, less likely to spur exit by the opponent, and more likely to be followed by subsequent split-the-difference offers if bargaining continues. We propose several theoretical frameworks to explain this behavior, including an inference argument under which split-the-difference offers can be viewed as an equal split of the potential surplus.

JEL Codes: C7, D8, D9

Keywords: Bargaining, negotiation, fairness, split-the-difference, incomplete information, inference, alternating offers

*We thank Panle Jia Barwick, Kalyan Chatterjee, Peter Cramton, Jack Fanning, Emin Karagözoğlu, Ted O'Donoghue, Marek Pycia, Al Roth, Caio Waisman, Alex Wolitzky, and seminar and conference participants at LSU, UCLA, NYUAD, IACM, and BEET for helpful comments and suggestions. Wendy Yin provided excellent research assistance. We thank an anonymous real estate brokerage, as well as Iñigo Hernandez-Arenaz and Nagore Iriberri, for providing data.

†Keniston: Louisiana State University, Department of Economics and NBER; dkeniston@lsu.edu; Larsen: Washington University in St. Louis, Olin School of Business and NBER; blarsen@wustl.edu; Li: Harvard University, Department of Economics; shengwu_li@fas.harvard.edu; Prescott: University of Michigan Law School; jprescott@umich.edu; Silveira: UCLA Department of Economics; silveira@econ.ucla.edu; Yu: Stanford University, Department of Economics; chuanyu@stanford.edu

1 Introduction

The role of fairness notions in bilateral bargaining has been widely accepted by practitioners and explored in depth in laboratory experiments. A number of studies document such influences, in particular demonstrating a bias toward an equal split of a known pie; see Camerer (2011) for a review. Little is known, however, about how these norms play out in the field, where assumptions of complete information common in laboratory experiments are unlikely to hold. In this paper, we document a largely understudied fact: agents in real-world sequential bargaining settings favor offers that split the difference between the two most recent offers on the table.

The term *split the difference* appears in previous research (especially experimental work) on bargaining with different meanings. Binmore et al. (1989), for example, use it to describe bargaining outcomes in which the players divide a known surplus net of their respective outside options. In the present paper, similar to Bochet et al. (2023) and Backus et al. (2020), splitting the difference refers instead to a notion in a sequential-offer game where the current bargaining offer lies at the midpoint between the two most recent offers (the most recent offer of the proposer and the most recent offer of the counterparty). As an example, suppose that in an alternating-offer game, at some point when it is the seller's turn, the seller proposes \$100, then the buyer proposes \$50, and then it is the seller's turn again. We refer to the \$100 and \$50 offers as the *two most recent offers*, and, if the seller next proposes \$75, we refer to this as a *split-the-difference* offer.

Our empirical evidence comes from novel, detailed data on sequential offers from several vastly different bargaining contexts: business-to-business negotiations for used cars in the U.S., pre-trial settlement bargaining on insurance injury claims in the U.S., street negotiations from a quirky TV game show in Spain, bargaining for auto rickshaw rides in India, international trade tariff bargaining, online retail negotiations from eBay.com, and bargaining over housing through a real estate dealer. In each setting, we find strong evidence of split-the-difference offers—a clear mode at the 50-50 point between the two most recent offers. To our knowledge, the widespread nature of this phenomenon—across very different bargaining settings—was previously undocumented.[1]

[1]Our motivation for selecting these particular settings is simple: we use every dataset available to us in which all back-and-forth offers are recorded, because only with such data can we examine split-the-difference behavior. Only recently did researchers start using this type of detailed bargaining data. As such, split-the-difference offers have been naturally understudied thus far, and several of the datasets we study are

We find that split-the-difference offers are more likely to be accepted: as an agent concedes more to the opponent, the agent's offer is generally more likely to be accepted; this probability, however, jumps discontinuously at offers that lie halfway between the two most recent offers, to the extent that a split-the-difference offer is even more likely to be accepted than a slightly more generous offer that concedes more than half of the distance between offers.[2] Similarly, offers that split the difference are less likely to result in the opponent walking away, ending the negotiations.

In the same vein, we show that certain agents are more likely than others to make split-the-difference offers, and that a seller is less likely to propose a 50-50 split if one of the two most recent offers is an extremely low buyer offer—suggesting that the 50-50 norm is constrained by the disparity between the two previous offers.

We also demonstrate that it is indeed the two most recent offers in a bargaining sequence to which players gravitate: proposing an offer at a later stage of the game that splits the difference between offers other than the previous two is far less common. Similarly, split-the-difference offers are substantially more prevalent than offers based on anchor points that are privately known to the proposer, such as the secret auto-accept and auto-decline prices set by the seller in the eBay setting or the loss estimate set by the defendant in the insurance settlement setting. That is, agents systematically split the difference between the previous two offers, which are common knowledge, and thus engage in a behavior that is subject to the scrutiny of all agents involved.

Having demonstrated that split-the-difference offers are ubiquitous, we ask what model could explain this behavior. The prevalence of equal splits of a *commonly known surplus* has been widely established in lab experiments. Researchers often attribute these patterns to fairness-related norms.[3] But it is far from obvious that the patterns that we document are

new to the literature. The housing dataset is completely new. The used-car dataset is largely new; Larsen (2021) uses some observations from this same setting, but here we take advantage of a larger dataset and more information on alternating offers that Larsen (2021) does not use. The data on pre-trial settlement negotiations is also largely new; Prescott et al. (2014) study a small subset of this data to examine high-low contracts in litigation. Data on auto rickshaw rides comes from Keniston (2011), data on the Spanish TV show comes from Hernandez-Arenaz and Iriberri (2018), international tariff negotiation data comes from Bagwell et al. (2020), and the eBay data was first analyzed by Backus et al. (2020).

[2]This finding bears a close resemblance to results from the experimental literature on ultimatum games showing that (i) respondents are more prone to accept equal divisions of the pie than divisions that are more favorable to the respondent (Bellemare et al., 2008); and (ii) respondents are discontinuously more likely to accept 50-50 offers (Lin et al., 2020).

[3]50-50 splits are a common outcome in relatively simple experiments, such as those involving the dictator and ultimatum games, as well as in less structured bargaining studies allowing the subjects to freely negotiate over the division of a pie. For excellent reviews of this vast literature, see Roth (1995) and Camerer (2011).

associated with fairness, as the surplus in our setting depends on agents' *private information*. In principle, split-the-difference offers are "fair" only in the sense that they lie halfway between the possible range of likely subsequent offers given the current offer history; they do not necessarily constitute an equitable split of the actual surplus available to the agents. In fact, unless both parties have made offers that differ by exactly the same amount from their true valuations, splitting the difference between offers cannot yield an equal split of the surplus.

We introduce several theoretical results to explore this question in the context of a bargaining game with alternating offers and two-sided private information. In this context, we first show that a pattern of split-the-difference offers may arise in a perfect Bayesian equilibrium even when the players have no preferences for "fair" outcomes. However, this equilibrium is far from being unique. Thus, while split-the-difference offers are consistent with "standard" game theory, they are by no means *predicted* by it.

We then propose an explanation for how agents may indeed view an equal split of the two most recent offers as "fair." Our argument relies on agents making inferences about the support of their opponent's valuation based on the fact that the opponent did not accept the agent's previous offer. Consider a case where a seller initially proposes a price of \$100. The buyer rejects this offer and counters at \$50. What can the seller infer about the buyer's valuation from the fact that the buyer *rejected* \$100? We show that, for all buyer types (i.e., buyers with some valuation) below \$100, rejecting the \$100 offer is a *sequential best response* (see Battigalli and Siniscalchi 2002) to *every* belief that the buyer may hold about the seller's type. And we show that, for all buyer types above \$100, accepting is a sequential best response to *some* belief a buyer could hold. In this sense, \$100 is the highest buyer type that the seller cannot rule out (1) based on the buyer's rejection of \$100, and (2) under the assumption that the seller knows the buyer is behaving rationally according to some belief. A similar argument applies to the buyer's inference about the seller's valuation if the seller rejects the buyer's counteroffer of \$50. Together, these arguments offer an explanation for how a buyer and a seller might jointly view the gap between the two most recent offers—$[50, 100]$—as the *potential surplus*; and hence an offer of \$75, halfway between the two most recent offers, can be viewed as "fair," an equal split of the potential surplus. We also discuss a number of other possible explanations for split-the-difference behavior.

Two existing studies within economics analyze a similar notion of splitting the difference.[4] Backus et al. (2020) provide descriptive evidence of several facets of eBay bargaining, documenting the prevalence of split-the-difference offers being made and accepted on eBay. We incorporate their data as one of our seven settings, and we adopt their framework as a starting point for our analysis.[5] Bochet et al. (2023) design a lab experiment in which agents with private information negotiate over multiple issues. The authors find that an alternating-offer regime arises endogenously, with proposals frequently splitting the difference between the two most recent offers whether agents negotiate over single items or bundles. We complement the analyses of these two papers by establishing that split-the-difference offers constitute a convention—not an ironclad convention, but a notable convention—among bargainers in the field, and that this behavior is not unique to consumer interactions nor to U.S. negotiation culture. Similarly, this is not a feature of small-stakes negotiations only; it also appears in home sales, the largest lifetime purchase for many buyers. Moreover, our paper establishes a possible grounding for why such behavior could be viewed as "fair" in some sense—even though the two most recent offers are themselves endogenous equilibrium objects and thus the halfway point does not actually correspond to an equal split of the pie. But we stress that we do not view this model as the only possible explanation for the behavior.

Outside economics, the negotiation literature broadly acknowledges the norm of "even-split" offers, although it contains differing assessments of the effectiveness of the strategy. Some authors (Babcock and Laschever, 2008) argue that "split the difference is a tactic that often works extremely well" (p. 212), while others (Thompson, 2020) discourage even-split offers, a perspective exemplified by the best-selling guide to negotiation titled *Never Split the Difference* (Voss and Tahl, 2016). Nevertheless, the negotiation literature universally argues that split-the-difference offers operate through an appeal to fairness that may increase the

[4]Many experiments in the lab involve alternating-offer bargaining, but they do not address the questions we study in this paper; see, for example, Binmore et al. (1985), Ochs and Roth (1989), and Binmore et al. (1989). These experiments have investigated issues such as time discounting and the relevance of outside offers in bargaining. Andreoni and Bernheim (2009) focus on bargaining with complete information and offer a model and experimental results demonstrating how an equal split of a known pie can arise from a preference to appear equitable (a concern for *social image*). Other related studies include Roth and Malouf (1979) and Roth (1985).

[5]Several other studies have also examined this same eBay bargaining data, such as Green and Plunkett (2022), who analyze how well human agents perform relative to reinforcement learning bots coded to respond optimally to observed actions in the data, and Freyberger and Larsen (2021), who estimate a structural model bounding private valuation distributions and bargaining inefficiency.

probability of an offer being accepted by one's bargaining partner.[6] Popular negotiation advice echos this motivation: Kwame Christian, lawyer and founder/CEO of the American Negotiation Institute, a consulting firm offering training and coaching in negotiation, argues that the primary benefit of a split-the-difference offer is that "it can trigger reciprocity and convey a sense of fairness."[7]

Overall, we view our primary contribution as documenting a stylized fact that may discipline the bargaining literature in developing more reasonable models. Theoretical bargaining models in economics make predictions about a range of things, such as the split of surplus, delay, and the sequence of offers. Until recently, there have been few (or no) real-world datasets containing back-and-forth offers with which a researcher could assess the most important stylized facts regarding the path of play. Our paper documents one particular set of patterns and demonstrates that they are ecologically robust, in the sense that they emerge in many, very different situations. Given this breadth of applicability, and given that canonical equilibrium models do not predict these patterns, we suggest that it would be wise for bargaining theory to attempt to account for them. Likewise, most existing structural estimation analyses of bargaining models omit systematic behavioral patterns such as preferences for splitting the difference. This is the case of workhorse empirical models that rely on various forms of Nash bargaining (e.g., Crawford and Yurukoglu, 2012, and subsequent papers), as well as of the majority of studies that bring to the data models of incomplete information bargaining (e.g., Keniston, 2011; Silveira, 2017; Larsen, 2021). By accounting for the split-the-difference behavior that we document, future iterations of these analyses could yield more realistic estimates and increase the accuracy and reliability of policy takeaways.

Our study also relates to recent work by Camerer et al. (2019) and Huang et al. (2020), which examine incomplete-information bargaining in an experimental setting. We elaborate in Section 5.3 how our results shed light on their findings, offering insight as to how, even in incomplete information settings, agents' behavior is consistent with a notion of

[6] Thompson (2020)'s negotiation textbook captures the tension between even-split offers and an even division of surplus: "even-split between whatever two offers are currently on the negotiation table ... has an appealing, almost altruistic flavor to it. So what is the problem with even splits? ... the pattern of offers up until that point was not 'equal."'

[7] https://www.forbes.com/sites/kwamechristian/2023/03/26/splitting-the-difference-in-negotiation-a-double-edged-sword/?sh=30ecfa152db5. Like the academic literature, popular advice differs when it comes to splitting the difference. Christian writes, "Although some may say never split the difference, there may be situations where the maneuver has validity."

fairness. Finally, our analysis contributes to a literature that studies behavioral phenomena in bargaining using real-world data, such as Pope et al. (2015), studying focal points, or Jiang (2022), studying left-digit bias.

The rest of the paper is organized as follows. Section 2 describes each of our data settings, and Sections 3–4 contain our empirical results documenting the prevalence of split-the-difference offers and related patterns. Section 5 explores several theoretical results and Section 6 concludes.

2 Description of Field Settings

We now introduce the field settings from which we obtain our bargaining data. A benefit of the question we study in this paper—how the current offer in a sequential bargaining game relates to the two most recent offers—is that we can address the question in each of these field settings even though the products or outcomes over which agents negotiate differ drastically from setting to setting. In particular, in each setting, we observe data on many bargaining *sequences* (which we also refer to as *threads*) and, for each sequence, we observe the full set of sequential offers between negotiating parties. We consider an *observation* to be an offer triple: the current offer and two preceding offers.

We drop from every dataset any bargaining sequences in which there are fewer than three offers. We drop any threads that continue beyond the point where an agent makes an offer exactly equal to the opponent's previous offer (which logically should have ended the game in agreement). We also drop any sequences in which an offer lies outside the two most recent offers or cases in which a seller's offer lies below a buyer's offer.[8] We describe additional cleaning steps in Appendix B.

2.1 Business-to-Business Used-Car Bargaining

The first dataset comes from the U.S. wholesale used-car industry. In this market, owners of used-car dealerships buy vehicles from other dealerships as well as from large companies, such as Hertz (rental cars), Wheels (a fleet company), Bank of America (selling off-lease

[8]Such behavior likely corresponds to misrecorded data or to cases where some feature of the bargaining environment changes prior to the current proposed offer, such as the arrival of additional information or a new outside option for an agent. Dropping such threads eliminates fewer than 2% of observations in most data settings, but as many as 24% in some data settings where the arrival of new information or misrecorded offers may be more prevalent. We discuss this in Appendix B.

or repossessed vehicles), or Ford (selling lease buy-back cars). All negotiating agents are professionals or businesses experienced in these negotiations. This $80 billion industry underlies the supply side of the used-car market in the U.S., trading 15 million cars annually; similar platforms exist internationally. For each car, an auction house runs a secret-reserve-price ascending auction, followed by bilateral bargaining between the seller and the highest bidder if the auction price falls short of the reserve price. The data we use in our analysis is generated during this post-auction bargaining stage.

The dataset comes from six auction houses from January 2007 to March 2010. It records each distinct attempt to sell the vehicle through the mechanism, and, for each attempt, every alternating offer proposed by either the buyer or the seller, as well as the outcome of the bargaining. This data overlaps in part with the data used in Larsen (2021), but it also contains additional bargaining sequences that were dropped in that analysis. This new inclusion highlights a major benefit of our empirical approach: the structural exercise of Larsen (2021) required careful data cleaning and controlling for heterogeneity in the items over which the parties negotiated, whereas our approach only requires looking at split-the-difference patterns between offers in a given bargaining sequence, regardless of heterogeneity across items.

Descriptive statistics for this dataset are shown in the first column of Table 1. This dataset consists of 21,734 total bargaining sequences and 33,356 observations (offer triples). Bargaining in this market begins if the auction price is below the reserve price, in which case the auction price becomes the first offer in an alternating-offer bargaining game. The seller can choose to accept the auction price, propose a counteroffer, or quit (ending the game). The bargaining process is typically wrapped up within a day, with an average of several hours between each offer. The average first offer (auction price) is $7,444, followed by an average counteroffer from the seller of $8,918. The average number of offers in a sequence in the used-car sample is 3.53 and the average last price offered is $8,072.[9] The negotiation ends in agreement 59% of the time, at an average accepted price of $7,987. The γ_t objects reported in Table 1 are defined and discussed in Section 3.

[9] As highlighted at the beginning of Section 2, for every data setting, our analysis conditions on bargaining sequences that include at least three offers.

Table 1: Descriptive Statistics of Bargaining Settings

	Cars	Settlement	TV Show	Rides	Housing	Trade	eBay
# Threads	21,734	74,356	204	2,058	176	44,048	6,976,776
# Offer Triples	33,356	208,463	714	2,986	176	46,985	9,789,903
Rounds	3.53	4.80	5.50	4.32	3.00	3.07	3.40
Pr(Agree)	0.59	0.94	0.91	0.39	0.73	0.08	0.29
First Offer	$7,444	$36,391	€19.50	₹51.69	$460,849	0.00	$151.06
Second Offer	$8,918	$64,042	€124.27	₹36.73	$430,071	69.60	$88.71
Final Offer	$8,072	$24,728	€56.67	₹39.50	$451,248	38.30	$122.70
Accept Price	$7,897	$21,838	€55.37	₹42.54	$469,658	24.29	$91.40
γ_3	0.39	0.29	0.32	0.28	0.71	0.44	0.42
γ_4	0.18	0.43	0.31	0.24		0.26	0.38
γ_5	0.38	0.29	0.25	0.18		0.33	0.23
γ_6	0.14	0.39	0.28	0.16			0.31
γ_7	0.33	0.30	0.19	0.17			0.19
γ_8	0.12	0.38	0.34	0.27			
γ_9	0.45	0.32	0.09				
γ_{10}	0.00	0.38	0.38				

Notes: Table shows the number of sequences/threads and number of offer triples in each data setting, as well as averages for several variables in the data. For $t \geq 3$, the variable γ_t denotes the average concession weight in bargaining round t, as defined in Section 3. The units for the average first, second, final, and accept prices are euros for the TV Show and Indian rupees for the Rides setting. The trade setting combines offer sequences in which the units are percentages and sequences where the units are a currency. Units are unimportant for our analysis, which relies on the unit-less γ_t weights.

2.2 Pre-trial Settlement Bargaining from Insurance Claims

The second dataset comes from pre-trial settlement bargaining over injury claims made under U.S. auto and general liability insurance policies.[10] This dataset consists of extensive proprietary information about all claims made to a large national auto and general liability insurer that closed between January 1, 2004, and March 31, 2009. The data contains details about the underlying accident, the alleged injury, the involved parties, the insurance contract, and all attempts by the parties to resolve the associated dispute.

In these insurance cases, an injured person alleges that an injurer caused harm covered by the insurer's policies. If a claimant asserts damages within policy limits or declines to pursue the insured individually for any excess—or if the insurer agrees to cover any damages in excess of the insurance contract's policy limit, which effectively it must do to act in "good faith" if it turns down a demand by the claimant for the policy limit or less—the insurer effectively replaces the injurer as the claimant's counterparty in any dispute, which occurs

[10] Prescott et al. (2014) study a small subset of this information, but the data we use in this paper—particularly the bargaining threads—remains largely unexplored.

in virtually all cases. Under these circumstances, the claimant and the insurer bargain over the amount that should be paid by the insurer to the claimant. This process can take many months to reach a conclusion. If the two parties cannot reach an agreement, the claimant will often pursue litigation, with the claimant (now plaintiff) filing a complaint in court demanding damages against the insurer (now defendant). The parties may also negotiate and settle the claim during the litigation stage. Importantly, the insurer records the amount of each back-and-forth proposal made by either the insurer or the claimant and whether the parties reach an agreement in these negotiations.

After cleaning, our sample contains 74,356 bargaining threads and 208,463 offer triples. Table 1 shows that the average first offer is \$36,391, followed by a counteroffer of \$64,042 from the opposing party, and an average final offer of \$24,728. The negotiations include 4.8 offers on average and end in agreement 94% of the time, at an average accepted price of \$21,838. This number is not between the average first and second offers because of a feature of negotiations in this setting: the first offer may come from the claimant or the insurer, and, in computing the average first and second offers, Table 1 pools over these two cases.[11] Appendix B discusses this point in more detail and describes how we clean the data to form alternating-offer sequences.

2.3 Street Bargaining from a TV Game Show in Spain

The third dataset comes from a TV game show in Spain titled *Negocia Como Puedas* (roughly, "Bargain However You Can") analyzed in Hernandez-Arenaz and Iriberri (2018). This data was generated in the streets of several major Spanish cities in summer 2013. In a typical episode of the show, the host approaches individuals in the street and invites them to participate in the game. Upon acceptance, an individual (the *proposer*) is endowed with a potential pie of 100 euros and is asked an easy question. The proposer is not allowed to answer the question herself. Instead, she must, within a three-minute limit, (i) find a passer-by (the *responder*) able to provide an answer to the question that the proposer finds satisfactory, and (ii) negotiate a price that the proposer will pay the responder to be able to use that answer.

If the negotiations succeed and the responder's answer to the original question is correct

[11] Examining these two cases separately, the average accepted price is indeed between the average first and second offers.

(as determined by the host), the proposer pays the responder the agreed amount. The proposer then moves on to a new question, referred to as a new stage of the game, where the process is repeated. In the new stage, the size of the pie increases (by 200 euros in the second stage, 300 in the third, and 1,000 in the fourth). In any stage of the game, if the proposer does not reach an agreement within the three-minute time limit, the game ends, and the proposer gets nothing. Throughout the game, the size of the pie is only known to the proposer, not the responder. In any given stage of the game, the bargaining is unstructured, but the negotiations typically follow an alternating-offer structure, with the proposer making the first offer. We only keep those threads in which offers clearly alternate between parties. Additional details on the cleaning of the data are found in Appendix B.

In the data, we have 204 sequences and 714 offer triples. Table 1 shows that the average first offer is 19.5 euros, followed by an average second offer of 124.3 euros and an average final offer of 56.7 euros. The game ends in agreement most of the time (91%), after an average of 5.5 rounds and at an average accepted price of 55.4 euros.

2.4 Auto Rickshaw Rides Bargaining in India

The fourth dataset comes from the local transportation market by auto rickshaw in Jaipur, India. An auto rickshaw is a form of three-wheeled mini-taxi, officially capable of carrying three passengers in a semi-enclosed back seat. Auto rickshaws are the primary means of rented transportation in Jaipur. During the period in which Keniston (2011) collected the data (January 2008 to January 2009), all prices were set by negotiation.

The data was collected by surveyors (whom we also refer to as buyers) who followed one of two possible protocols. In *real* bargaining, buyers were assigned to travel by auto rickshaw along fixed routes through the city, bargaining for the price of each ride. At the beginning of each route, buyers were paid a lump sum slightly higher than the expected cost of the route and were allowed to keep any money not spent on auto rickshaw fares. At the end of their assigned trip, they were free to return to their homes or alternate employment. Thus, their financial incentives and cost of time were similar to real trips taken for personal purposes. In *scripted* bargaining, buyers negotiated with sellers according to a written bargaining script (prepared by Keniston) consisting of a sequence of pre-determined counteroffers. Scripted surveyors were instructed to act as if they were bargaining in a realistic manner so that drivers (whom we also refer to as sellers) would respond as naturally as possible. After the conclusion of the bargaining, surveyors wrote down the series of offers made by the drivers and themselves. Drivers were not aware that they were part of a field experiment. The average negotiation took 55 seconds to complete.

In our analysis, we exclude any offers or responses that come from scripted surveyors, as they do not represent actual reactions. Additional details on data cleaning are in Appendix B. Our main sample consists of 2,058 bargaining threads and 2,986 offer triples. Table 1 demonstrates that the average first offer is 52 rupees, followed by an average counteroffer of 37 rupees and a final offer of 40 rupees. Negotiation concludes after an average of 4.32 offers, ending in agreement 39% of the time at an average accepted price of 43 rupees.

2.5 Bargaining in Residential Real Estate

The fifth dataset we use is new to the literature and comes from a growing residential real estate brokerage company that offers discounted agency fees of 2–3% rather than the traditional 6%. We collected this dataset in collaboration with the company, covering a number of houses on the market from 2015 to 2019 in Colorado. This dataset differs from the others we analyze in that we only observe offers placed by potential buyers, not by the seller. The company informs us that seller counteroffers are indeed quite rare in this market. Rather, the typical negotiation proceeds with a seller list price (which we treat as the first offer) followed by sequential offers from the buyer. For a given home, we observe the seller's list price and each offer placed by potential buyers. The time on the market for a given home can be several weeks or several months. In this setting, we consider split-the-difference behavior to be cases in which a buyer makes an offer and then, if that offer is rejected, subsequently makes an offer that splits the difference between the list price and the buyer's initial offer.

Table 1 shows that our sample has 176 threads and the same number of offer triples. This feature is by construction: our main object introduced in Section 3 cannot be defined in cases where one party makes three consecutive offers; thus, we only keep the seller's list price and the first two offers from the buyer, i.e., one offer triple for each bargaining sequence. These bargaining sequences end in agreement 73% of the time. Bargaining begins with an average list price of \$460,849, followed by an average second offer of \$430,071, and a final offer of \$450,248. When parties agree, they end at an average price of \$469,658. Additional details on data cleaning are found in Appendix B.

2.6 International Trade Tariff Bargaining

The sixth dataset contains detailed information on international trade negotiations recently declassified by the World Trade Organization (WTO). In these negotiations, countries bargain over commitments on their respective import tariffs. Despite the multilateral nature of both the WTO and its predecessor, the General Agreement on Tariffs and Trade (GATT), the negotiations are mostly bilateral, with individual country pairs making *requests* and *offers* over the tariff for a specific tariff-line (a product code).

We use the dataset of Bagwell et al. (2020), which comprises the Torquay Round (1950–1951). This data includes 298 bilateral bargaining pairs from 37 countries, negotiating tariffs over thousands of tariff-line products. A bargaining thread is defined as two countries (*proposer* and *target*) negotiating a tariff over a tariff-line product. In each thread, we observe all requests from the proposer and offers from the target. As documented in Bagwell et al. (2020), relatively few back-and-forth offers and counteroffers take place in any given thread. For our analysis below, we consider whether a proposer requests a tariff that splits the difference between a zero tariff and the existing tariff (the status quo before the negotiations). To map the possibilities of a zero tariff and existing tariff into the same framework as the other data settings, Table 1 considers a zero tariff as the de facto initial request from the proposer (and, indeed, a zero tariff is frequently an actual request made in this setting). Similarly, Table 1 considers the existing tariff as the initial offer from the target (the row of "second offer"). We discuss this point further in Section D. All subsequent requests and offers are also recorded in the data.

Table 1 shows that the sample consists of 44,048 bargaining sequences and 46,985 offer triple observations. The game proceeds for 3.07 offers on average, ending in agreement 8% of the time. Note that here we combine sequences in which the units are percentages and sequences in which the units are a currency; thus, we provide no units on these averages of offers in Table 1 and their values are hard to interpret. However, units are unimportant for our analysis, which relies on the unit-less γ_t weights (described in Section 3). Additional details on data cleaning are found in Appendix B.

2.7 Bargaining on eBay's Best Offer Platform

Our final field setting comes from eBay's Best Offer negotiation platform. eBay is well known as a platform for buying and selling via auctions or fixed prices. Less well known is the bargaining mechanism on eBay, through which a buyer and seller negotiate via alternating offers (limited to three offers by each party). The game begins with the seller posting a list price. An interested buyer can pay this price or make an offer. Offers are sent through the eBay platform, and the receiving party has 48 hours to respond by either accepting, declining, or proposing a counteroffer. Our data comes from internal data collected by Larsen and coauthors for a separate project (Backus et al. 2020), and it consists of all bargaining sequences placed by buyers and sellers on eBay (regardless of the product) from June 2012 through May 2013.

The sample consists of 6,976,776 bargaining sequences, comprising 9,789,903 offer triple observations. These sequences contain an average of 3.4 offers. The average list price is \$151, the average second offer is \$89, and the average final offer is \$123. When trade occurs (which happens 29% of the time), the final agreed-upon price is \$91.

3 Split-the-Difference Offers are Widespread

We now demonstrate a key empirical pattern: agents favor offers that split the difference between the two most recent offers. To show this, we first introduce some useful notations. For each data setting, we organize each sequential bargaining thread in the following way, as in Backus et al. (2020). For each round $t = 1, 2, ...$ in the bargaining thread j, we observe the proposed amount, $p_{j,t}$. This proposed amount is an offer made by the seller/buyer, insurer/claimant, proposer/respondent, driver/surveyor, proposer/target, etc. If $p_{j,t}$ comes from one player, $p_{j,t+1}$ must come from the opponent—with the exception of threads in the housing dataset, as we note above and explain further later in this section.

We can write the proposed amount in round $t \geq 3$ as a weighted average of the proposed amount in the previous two rounds: $p_{j,t} = \gamma_{j,t} p_{j,t-1} + (1 - \gamma_{j,t}) p_{j,t-2}$. Therefore, $\gamma_{j,t}$ represents, for bargaining thread j, the weight that the player in round t places on the opponent's previous offer, and $1 - \gamma_{j,t}$ represents the weight the player places on her own previous offer. We can think of $\gamma_{j,t}$ as how much a player *concedes* to her opponent when she is making a counteroffer, or her *concession weight*. Rearranging to solve for $\gamma_{j,t}$ yields

$$\gamma_{j,t} = \frac{p_{j,t} - p_{j,t-2}}{p_{j,t-1} - p_{j,t-2}}. \tag{1}$$

In particular, a split-the-difference offer corresponds to $\gamma_{j,t} = 0.5$.

As highlighted in Section 2.5, in the housing dataset, we observe a seller's list price and a given buyer's offers, and in no bargaining thread does the number of offers from the same buyer exceed two. In this setting, we only define the concession rate for the second buyer's offer. Specifically, we let

$$\gamma_{j,3} = \frac{p_{j,3} - p_{j,2}}{p_{j,1} - p_{j,2}}, \tag{2}$$

where $p_{j,1}$ is the list price, $p_{j,2}$ is the buyer's first offer, and $p_{j,3}$ is the buyer's second offer. Thus, in the housing dataset, $\gamma_{j,3}$ represents the weight a buyer places on the list price, and $1 - \gamma_{j,3}$ represents the weight she places on her own initial offer.

The advantage of focusing on concession weights is that they are unit-less and do not require considering any heterogeneity across negotiation threads. Table 1 shows the average $\gamma_{j,t}$ for each round of the game, beginning with $t = 3$. We observe as many as ten $\gamma_{j,t}$'s in some settings,[12] and, as highlighted above, only the round-3 $\gamma_{j,t}$ in the housing setting. The average $\gamma_{j,t}$ in each round is below 0.5 (other than for housing, where it is 0.71), suggesting that most agents make offers closer to their own previous offers than to their opponent's. In the settlement data, the average $\gamma_{j,t}$ is roughly 0.3–0.4 in every round of the game, whereas in the cars setting, $\gamma_{j,t}$ tends to be much smaller in even rounds, which correspond to the seller's turns, suggesting that sellers generally concede less than buyers in this context.

The main pattern we wish to examine is whether agents in each of these settings exhibit a tendency to make offers that lie halfway between the two most recent offers on the table (i.e., $\gamma_{j,t}$ close to 0.5). Figure 1 plots a histogram of these concession weights, with each panel corresponding to one data setting. We pool together $\gamma_{j,t}$ for all $t \geq 3$; our results are similar if we analyze $\gamma_{j,t}$ from each round t separately.[13] An interesting pattern stands out in all datasets: there is a mass point at 0.5—counteroffers that are halfway between the

[12] Some settings have bargaining threads exceeding ten rounds, and we truncate them to the first ten rounds in Table 1.

[13] For the eBay setting, Figure 1.G is not new to the literature, capturing a result that is illustrated in several panels in Figure 8 of Backus et al. (2020), but here the result is pooled across rounds. Similarly, our Figure A4.D and column 6 of Table 2 capture the same information as Figure 9 and Table 8 of Backus et al. (2020). All other results in this study are new.

previous two offers, or split-the-difference offers.[14]

There are a number of differences across our data settings. The trade tariff setting, for example, likely involves agents bargaining over multiple issues simultaneously, unlike, say, the rides setting. In practice, agents in a multi-issue setting might be expected to fully concede on one issue in return for the other party's concession on a separate issue. These possibility makes it all the more striking that split-the-difference offers within a given thread appear even in the trade setting.

The various settings certainly differ in the relative prevalence of split-the-difference behavior. Figure 1 shows that housing and settlement appear to have lower frequencies of split offers relative to non-split offers.[15] We note that both settings often involve principals and agents on each side, with the agent formally bargaining with the counterparty's agent while the principals, in theory, make decisions about what to offer and when to accept on the basis of communications from their agents. This insulated arrangement may dampen some effects of principals' preferences, such as a preference for fairness. Also, parties bargaining over the resolution of an insurance claim may anticipate bargaining over more rounds on average. Early offers and counter-offers may serve some other purpose, whereas split-the-difference offers may represent a genuine attempt to resolve the dispute without (or with very few) further rounds. There are other candidate explanations for this variation as well. In the housing setting, for instance, the split-the-difference bargaining may be less likely when sellers have a particularly strong position in the market (with few sellers and many buyers), such that a split-the-difference offer corresponds to too much concession from the seller's standpoint.

Despite the drastic differences in the environments, products, outcomes, or agents, the presence of a split-the-difference pattern is clear in all seven settings. To our knowledge, the widespread nature of this split-the-difference pattern in negotiations in the field has not previously been documented. Unlike laboratory experiments, where prior research shows that players split a *commonly known* pie, in our settings, no player knows the true valuation of her opponent. A challenge is then to explain why, in real-world negotiations with private information, players propose offers that split the difference between endogenous *offers*.

[14]Another common mass point in these histograms is at zero, representing cases where a player does not budge at all. In the housing dataset, a mass point at 1 is also common, representing that an agent fully concedes to the seller's list price.

[15]This can also be seen in the "split rate" row in Table 2, discussed in the following section.

Figure 1: Distribution of Concession Weights $\gamma_{j,t}$

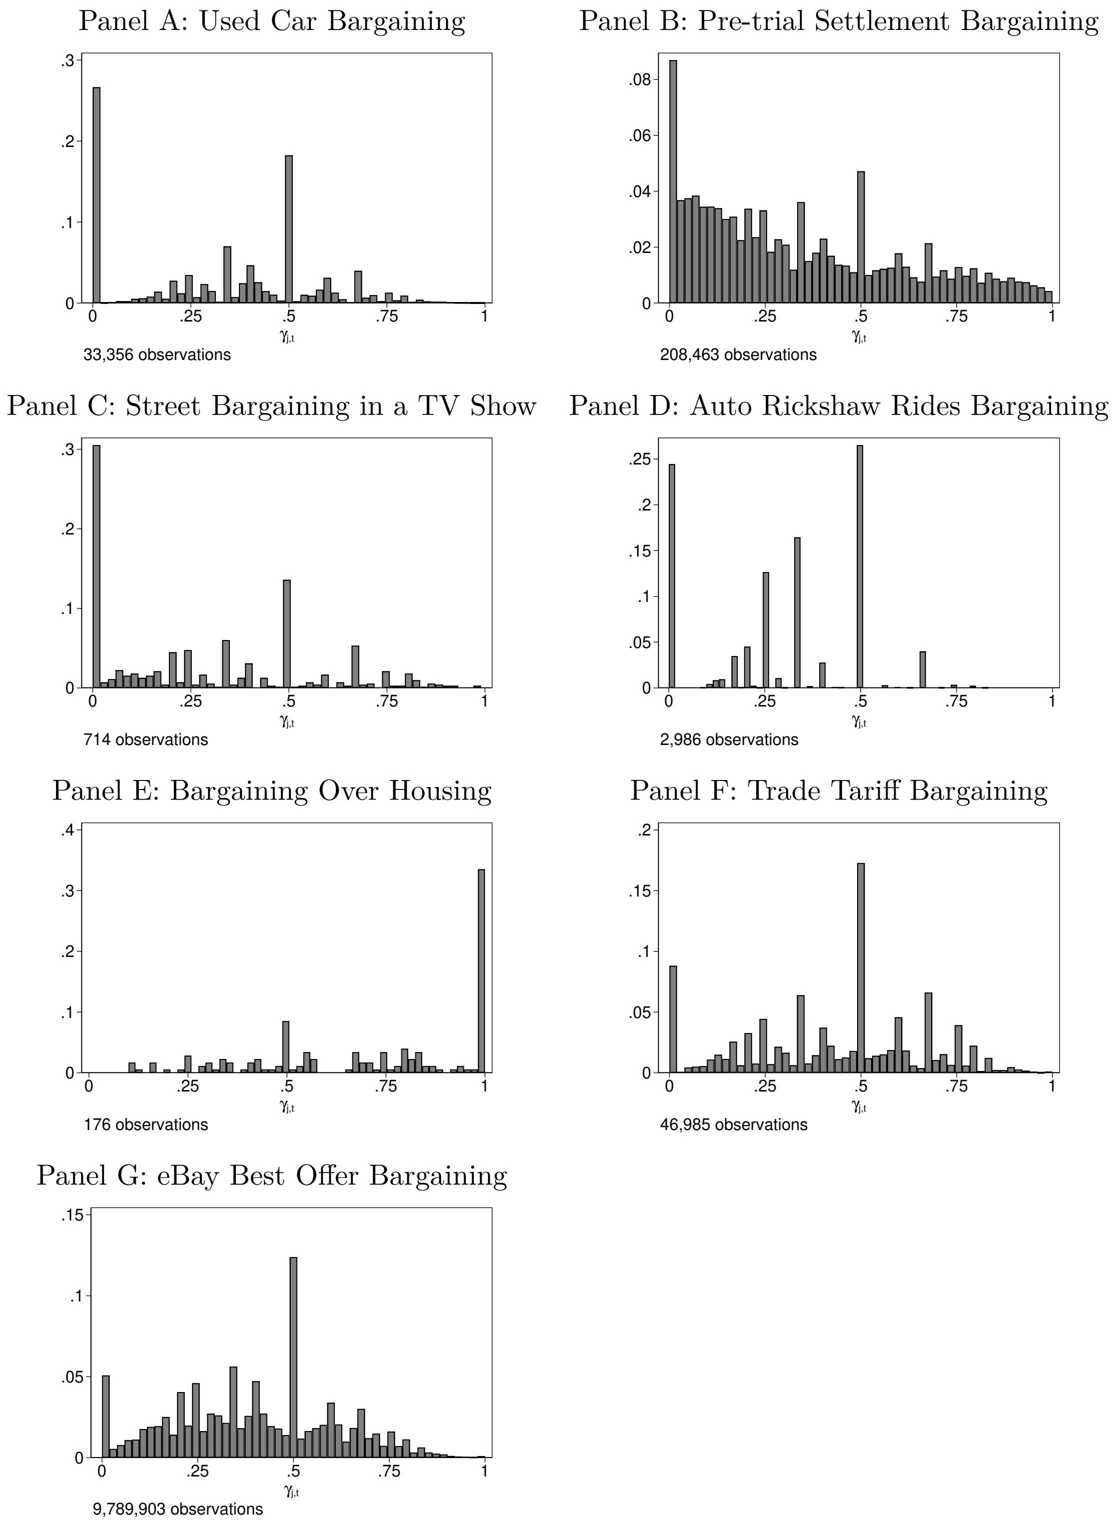

Notes: Each panel shows a histogram of $\gamma_{j,t}$ in a given data setting, as defined in Section 3.

4 Split-the-Difference Offers as a Norm

We now present evidence that the split-the-difference pattern that we document in our empirical settings constitutes a social norm. Following Fehr and Gächter (2000), we define a social norm as a behavioral pattern that (i) relies on a socially shared belief about what constitutes appropriate behavior; and (ii) triggers the enforcement of the prescribed behavior by informal social sanctions. This characterization of split-the-difference as a social norm is echoed by veteran business and hostage negotiator Chris Voss, who writes, "The traditional negotiating logic that's drilled into us from an early age, the kind that exalts compromise, says 'Let's just split the difference...Then everybody's happy"' (Voss and Tahl, 2016, p.115). Below we show results suggesting that agents that make split-the-difference offers are systematically rewarded by the opposing party (or, conversely, that the failure to make a split-the-difference offer does not yield this reward), confirming the view that these offers, as part of a social norm, are enforced by informal rewards. These rewards come in the form of higher probabilities of acceptance and higher probability of subsequent split-the-difference offers within the same bargaining sequence. Buyers who do not comply with the norm experience a higher probability of their opponent exiting the negotiations, which could impose financial or time costs, depending on the setting.

We also show that the two most recent offers are special: offers that split the difference between other anchor points as less prevalent. These other anchor points include earlier offers in the game other than the two most recent offers or privately known reference points (observable only to one bargaining party and not the other), which we observe in some of our settings. This latter results suggests that agents gravitate to a behavior that is easily verifiable by all parties, consistent with the idea that split-the-difference offers are subject to social enforcement. Additionally, we show that an agent's offer is more likely to split the difference if either the previous offer by the opponent or that agent's own previous offer was a split offer. Split-the-difference offers are also more likely to occur if the opponent's previous offer was arguably reasonable. For example, a seller is more likely to propose a split-the-difference offer at the third round of the game if the buyer's offer at round 2 was not well below the seller's initial offer (i.e., the buyer was not low-balling the seller).

4.1 Split-the-Difference Offers Are More Likely to be Accepted, Less Likely to Cause Exit

In this section, we explore how agents *respond* when they receive split-the-difference offers compared to when they do not. We show that split offers are discontinuously more likely to be accepted than non-split offers, suggesting that splitting the difference is not merely a heuristic used by the *offering* party, but rather that it is viewed as preferable (arguably fair) behavior by both the offerer and the receiver. Consistent with this interpretation, after agents make split offers, their opponents are less likely to abandon negotiations. This suggests that the choice to exit a bargaining interaction may be motivated not just by the lack of expected surplus, but also by the perception that the opponent is not being "fair."

Specifically, we examine how a player's choice of offer, as measured by the concession weight, $\gamma_{j,t}$, relates to the probability that the offer is accepted or leads to a breakdown of negotiations. We create a measure for whether the offer is a "split" offer by creating an indicator $Split_{j,t}$ that is equal to one if $\gamma_{j,t}$ is equal to 0.5 (after being rounded to the nearest hundredth, or $\gamma_{j,t} \in [0.495, 0.505]$) for each $t \geq 3$. We then estimate the following linear probability regressions:

$$Accept_{j,t} = \beta Split_{j,t} + f(\gamma_{j,t}) + \tau_t + \epsilon_{j,t} \tag{3}$$

$$Exit_{j,t} = \beta Split_{j,t} + f(\gamma_{j,t}) + \tau_t + \epsilon_{j,t}, \tag{4}$$

where $Accept_{j,t}$ is an indicator for whether the offer is accepted, $Exit_{j,t}$ is an indicator that the opponent exits before period $t+1$, τ_t is a round fixed effect, and $f(\gamma_{j,t})$ is a flexible function of $\gamma_{j,t}$. We specify $f(\gamma_{j,t})$ as a third-order polynomial of $\gamma_{j,t}$.[16] The results are reported in Table 2, where each column corresponds to one dataset. We also report the frequency of acceptance, exit, and split offers. The acceptance rate varies across settings from 7% to 73%, the exit rates are between 3% and 86%, and the fraction of split offers ranges from 4% to 19%.

In Panel A, we see a positive coefficient before the "split" offer indicator in all of our datasets, and this is statistically significant in all columns except the cases of auto rickshaw rides and housing (columns 4 and 5). This means that an offer in bargaining is more likely to

[16] The results are not sensitive to this choice; we find similar results with second-, fourth-, or fifth-order polynomial approximations. We also find similar results defining "split" offers using other bandwidths, including 0.01 (i.e., $\gamma_{j,t} \in [0.49, 0.51]$) and 0.05 (i.e., $\gamma_{j,t} \in [0.45, 0.55]$).

Table 2: Acceptance and Exit Following Split Offers

	(1) Cars	(2) Settlement	(3) TV Show	(4) Rides	(5) Housing	(6) Trade	(7) eBay
Panel A: Acceptance							
Split	0.120***	0.219***	0.151**	0.0569	0.158	0.0676***	0.0849***
	(0.00830)	(0.00555)	(0.0615)	(0.0354)	(0.119)	(0.00341)	(0.000484)
Panel B: Opponent Exit							
Split	-0.0402***		-0.00977	-0.103***		-0.0511***	-0.0198***
	(0.00577)		(0.0122)	(0.0392)		(0.00453)	(0.000504)
N	33356	208463	714	3010	176	46985	9789903
Order of $\gamma_{j,t}$	3	3	3	3	3	3	3
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Accept rate	0.38	0.34	0.26	0.20	0.73	0.07	0.20
Exit rate	0.27		0.03	0.41		0.86	0.51
Split rate	0.18	0.04	0.14	0.19	0.08	0.17	0.12
R^2	0.146	0.462	0.0295	0.160	0.0336	0.110	0.113

Notes: Table shows the estimated coefficient on the split indicator from the regressions described by equations (3) and (4). Each column corresponds to a separate data setting. The accept and exit rates are the means of the dependent variables, and the split rate is the mean of the split indicator. Standard errors are shown in parentheses. The number of observations for the rides setting differs from that in Table 1 because some scripted bargaining acceptance must be dropped and some scripted bargaining offers can be included here. The number of observations for the settlement setting differs from that in Table 1 because some sequences end at round 10, and we have no data on whether round 10 offers were accepted. See Appendix B for details. $*: p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$

be accepted if it is a split offer than if it is not. This effect is surprisingly large in magnitude, varying from 5.7% to 21.9%. Conversely, in Panel B, "split" offers are negatively associated with the opponent exiting the bargaining interaction. While sizes of the effect on exit are usually smaller than those on acceptance, they are significant in all cases except for the TV show, where exit is rare. The structure of the data on legal settlements and housing (columns 2 and 5) do not allow us to examine exit as an outcome in these contexts.

Our result in the housing setting is likely insignificant due to the small sample size (176 offer triples). One possible explanation for the lack of a significant effect on acceptance in the setting of auto rickshaw rides is, this specification compares the acceptance rate of a split offer with nearby offers and, as shown Figure 1, the nearby offers in this dataset are quite sparse. This is also reflected in the high rate of split offers in the fourth column (19%). Therefore, we likely lack the power to detect a significant effect in this dataset. This dataset, however, has a unique subset—the scripted bargaining offers that Keniston (2011) randomly assigned—in which we can obtain estimates that are closer to a *causal* effect of split offers on acceptance rates. As highlighted in Section 2.4, buyers in these scripted sequences make offers that are assigned by the experiment designer rather than arising endogenously. When we estimate equation (3) in this subset of the data, shown in columns 4–5 of Appendix Table A2, we find positive point estimates, and a particularly large, positive effect in those sequences that begin with a buyer offer. Appendix Table A3 presents the results of a similar subsample analysis for equation (4).

In Figure 2, we offer an even more flexible approach to this question, plotting a weighted local linear fit of acceptance and $\gamma_{j,t}$, using observations where $\gamma_{j,t}$ is *not* a split offer. We also plot in Figure 2 the average acceptance probability for observations that *are* split offers, along with the 95% confidence bound around this mean. We find that the underlying relationship between the acceptance and $\gamma_{j,t}$ is monotonic in most regions, and split offers are substantially more likely to be accepted than nearby offers with similar $\gamma_{j,t}$ in all datasets except the settings of auto rickshaw rides (where we again lack power locally around 0.5) and housing (where the number of observations is small). In the latter two datasets, the point estimates are still higher at split offers.

The striking implication overall is that, even across these widely varying field settings, split-the-difference offers are more likely to be accepted than even a slightly *more favorable* offer. This suggests that a preference toward splitting the difference between the two most recent offers is a norm followed not only by the proposer of these offers but also by the receiver. These results are reminiscent of findings in a wide range of laboratory ultimatum games (e.g., Roth et al. 1991), which show receivers frequently rejecting offers of less than half of the surplus, but accepting "fair" offers of a 50-50 split of the surplus nearly 100% of the time.

4.2 Split Offers Are Frequently Followed by Split Offers

In this section, we address two empirical questions. First, do split offers in period t by one party tend to be followed by the *opponent* proposing a split offer in period $t+1$? Second, do split offers in period t tend to be followed by split offers by the *same* party in period $t+2$? Both of these points speak to the question of whether splitting the difference between the two most recent offers is a norm that is perhaps followed more consistently by some agents than others and that, when invoked by one agent, tends to be adopted by the opponent.

To examine this question, we analyze the following linear probability model:

$$Split_{j,t} = \beta Split_{j,t-1} + \epsilon_{j,t}, \tag{5}$$

Figure 2: Probability of an Offer Being Accepted

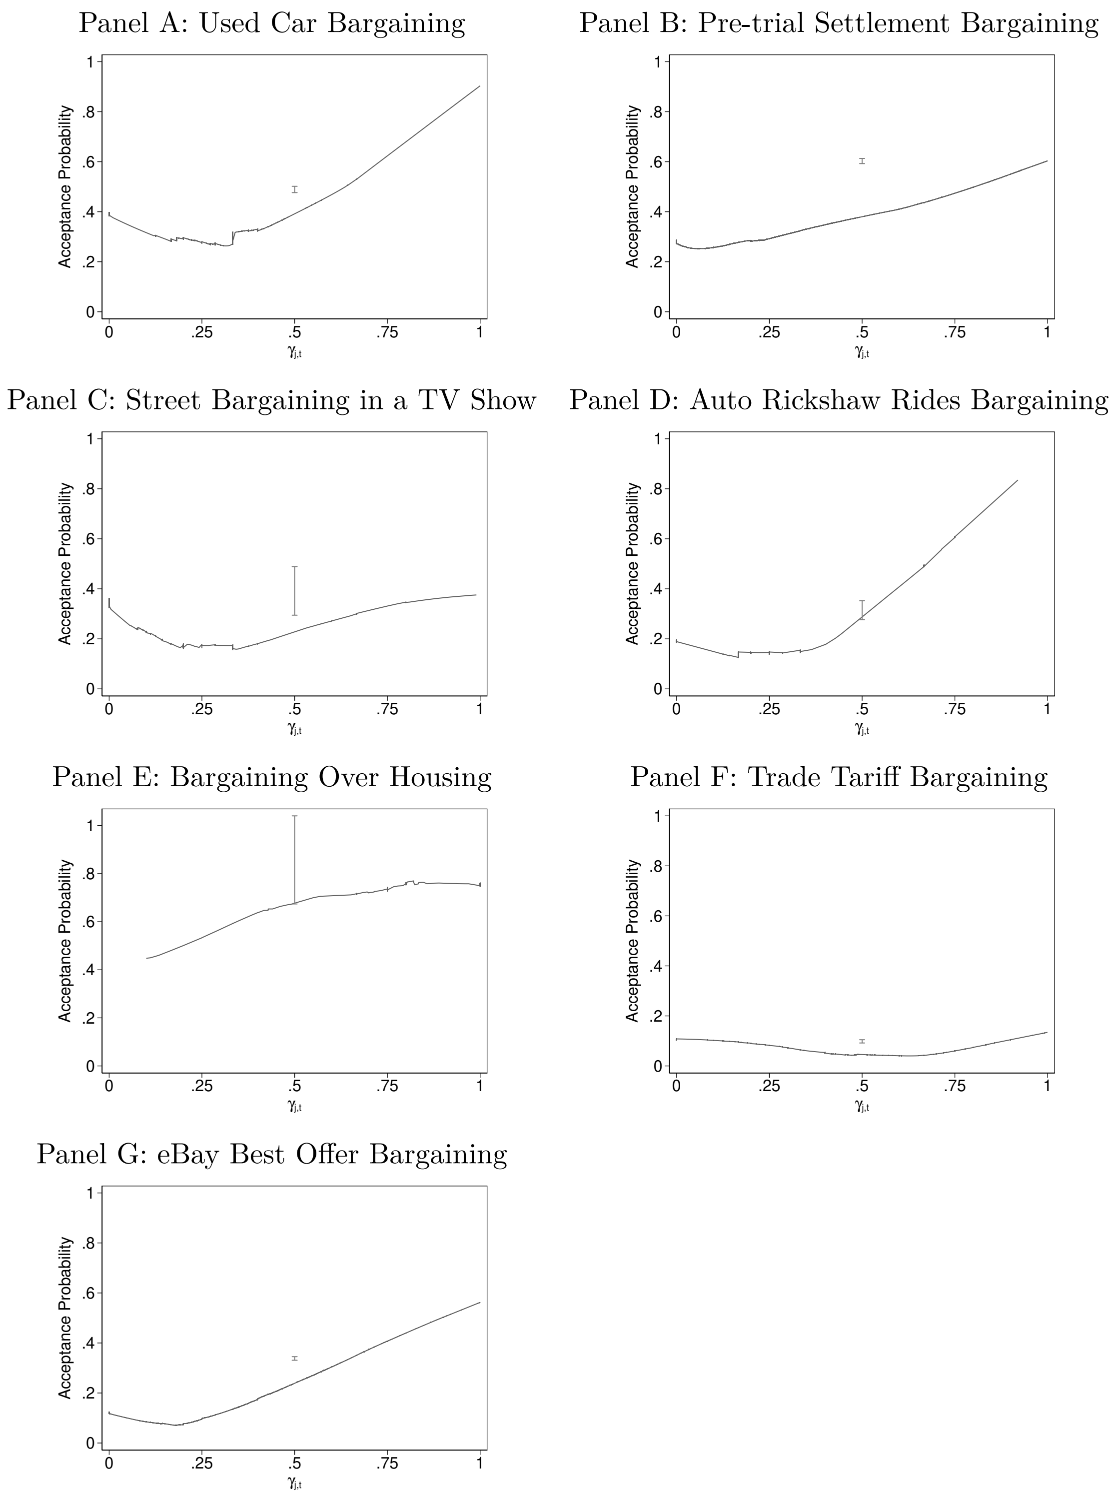

Notes: Each panel shows a local linear fit of the acceptance probability as a function of $\gamma_{j,t}$ along with the acceptance probability of split offers (plus the 95% confidence interval around this point). The fitted values are estimated using locally weighted least squares with a tricube weighting function. To facilitate computation of the local linear estimator in the eBay setting, we use a random sample of 100,000 threads.

where we regress the indicator of whether, in period t, an agent proposes a split offer, $Split_{j,t}$, on the indicator of whether the most recent offer, $Split_{j,t-1}$ (which naturally comes from the opponent), also corresponds to a split offer. Panel A in Table 3 shows the results.[17] We observe positive point estimates in each setting, and these estimates are significant in most columns, suggesting that agents are more likely to propose a split offer when the opponent has just done the same.

In panel B, we consider a version of equation (5) in which we use $t-2$ actions on the right-hand side rather than $t-1$, allowing us to examine whether agents who make split-the-difference offers earlier in the game are more likely to do so again. We find evidence of this effect for the case of settlement and eBay negotiations and a marginally significant effect in the case of auto rickshaw rides.

Table 3: Repeat Split-the-Difference Behavior

	(1) Cars	(2) Settlement	(3) TV Show	(4) Rides	(5) Trade	(6) eBay
	A: Effect of Opponent's Split Offer					
$Split_{j,t-1}$	0.00304	0.144***	0.117**	0.00358	0.0928***	0.125***
	(0.00763)	(0.00672)	(0.0591)	(0.0193)	(0.0188)	(0.000840)
N	11622	134107	510	2714	2937	2813127
	B: Effect of Agent's Own Split Offer					
$Split_{j,t-2}$	0.0440	0.0807***	0.0203	0.0708*	-0.000644	0.0699***
	(0.0276)	(0.00991)	(0.0724)	(0.0382)	(0.0519)	(0.00137)
N	2155	83296	343	1121	250	1100490

Notes: Panel A shows the estimated coefficient on the one-period-lagged split indicator from the regression described by equation (5). Panel B shows results instead using the two-period-lagged split indicator. Each column corresponds to a separate data setting. The housing data setting is omitted because no sequence contains more than three offers. $*: p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$

We examine this latter effect further by exploring whether splitting the difference is a norm followed by certain agents more than by others. To do so, we take advantage of agent identifiers, which we observe in the used car, trade, and eBay bargaining settings.[18] In a given dataset, we sort agents by the fraction of split offers among all offers they make. We then plot, on the horizontal axis in each panel of Figure 3, the cumulative share of offers made by agents. The solid line corresponds to the cumulative share of offers by these agents

[17] We cannot examine this question in the housing dataset as we only observe a single offer triple in each sequence in that setting.

[18] In our other data settings, we have no consistent means of tracking an agent across different bargaining sequences.

that are *split* offers. If the propensity to propose split offers is roughly equal across all agents, and if each agent makes a large number of offers, the solid line should be close to the 45-degree line (the dotted one). This comparison can thus be considered a modified Lorenz curve that measures the "inequality" of split offers among agents.

Figure 3: Some Agents More Likely to Make Split Offers

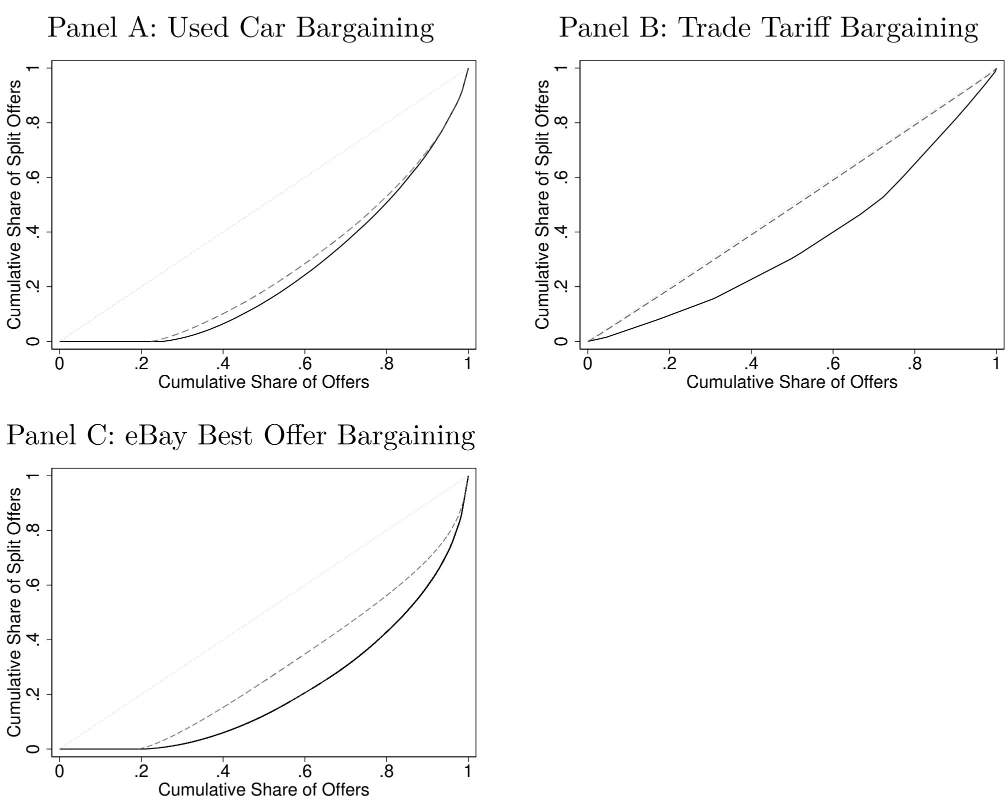

Notes: Each panel ranks agents by the fraction of split offers among all offers they make and plots their cumulative share of total offers on the x-axis and the cumulative share split offers on the y-axis. The solid lines use the real data. The dashed lines use simulated split indicators assuming every agent has the same propensity to propose split offers. The dotted line indicates the 45-degree line.

Figure 3 indeed shows a gap between the solid and dotted lines in each panel. In reality, however, we only observe a few offers made by a given agent, and hence a large part of the area between the solid line and the 45-degree line is due to sampling noise.[19] To construct a more meaningful benchmark, we consider a case where each agent has the same probability to propose a split offer, with this probability given by the split rate reported in Table 2.

[19]For example, suppose every agent has the *same* propensity q to propose a split offer, but each agent only ever proposes one offer in the data. We would observe roughly a fraction of q players proposing one split offer and a fraction of $1-q$ players proposing no split offers, and the curve will be far off the 45-degree line.

Using this split rate, we simulate a fake split indicator for each observation following a Bernoulli distribution. We plot our cumulative share of split offers based on these fake indicators using the dashed line in Figure 3, which should lie between the 45-degree line and the curve plotted using the real data.

Comparing the dashed and solid lines in Figure 3, we find evidence that some agents have a stronger proclivity toward split-the-difference than others. This is particularly the case in the eBay and trade settings, where the dashed line is farther from the solid line. In the used-car setting, the solid and dashed lines are close, indicating that the propensity to make split-the-difference offers is roughly uniform across agents.

4.3 The Two Most Recent Offers Are Special

Our analysis in Section 3 demonstrates robust evidence across a wide spectrum of settings that a modal strategy in real-world bargaining is to make offers that split the difference between the two most recent offers. Here, we explore whether it is indeed the two most recent offers that serve as the most prominent anchor points that players tend to split in half, or whether other anchor points are equally common. For example, it is possible that offers splitting the difference between *earlier offers* in the game (prior to the two most recent offers) are also common. For example, in a sequence $\{100, 50, 90, 70\}$, a subsequent offer splitting the difference between the two most recent offers would be 80, but it may be that 75 (which splits the difference between the *first two* offers) is also a focal point for players in this game.

To examine this possibility, we define *placebo concession* by treating the proposed amount in round t as a convex combination of offers from *earlier* rounds of the game. For example, for $t = 4$, we can treat the proposed amount as a convex combination of offers from the first and second rounds, which we define as $\gamma^3_{j,4} = \dfrac{p_{j,4} - p_{j,1}}{p_{j,2} - p_{j,1}}$. In general, for $t \geq 4$ and $s < t$, the placebo concession is defined as follows:

$$\gamma^s_{j,t} = \frac{p_{j,t} - p_{j,s-2}}{p_{j,s-1} - p_{j,s-2}}, t \geq 4, 3 \leq s < t, \tag{6}$$

where t is the round in which the current offer is evaluated, and $s < t$ is the round in which offers from rounds $s - 1$ and $s - 2$ actually were the two most recent offers. If negotiating agents care most about fairness as defined by an equal split of the two most recent offers, we expect less mass at 0.5 for the placebo concession than in our main results from Section 3.

For this analysis, we focus on the three datasets for which we have the longest sequences, as they allow us to construct the placebo concession metric: used car bargaining, pre-trial settlement bargaining, and eBay bargaining. Figure 4 plots the distribution of true concession and placebo concession in these three datasets. In the left column, we replicate the histograms of true concession weights from Figure 1. In the middle column, we plot histograms of the placebo concession weights based on bargaining offers in earlier rounds, $\gamma_{j,t}^{s}, t \geq 4, 3 \leq s < t$, as defined in equation (6).

Figure 4: Distribution of True Concession and Placebo Concession

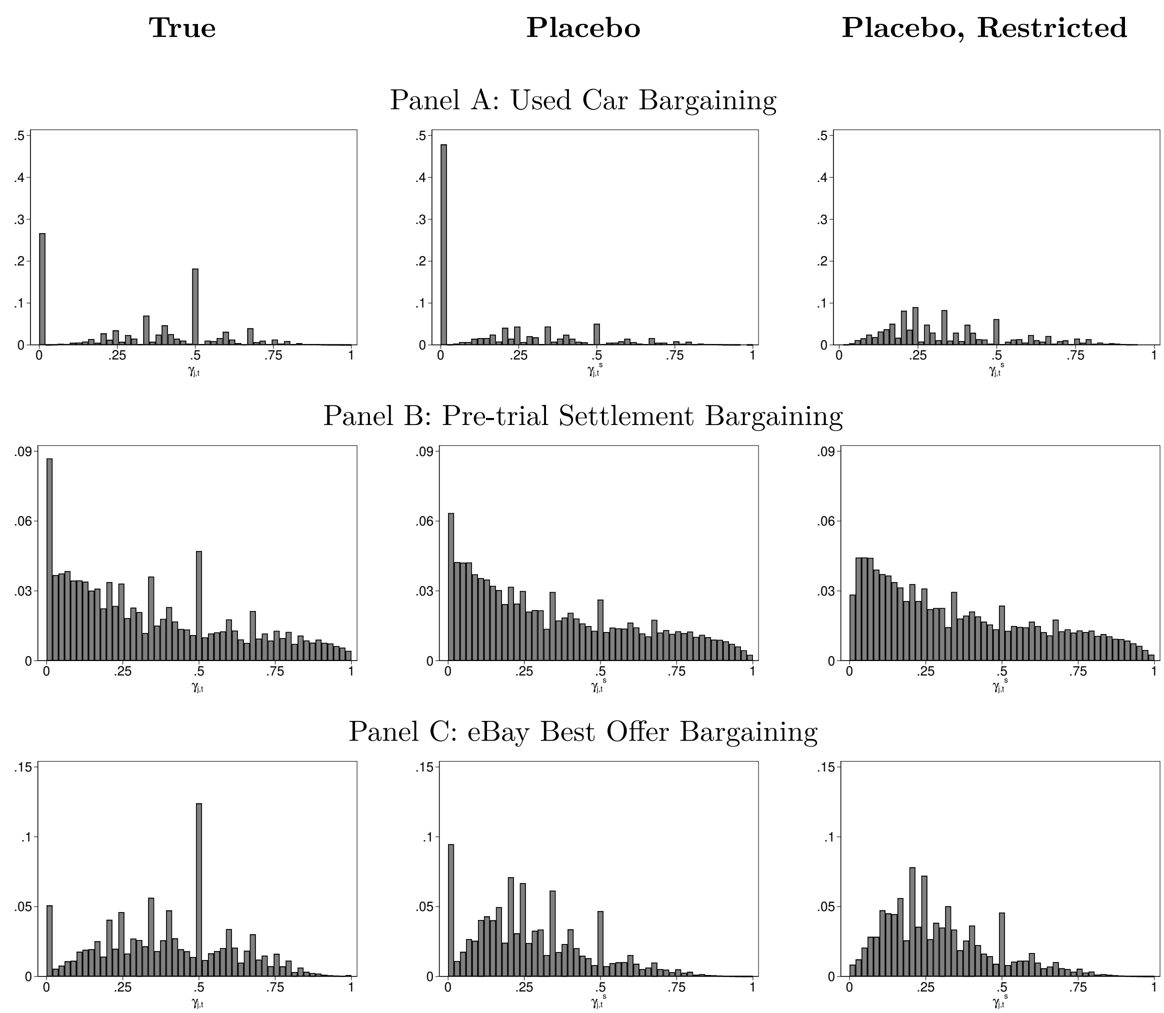

Notes: Figure shows, in left plots, histograms of true concession weights ($\gamma_{j,t}$) as in Figure 1 for the cars (panel A), settlement (panel B), and eBay (panel C) settings. In the middle plots, we show histograms of the placebo concession weights. In the right plots, we show histograms of the placebo concession weights restricting the sample to exclude observations that are mechanically equal to 0.5.

In the third column of Figure 4, we focus on a restricted sample in which we drop cases that can mechanically lead to mass points at 0.5 even in the placebo concession weights. As an example, consider an offer sequence with the first four offers being $\{100, 60, 90, 70\}$. A split-the-difference offer at the fifth round would be 80, but this offer would also represent splitting the difference between the earlier offers of 100 and 60. Our restricted sample excludes placebo concession weights that are exactly equal to the true concession weights for a given round.

Panels A and C of Figure 4 demonstrate that counteroffers occur halfway between *earlier* offers of the game (the middle and right columns) less frequently than between the two most recent offers (the left column). In panel B, the pre-trial settlement bargaining, placebo split-the-difference offers are also frequent. However, in both the middle and right figures in panel B, the mass is more uniformly distributed across concession weights in the placebo cases than in the left column, suggesting that the *relative* likelihood of split-the-difference offers is the highest when considering the two most recent offers.

To formally test whether the masses at 0.5 are different in the true vs. placebo concession weights, we calculate the fraction of split-the-difference offers (as in Section 4.1) for the true concession weights, placebo concession weights, and their differences, as well as standard errors on each of these.[20] We also compute the fraction of split offers among placebo concession weights in the restricted sample. Table 4 presents the results. In each dataset, we detect a (statistically significantly) larger fraction of split offers using the true concession weights than the placebo concession weights, suggesting that players indeed rely more strongly on the two most recent offers than on earlier offers in determining a 50-50 split.

4.4 Splitting Based on Private Information

We also analyze whether agents tend to propose offers that split the difference between a publicly known threshold (a previous offer) and a *privately* known quantity that relates to agents' values. For this analysis, we exploit a number of privately known variables specific to several of our settings, including the secret reserve price known only to the seller in the used-car setting; the reserve/loss estimate known only to the "buyer" (the insurer) in the settlement negotiation setting; and the auto-accept or auto-decline prices that are known only to the seller in the eBay setting. The importance of these variables is that they are

[20]Standard errors on the difference are constructed from 100 nonparametric bootstrap-sample estimates of the true and placebo rates, sampling at the thread level.

Table 4: Fraction of Split Offers in True vs. Placebo Concession Weights

	True	Placebo	Difference	Placebo Restricted	Difference
		Panel A: Used Car Bargaining			
Split	0.1823	0.0503	0.1319	0.0614	0.1208
	(0.0019)	(0.0020)	(0.0023)	(0.0033)	(0.0034)
N	33,356	14,721		6,674	
		Panel B: Pre-trial Settlement Bargaining			
Split	0.0435	0.0218	0.0217	0.0190	0.0245
	(0.0004)	(0.0003)	(0.0005)	(0.0003)	(0.0005)
N	208,463	313,216		292,470	
		Panel C: eBay Best Offer Bargaining			
Split	0.1176	0.0437	0.0739	0.0424	0.0753
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
N	9,789,903	4,353,489		3,880,706	

Notes: Table shows the fraction of split offers among the true concession weights (first column) vs. the placebo concession weights (second column), as well as the difference of these means (third column). The fourth column shows the fraction of split offers among the placebo concession weights in the restricted sample, and the fifth column shows the difference between this mean and the true concession weight fraction. Standard errors are shown in parentheses. Panel A shows this analysis for the cars sample, panel B for the pre-trial settlement sample, and panel C for the eBay sample.

each privately known to only one side; thus, these variables allow us to examine whether agents' behavior is more consistent with equally splitting some quantity that is based on the public information contained in agents' previous offers or whether instead agents appear to offer prices that constitute an equal split relative to their *own private* information.

For this analysis, we follow our construction of *placebo* concession weights from Section 4.3, replacing some offer information with these privately known quantities. We describe this analysis and results in detail in Appendix C. The bulk of the evidence from this exercise suggests that split-the-difference behavior is less related to such privately known values and more related to the two most recent offers.

4.5 When are Split-the-Difference Offers Made?

It is possible that the set of cases in which split-the-difference offers might occur is constrained, either because the mid-point between the two previous offers might be lower than the seller's true willingness to sell or higher than the buyer's willingness to pay, or because a very demanding offer from the opponent might be seen as "unfair" (a buyer low-balling the seller, for example) and thus not deserving of a "good faith" split-the-difference counteroffer. Both of these arguments suggest that when a buyer's offer is only a small fraction of the seller's initial offer, we are unlikely to observe a subsequent split-the-difference offer, with the same being true for relatively high seller offers. Conversely, when the difference between previous offers is very small—for instance, within a few percentage points—the size of potential surplus may be small, and fairness concerns might be less relevant.

Figure 5 examines these patterns, showing the probability that an agent makes a split offer in round $t = 3$ as a function of the ratio of the first two offers. Regardless of the order of the moves (that is, whether the seller or the buyer make the first offer in the thread), this is the ratio of the buyer's first offer to the seller's first offer; thus, the ratio is always between zero and one. We cannot examine this in the trade setting, where the first offer is zero.

In all six datasets in Figure 5, we find that the probability of a split offer is lower when the buyer's offer is a very small fraction of the seller's (that is, when the gap is relatively large), and then increases initially. When the buyer's and seller's previous offers are relatively close, the chance of a split counteroffer decreases again in the TV, rides, and housing cases. In these cases, split offers are most likely when the ratio of previous offers lies in the 60%-90% range. For these values, where the opponent's offer is not too "unreasonable," but the gains from bargaining may still be substantial, agents are most likely to adopt the fairness norm. For the other three settings—cars, settlement, and eBay, which also correspond to the cases where we have the most data—we see the frequency of split offers increase globally as the ratio increases. Interestingly, these results are related to the findings by Karagözoğlu and Riedl (2015) that, in the context of unstructured bargaining lab experiments, the difference between the first two offers in a negotiation (i.e., tension or conflict in opening offers) is a very good predictor of bargaining duration and the probability of disagreement.

Figure 5: Split Probability as a Function of Ratio of First Two Offers

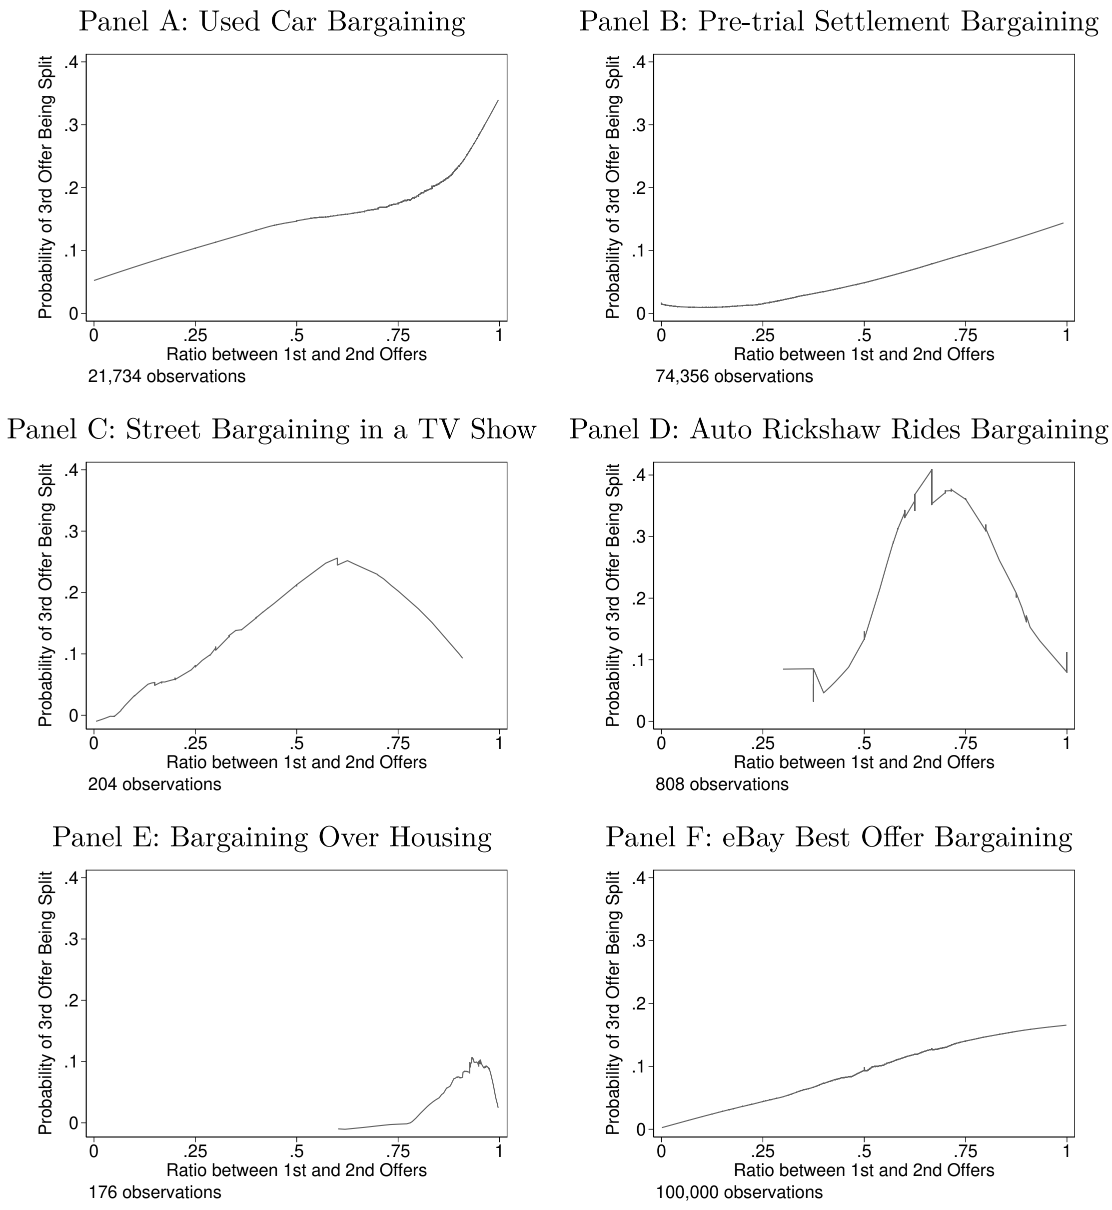

Notes: Each panel shows, on the vertical axis, the probability that the third offer in the sequence is a split offer. On the horizontal axis is the ratio of previous buyer offer and previous seller offer, regardless of order. The line in each plot is a weighted local linear fit. To facilitate computation of the local linear estimator in the eBay setting, we use a random sample of 100,000 threads.

5 What Model Could Generate This Behavior?

We now turn to the question of what theoretical models of bargaining might generate this type of behavior. In doing so, we offer several approaches. First, we consider whether such behavior can arise in a standard game-theoretic framework—perfect Bayesian equilibrium (PBE). The answer does not follow from prior work, because there is no known characterization of the PBE offer sequences under two-sided incomplete information. We prove that there indeed exists a PBE with split-the-difference behavior, but there are many other equilibria without it. Hence, while PBE is not falsified by splitting the difference, it neither explains nor predicts the behavior, nor does it address why splitting the difference could be viewed as fair by agents with private information.

In response to these shortcomings of a PBE approach, we propose an alternative argument that, without imposing a complete equilibrium model, addresses what we see as the most important question: why might agents view an equal division of the difference between the last two offers as a fair outcome, even though in reality it does not correspond to realized surplus? We propose a notion of inference that addresses this question.

We then discuss several other possible motivations for split-the-difference behavior. Each of these explanations—including those that we flesh out in more detail, the PBE and inference argument—are speculative in nature. They do not yield testable predictions that allow us to confidently reject one explanation over the other. We view our theoretical contribution in this Section as providing a first pass at incorporating into bargaining models the (empirically relevant) notion of splitting the difference between the two most recent offers.

5.1 Model Preliminaries

We consider an alternating-offer game with two-sided incomplete information: neither the buyer nor the seller knows the other party's valuation for the good.[21] The buyer has a value b and the seller a value s, each of which are in $[0, 1]$, drawn independently according to CDFs $F_b : [0, 1] \to [0, 1]$ and $F_s : [0, 1] \to [0, 1]$. Time is discrete, and the discount factor for both players is $\delta \in (0, 1)$. Upon receiving an offer at time t, the receiver either accepts, quits, or proposes a counteroffer for time $t + 1$. If an offer of price p is accepted at time t, the buyer's present discounted utility is $\delta^{t-1}(b - p)$ and the seller's is $\delta^{t-1}(p - s)$. Without

[21] As in our empirical analysis, our theoretical analysis does not consider bargaining over multiple issues at once.

loss of generality, we focus on a game in which the seller is the first proposer.

5.2 A PBE with Split-the-Difference Behavior

We construct a PBE of the bargaining game that has split-the-difference behavior at every history on the path of play. For ease of construction, we assume that both values are uniformly distributed.

Proposition 1. *Suppose $F_s(v) = F_b(v) = v, v \in [0, 1]$. There exists $\underline{\delta} < 1$ such that, if $\delta > \underline{\delta}$, then there is a PBE of the alternating-offer bargaining game such that, for every realization of s and b, the resulting sequence of offers satisfies $p_t = 0.5p_{t-2} + 0.5p_{t-1}$ for all $t \geq 3$.*

The proof of this proposition, and all other proofs, are found in Appendix A. The proof is quite involved, requiring that we separately handle possible deviations at the first offer and those at later offers, and that we construct an appropriate punishment scheme to prevent such deviations for buyers or sellers.

Proposition 1 demonstrates that split-the-difference behavior is at least consistent with a PBE. However, the equilibrium we construct is somewhat *post hoc*: the behavior is built in, with no explanation of *why* agents might gravitate toward it.[22] Furthermore, splitting the difference is by no means the unique PBE prediction. For bargaining under incomplete information, it is widely acknowledged that even sequential equilibrium has only weak implications for on-path behavior. Gul and Sonnenschein (1988) write that, with one-sided incomplete information, "almost any pair of strategies that is sequentially rational *along the equilibrium path* can be supported as a sequential equilibrium." The equilibrium we have constructed is instead for the case of two-sided incomplete information, but the same point applies: the construction can be easily modified to support on-path offers that are *unequal* splits between the most recent two.

[22] For example, one feature of this PBE is that split offers, if accepted, generically do not result in an equal split of the underlying *surplus* (either in ex-ante or ex-post terms); one player inevitably takes home more of the surplus than the other, and there is a shift in *which player* gets more expected surplus depending on which round the game ends at when it ends in agreement. To see this, suppose that the bargaining ends in agreement at round t, for $t \geq 3$, given the behavior specified in the proof of Proposition 1. The expected surplus of the offering player is $\frac{1}{2^{t-1}}\left(\frac{1}{6} + \frac{\alpha}{2}\right)$, and the expected surplus of the receiving player is $\frac{1}{2^{t-1}}\left(\frac{5}{12} - \alpha\right)$. For $\delta > 0$, the receiving player has strictly more expected surplus, and the inequality is strict even in the limit as $\delta \to 1$.

5.3 How Could Agents Possibly View Splitting the Two Most Recent Offers as "Fair"?

It seems plausible that, rather than some equilibrium notion, split-the-difference offers are supported by fairness norms, but such an explanation is less straightforward than it appears. Evidence from laboratory experiments indicating that subjects often favor equal divisions of a *known* pie has helped motivate the development of theoretical models involving social preferences, which account for agents' concerns for issues such as equity, social welfare, reciprocity, or social image (Bolton, 1991; Bolton and Ockenfels, 2000; Andreoni and Miller, 2002; Andreoni and Bernheim, 2009). However, standard behavioral economic theories that incorporate fairness concerns do not apply directly to our setting, because they are defined under complete information (Rabin, 1993; Fehr and Schmidt, 1999; Andreoni and Bernheim, 2009).[23] When bargaining under incomplete information, the buyer's last offer need not be equal to the seller's value, and *vice versa*, so splitting the last two offers does not in general split the surplus; thus, the behavior also does not correspond to any previous notion of fairness. Moreover, in the equilibrium constructed for Proposition 1, splitting the last two offers does not even split the *expected* surplus equally, even though both buyer and seller values are independently and identically distributed.

To illustrate the challenge of attempting to attribute split-the-difference offers to fairness, consider one natural line of inference agents might invoke in a negotiation that begins with a seller asking for \$100 and a buyer countering at \$50. Assuming that no player makes an offer that would be unprofitable if accepted, the buyer may infer that the seller values the item weakly less than \$100, and the seller may infer that the buyer values the item weakly more than \$50. These arguments alone offer no useful motivation for how a split-the-difference offer (\$75) could be viewed as equitable—the arguments suggest only that the size of the pie is at least -50. The arguments are not wrong; they are just not useful for connecting \$75 to some notion of fairness. It could be that the seller's value is \$0 and the buyer's value is \$500, in which case \$75 is far from the midpoint of the surplus.

How, then, could agents view splitting the two most recent offers as fair? What is missed in the above line of reasoning is the information contained in the fact that neither the \$100

[23] An exception is Kohler (2012), which develops an alternate-offers bargaining model in which one of the parties is privately informed about her social preferences. The size of the surplus is still commonly known to the players in that model, distinguishing it from the setting that we consider here.

nor the \$50 offer were accepted—a fact that is publicly available information to both agents (unlike their values). We demonstrate how the last offer that the buyer rejects can be viewed as an upper bound for the buyer's value in a sense that will be formalized below. In the context of the example above, rationality alone implies that any buyer with value less than \$100 should not accept the seller's offer. If the buyer's value is strictly above \$100, then rationality alone does not imply the buyer will reject. That is, there exist buyer beliefs about the seller's strategy that would rationalize accepting an offer of \$100. Hence, if the seller believes that the buyer is rational, she cannot rule out that she is facing a buyer with a value below \$100. But, depending on which rational strategy the seller believes the buyer is playing, she might be able to rule out values above \$100. Therefore, \$100 is the highest buyer value that the seller cannot rule out (under any beliefs) on the basis of the buyer's past rejection behavior.

The preceding argument applies symmetrically to the inferences that the buyer can make when the seller rejects an offer. Hence, under this inference, the last offer that the buyer rejects can be viewed as an upper bound for his value, and the last offer that the seller rejects can be viewed as a lower bound for her value. By themselves, the \$100 and \$50 prices carry little information about the size of the pie, but once coupled with the fact that they were rejected, an equal split can be explained as a fairness norm applied to these bounds on the *potential surplus*.

To formalize this argument, we introduce a few definitions. A **narrow strategy** σ_i for player $i \in \{s, b\}$ is a function from public histories to i's available actions.[24] A **belief system** for player i is a conditional probability system defined on $H \times \Sigma_{-i} \times \Theta_{-i}$, where H is the set of public histories, Σ_{-i} the set of opponent narrow strategies, and Θ_{-i} the set of opponent types. A pair (σ_i, θ_i) is **rational** if there exists a belief system μ such that σ_i is a sequential best reply to μ for type θ_i.[25] An offer sequence $\{p_t\}_{t=1}^{T}$ is **monotone** if the buyer's offers are strictly increasing, the seller's offers are strictly decreasing, and the highest buyer offer is no more than the lowest seller offer.[26]

[24] We call this *narrow* because it does not depend on i's type.

[25] Rationality is a weaker condition than extensive-form *rationalizability* (Pearce, 1984). Rationality only requires that agents play sequential best responses to some belief system. But we could require more strategic sophistication, stipulating that each agent believes that his opponent is rational, and believes that his opponent believes that he is rational, and so on. Extensive-form rationalizability requires *strong belief* in rationality, meaning that, at every history, each agent attributes to her opponent the highest level of strategic sophistication consistent with what has already occurred. See Battigalli and Siniscalchi (2002) for details.

[26] Our restriction to monotone sequences simplifies the theory. Appendix B demonstrates that sequences

Proposition 2. *Let h be a history with a monotone offer sequence $\{p_t^*\}_{t=1}^T$, with the seller making the latest offer $p_T \in (0,1)$.*

1. *For all buyer values $b < p_T^*$, for any buyer strategy $\sigma_B \in \Sigma_B(h)$ such that (σ_B, b) is rational, σ_B rejects p_T^* at h.*

2. *There exists $\underline{\delta} < 1$ such that for all discount factors $\delta \geq \underline{\delta}$, for all buyer values $b > p_T^*$, there exists $\sigma_B \in \Sigma_B(h)$ such that (σ_B, b) is rational and σ_B accepts p_T^* at h.*[27]

And symmetrically for histories at which the buyer makes the latest offer.

Proposition 2 provides an explanation for why largely rational agents could gravitate toward offers that are halfway between the two most recent offers. We take as given, from the vast evidence in the literature, that people value fairness, and in Proposition 2, we describe how split-the-difference offers could be viewed as fair. The dimension along which our argument weakens standard game-theoretic concepts is that we consider agents making inferences about an opponent's value based only on the history of offers that the opponent has *rejected*, not proposed. We believe this injects an element of realism into the theory: inference based on an opponent's binary decision to accept or reject a given offer is much less computationally intensive than the process of conditioning on the full history of the game and attempting to invert an opponent's continuously valued offer.[28]

Proposition 2 shows that, under this offer-rejection inference, a seller cannot rule out the possibility that the buyer's value is equal to (or less than) the most recent offer the buyer rejected, as any rational buyer would not have accepted an offer below his value. And for buyers with higher valuations, there exist rational strategies under which these buyers would have accepted the most recent offer of the seller. Therefore, the largest buyer type that the seller could not possibly rule out is a buyer with a value equal to the most recent offer rejected by the buyer, and analogously for the buyer's inference about the seller rejecting

violating a weak version of monotonicity are not common empirically.

[27] It is natural to ask whether the beliefs supporting σ_B are reasonable. With small modifications, our proof can be strengthened to show that there exists a common prior on values and costs, and a PBE of the bargaining game with that common prior, such that type b plays σ_B under the PBE.

[28] The idea of agents only making inferences based on rejections of previous offers, and not the levels of those offers, also arises in a number of previous theoretical analyses of incomplete-information bargaining such as Chatterjee and Samuelson (1988), Gul and Sonnenschein (1988), and Ausubel and Deneckere (1992). They examine equilibria in which the level of the offer made by a privately informed party is not informative to the counterparty; only the fact that the previous offer was rejected is informative.

the most recent offer. Thus, an agreement to split this difference equally can be thought of as a fair division of the potential surplus.

By this reasoning, a hypothetical buyer who has just rejected an offer of \$100 might make the following statement: "I just told you (the seller) that I won't pay \$100 for the item. You and I both know that, after such an action, I would never admit to really valuing it at more than \$100 (even if that were true), because I could always argue that if I valued it at more than \$100 I would have accepted. But we also know that I would certainly have rejected your offer if my value was \$100 or less." Under this argument, both agents know that \$100 is the *most optimistic* belief the seller could have about the buyer's value that the buyer could not dispute by appealing to rational rejection behavior. Because each player could make such a statement about the other's valuation, a split-the-difference offer could emerge as a "fair" resolution of these claims.[29]

A number of laboratory experiments find evidence that, in settings where agents have private information, they often take advantage of "moral wiggle room" that allows them to appear fair while, in actuality, they are reserving for themselves a larger share of the surplus (see Dana et al. (2007)). For example, in a dictator game where the proposer is privately informed of the size of the pie, Ockenfels and Werner (2012) and Dana et al. (2007) documents that the proposer often "hides behind the small cake," offering what would be an equal split of the pie if it were indeed small but actually retaining more money for himself. Huang et al. (2020) study lab subjects in ultimatum games under incomplete information.[30] In one treatment, the proposer (buyer) is unaware of the receiver's (seller's) cost, but the receiver knows both the buyer's value and receiver's cost. In their data, the modal proposer offer is exactly halfway between the proposer's value and the *most optimistic* belief that the proposer could hold about the receiver's cost (Huang et al., 2020, Figures B4.c and B4.d). In a separate study, Camerer et al. (2019) conduct an unstructured bargaining experiment in which one of the parties is privately informed about the size of the pie. The authors find that, when the uninformed party in the experiment is the first one to make a demand, she tends to ask for half of the largest possible value of the pie (Camerer et al., 2019, pp.

[29]Note that, in a setting of *complete* information, our theory—that agents favor an equal division of the most optimistic potential surplus that cannot be ruled out under any belief—corresponds to the standard notion of splitting a pie of known size.

[30]Note that ultimatum games and dictator games—which have been studied extensively in the lab—have starkly different strategic incentives compared to alternating-offer bargaining, and therefore not all of the treatment arms in Huang et al. (2020) or other related studies are easily comparable to our settings.

1878-1879). Our study of field data offers a similar insight to these experiments: agents in incomplete-information bargaining gravitate toward a notion of "fairness" that does not correspond to an equal split of the surplus but rather a split of the *potential* surplus under agents' most optimistic beliefs subject to this inference.[31]

Some of the empirical evidence from Sections 3–4 are consistent with this notion that agents consider the two most recent offers as a bound on surplus. Of particular note is the fact that the publicly available information contained in the two most recent offers is special: agents are more likely to make offers lying halfway between these offers than between earlier offers of the game or between privately know quantities. We also examine, in Appendix D, cases where only one previous offer exists (e.g. a seller's list price) and a buyer's value may reasonably be assumed to be larger than zero. In these cases, a similar idea of agents favoring an equal split of the largest potential surplus that cannot be rejected under any belief based on the current history of rejected offers would suggest that buyers in these cases would gravitate toward offers that split the difference between 0 and the seller's list price. We find this to be true. However, we do not view any of these empirical patterns as conclusive evidence in favor of one explanation or another.[32]

5.4 Alternative Explanations of Split-the-Difference Offers

In this section, we briefly consider several alternative theories for why a 50-50 split between offers is a modal outcome in real-world settings. First, rather than being driven by an equitable split of potential surplus, agents' split-the-difference behavior may arise simply because it is *easier* (cognitively) to select the midpoint between the past two offers than to compute the optimal (surplus-maximizing) offer. While this may be one reason why

[31] This insight is consistent with the view that agents have a concern for the *appearance* of fairness, rather than for fairness per se. Other experimental studies supporting such a view include Mitzkewitz and Nagel (1993), Güth et al. (1996), and Andreoni and Bernheim (2009). It may also be the case that agents wish to avoid negative outcomes (such as the opponent exiting) from making offers that the opponent may consider unfair; the economics literature on social norms stresses the role of informal enforcement as a key component of social interactions. See, for example, Elster (1989) and Fehr and Gächter (2000).

[32] An interesting avenue for future work would be to explore whether, in a laboratory setting, splitting the difference between previous offers occurs more frequently when agents have incomplete information than when they have complete information. If agents indeed favor an equal split of the surplus, splitting the difference between the two most recent offers should not occur as much in complete-information settings, where those offers reveal no information about the surplus, because the surplus is already known. (We thank Emin Karagözoğlu for this point.) An experiment in this direction by Navarro and Veszteg (2020) finds evidence that players in unstructured, complete-information bargaining games tend to agree on outcomes that equalize their payoffs. However, Navarro and Veszteg (2020) do not compare complete-information to incomplete-information environments and do not analyze the evolution of offers within bargaining threads.

50-50 offers are selected, this argument alone fails to explain much of the empirical patterns surrounding split offers. In particular, it cannot explain why split offers are *accepted* more frequently and cause less exit, nor would it explain why split offers are frequently followed by other split offers. Furthermore, in high-stakes negotiations—some of which may involve highly experienced professionals, such as insurance companies, car dealers, or trade negotiators—one would suppose that computational constraints are relatively less important.[33] Yet our results show that split offers are common in each of these settings.

A second alternative is that split-the-difference offers may be justified by a notion of regret minimization.[34] Like in the previous paragraph, suppose a seller behaves as though the current offer is the final one. The ex-post *regret* of the seller can be defined as the difference between that offer and the buyer's value; this corresponds to the amount of money the seller would have learned she had missed out on if she were to discover the buyer's value with certainty ex-post.[35] If the buyer's value is *known* to be between the two most recent offers p_t and p_{t-1}, and the seller's value is no more than $\min\{p_t, p_{t-1}\}$, then the split-the-difference offer $\frac{p_t+p_{t-1}}{2}$ minimizes regret. This theory by itself, however, does not offer an explanation for the restrictions on the supports of values (e.g., the buyer's value lying between the two most recent offers).

Yet another possibility is that agents indeed value fairness, but do not view the two previous offers as some notion of the surplus. Instead, agents may focus on the fact that, when two previous offers exist in a given negotiation, any price the agents agree on will most likely lie between those two most recent offers, involving some degree of concession from one party or the other relative to their most recent offers. A split-the-difference offer requires the same level of concession from both players, and may thus be viewed as fair from the standpoint of the concession size. While this explanation neglects any reference to utility theory, it is possible that behavioral agents' reliance on such a heuristic motivates split offers.

[33] Conversely, we do not necessarily expect fairness considerations to scale down with negotiation stakes. In fact, the evidence emerging from complete-information lab-in-the-field ultimatum game experiments involving large sums (from the subjects' perspective) is mixed regarding the effect of stakes on outcomes. Most experiments largely reproduce the typical results from the lab, suggesting that concerns about fairness persist when stakes are high (Slonim and Roth, 1998; Cameron, 1999). But see Andersen et al. (2011) for evidence that, when sums get extremely large, subjects might start proposing and accepting unequal splits of the surplus.

[34] We thank Alex Wolitzky for this point.

[35] This is the standard definition of the *minimax regret decision criteria*, dating back to Savage (1951); see Stoye (2011) for a recent treatment of the literature.

An alternative, more ambitious, approach would be to define a new kind of equilibrium that incorporates fairness concerns under incomplete information, and then prove that all the strategy profiles that it selects exhibit split-the-difference behavior. We have not pursued this approach, partly in the interest of brevity, and partly because our field settings are varied enough that it seems implausible to assume that all strategic behavior arises from a single equilibrium. Instead, with our two primary theoretical approaches in Sections 5.2–5.3, we have sought to demonstrate, in Proposition 1, that split-the-difference behavior can indeed arise in a traditional game-theoretic model, but with little predictive power. And we introduce Proposition 2 to explain one key conundrum: why a 50-50 split of the two most recent offers could be seen as equitable, and why anything else might therefore be seen as generous or greedy. We hope the combination of these exploratory models and the empirical evidence in this study will motivate continued work in incorporating these patterns into bargaining theory.

6 Conclusion

Our study provides extensive evidence from several unique empirical settings—cars, insurance claims, entertainment, transportation, housing, trade, and eBay—that negotiating agents gravitate toward offers that split the difference between the two most recent offers. Split offers are more likely to be accepted by the receiver, less likely to be followed by a breakdown of negotiations, and are more likely to be followed by subsequent split offers. We also show that it is the two most recent offers in particular that agents are most likely to favor splitting equally, and that split-the-difference behavior is constrained by the disparity between the two previous offers. Finally, we demonstrate that some agents are more likely than others to follow split-the-difference patterns of behavior.

From prior experimental work on ultimatum games and bargaining with complete information, it is reasonable to believe that 50-50 surplus-sharing is supported by fairness norms. But in our settings, where agents have *incomplete information*, it is not obvious *a priori* how they could possibly view a 50-50 split of the two most recent offers as a "fair" outcome. Our study demonstrates that this behavior can arise in a PBE, but such behavior does not actually represent an equal split of surplus between agents, and the equilibrium is far from unique. Thus, while split-the-difference behavior is compatible with a PBE, it is not uniquely predicted by it. As an alternative, we offer a new argument rationalizing split-the-difference behavior: the two most recent offers constitute bounds on potential surplus corresponding to the most optimistic inference that agents could make about one another when those inferences are based only on previous *rejections* and on the common knowledge that each agent is rational. We offer several other possible explanations of the behavior. While we provide some possible theoretical explanations, we do not believe these hold the last word on the subject. Rather, we hope they may help put some structure on future theories.

Our evidence suggests that the central role of fairness norms in popular press accounts of bargaining and laboratory experiments is empirically important in a wide range of real-world negotiations. We see this both as a promising window into the effects of behavioral economics in settings with information asymmetries, as well as a cautionary tale for the estimation of structural bargaining models: the majority of the existing approaches to bring bargaining models to the data omit systematic behavioral patterns, which may result in biased estimates. Conversely, the counterfactual effects of policy changes on the outcomes of negotiations are also likely to be affected by these policies' interaction with agents' preference for split-the-difference offers. Interesting avenues for future research would be to explore the welfare implications of these norms, which may be feasible given the richness of many of the datasets in our analysis.

References

Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., and List, J. A. (2011). Stakes matter in ultimatum games. *American Economic Review*, 101(7):3427–3439.

Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5):1607–1636.

Andreoni, J. and Miller, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753.

Ausubel, L., Cramton, P., and Deneckere, R. (2002). Bargaining with incomplete information. *Handbook of Game Theory*, 3:1897–1945.

Ausubel, L. M. and Deneckere, R. J. (1992). Bargaining and the right to remain silent. *Econometrica*, 60(3):597–625.

Babcock, L. and Laschever, S. (2008). *Ask for it: How women can use negotiation to get what they really want*. Bantam.

Backus, M., Blake, T., Larsen, B., and Tadelis, S. (2020). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *Quarterly Journal of Economics*, 135(3):1319–1361.

Bagwell, K., Staiger, R. W., and Yurukoglu, A. (2020). Multilateral trade bargaining: A first look at the GATT bargaining records. *American Economic Journal: Applied Economics*, 12(3):72–105.

Battigalli, P. and Siniscalchi, M. (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory*, 106(2):356–391.

Bellemare, C., Kröger, S., and Van Soest, A. (2008). Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities. *Econometrica*, 76(4):815–839.

Binmore, K., Shaked, A., and Sutton, J. (1985). Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75(5):1178–1180.

Binmore, K., Shared, A., and Sutton, J. (1989). An outside option experiment. *Quarterly Journal of Economics*, 104(4):753–770.

Bochet, O., Khanna, M., and Siegenthaler, S. (2023). Beyond the dividing pie: multi-issue bargaining in the laboratory. *Review of Economic Studies*, 90(5).

Bolton, G. E. (1991). A comparative model of bargaining: Theory and evidence. *American Economic Review*, 81(5):1096–1136.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193.

Camerer, C. F. (2011). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

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

Cameron, L. A. (1999). Raising the stakes in the ultimatum game: Experimental evidence from indonesia. *Economic Inquiry*, 37(1):47–59.

Chatterjee, K. and Samuelson, L. (1988). Bargaining under two-sided incomplete information: The unrestricted offers case. *Operations Research*, 36(4):605–618.

Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–685.

Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33:67–80.

Elster, J. (1989). Social norms and economic theory. *Journal of Economic Perspectives*, 3(4):99–117.

Fehr, E. and Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, 14(3):159–181.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868.

Freyberger, J. and Larsen, B. (2021). How well does bargaining work in consumer markets? A robust bounds approach. NBER Working Paper 29202.

Green, E. A. and Plunkett, E. B. (2022). The science of the deal: Optimal bargaining on ebay using deep reinforcement learning. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 1–27.

Gul, F. and Sonnenschein, H. (1988). On delay in bargaining with one-sided uncertainty. *Econometrica*, 56(3):601–611.

Gul, F., Sonnenschein, H., and Wilson, R. (1986). Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39(1):155–190.

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

Hernandez-Arenaz, I. and Iriberri, N. (2018). Women ask for less (only from men): Evidence from bargaining in the field. *Journal of Economic Behavior & Organization*, 152:192–214.

Huang, J., Kessler, J. B., and Niederle, M. (2020). Fairness concerns are less relevant when agents are less informed. *Experimental Economics*, forthcoming.

Jiang, Z. (2022). An empirical bargaining model with left-digit bias: A study on auto loan monthly payments. *Management Science*, 68(1):442–465.

Karagözoğlu, E. and Riedl, A. (2015). Performance information, production uncertainty, and subjective entitlements in bargaining. *Management Science*, 61(11):2611–2626.

Keniston, D. (2011). Bargaining and welfare: A dynamic structural analysis. Working Paper.

Kohler, S. (2012). Incomplete information about social preferences explains equal division and delay in bargaining. *Games*, 3(3):119–137.

Larsen, B. J. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies*, 88(2):851–882.

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

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

Navarro, N. and Veszteg, R. F. (2020). On the empirical validity of axioms in unstructured bargaining. *Games and Economic Behavior*, 121:117–145.

Ochs, J. and Roth, A. (1989). An experimental study of sequential bargaining. *American Economic Review*, 79(3):355–384.

Ockenfels, A. and Werner, P. (2012). 'hiding behind a small cake' in a newspaper dictator game. *Journal of Economic Behavior & Organization*, 82(1):82–85.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Pope, D. G., Pope, J. C., and Sydnor, J. R. (2015). Focal points and bargaining in housing markets. *Games and Economic Behavior*, 93:89–107.

Prescott, J., Spier, K. E., and Yoon, A. (2014). Trial and settlement: A study of high-low agreements. *Journal of Law and Economics*, 57(3):699–746.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302.

Roth, A. E. (1985). Toward a focal point theory of bargaining. In Roth, A. E., editor, *Game-Theoretic Models of Bargaining*, volume 259, pages 265–67. Cambridge University Press Cambridge, UK.

Roth, A. E. (1995). Bargaining experiments. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.

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

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and market behavior in jerusalem, ljubljana, pittsburgh, and tokyo: An experimental study. *American Economic Review*, 81(5):1068–1095.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67.

Silveira, B. S. (2017). Bargaining with asymmetric information: An empirical study of plea negotiations. *Econometrica*, 85(2):419–452.

Slonim, R. and Roth, A. E. (1998). Learning in high stakes ultimatum games: An experiment in the Slovak Republic. *Econometrica*, 66(3):569–596.

Stokey, N. L. (1981). Rational expectations and durable goods pricing. *Bell Journal of Economics*, 12(1):112–128.

Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6):2226–2251.

Thompson, L. L. (2020). *The Mind and Heart of the Negotiator.* Pearson, 7th edition.

Voss, C. and Tahl, R. (2016). *Never Split the Difference.* Random House Business.

For Online Publication Only

A Proofs

A.1 Proof of Proposition 1

The following preliminary lemma first establishes an equivalent way to represent sequences of split-the-difference offers.

Lemma 1. *Consider any sequence $\{x_t\}_{t=1}^{\infty}$, such that for $t \geq 3$, $x_t = 0.5x_{t-2} + 0.5x_{t-1}$. This is equal to the sequence $\{y_t\}_{t=1}^{\infty}$, defined by $y_t = \overline{x} + \left(-\frac{1}{2}\right)^{t-1}\alpha$, where $\overline{x} = \frac{1}{3}x_1 + \frac{2}{3}x_2 = \lim_{t\to\infty} x_t$ and $\alpha = x_1 - \overline{x}$.*

Proof. The first two terms are identical, since $y_1 = \overline{x} + x_1 - \overline{x} = x_1$ and $y_2 = \overline{x} - \frac{1}{2}(x_1 - \overline{x}) = \frac{3}{2}\overline{x} - \frac{1}{2}x_1 = x_2$. For any $t \geq 3$,

$$\frac{1}{2}y_{t-2} + \frac{1}{2}y_{t-1} = \overline{x} + \frac{1}{2}\left(-\frac{1}{2}\right)^{t-3}\alpha + \frac{1}{2}\left(-\frac{1}{2}\right)^{t-2}\alpha = \overline{x} + \left(-\frac{1}{2}\right)\left(-\frac{1}{2}\right)^{t-2}\alpha = y_t. \quad (7)$$

□

We now argue that, for δ close enough to 1, there exists a PBE in which every on-path offer is

$$p_t = \frac{1}{3} + \left(-\frac{1}{2}\right)^{t-1}\alpha, \quad (8)$$

where $\alpha = \frac{2(1-\delta)}{3(4-\delta)}$. We construct an equilibrium such that proposers are deterred from making offers not equal to equation (8) on the path of play.

For convenience, we will represent the game so that no player is called to play twice in a row; at each point, the receiver either accepts the current offer, chooses a counteroffer, or quits. This simplifies the application of the one-stage deviation principle, and has no substantial effect on the results.[36] The game proceeds as follows:

1. At $t = 0$, the **seller** makes an offer $p_1 \in \mathbb{R}$ or quits.

2. For all $t \in \{1, 3, 5, \ldots\}$, the **buyer** either accepts the offer p_t, chooses a counteroffer $p_{t+1} \in \mathbb{R}$, or quits.

[36] Collapsing these two adjacent decisions into one means that the one-stage deviation principle licenses us to directly compare the payoff from making any counteroffer p_{t+1} to the payoff from accepting p_t.

3. For all $t \in \{2, 4, 6, \ldots\}$, the **seller** either accepts the offer p_t, chooses a counteroffer $p_{t+1} \in \mathbb{R}$, or quits.

We construct an equilibrium with the following path of play, in which we denote $\phi \equiv \frac{4-\delta}{4(1-\delta)}$:

1. For all t, on the path of play, the offer p_t is as in equation (8).

2. At $t = 0$, the seller offers p_1 if $s \leq p_1$ and quits otherwise.

3. For all $t \in \{1, 3, 5, \ldots\}$, the buyer accepts if $b > \frac{1}{3} + \phi(p_t - \frac{1}{3})$, quits if $b < p_{t+1}$, and counters with p_{t+1} otherwise.

4. For all $t \in \{2, 4, 6, \ldots\}$, the seller accepts if $s < \frac{1}{3} - \phi(\frac{1}{3} - p_t)$, quits if $s > p_{t+1}$, and counters with p_{t+1} otherwise.

5. Whenever a player observes any off-path behavior, she believes her opponent is the weakest possible type ($b = 1$ or $s = 0$, respectively).

We proceed by the one-stage deviation principle, showing that no one-stage deviation is profitable in expectation for any type.

A.1.1 Deviations that Do Not Involve Off-path Offers

We start by proving that at any public history such that all previous offers are as specified by Equation (8), the active player cannot profitably deviate by accepting, quitting, or making the next offer specified by (8).

Consider the position of a buyer who receives an offer p_t in an odd period (a symmetric argument works for a seller who receives an offer in an even period). The seller made an offer of p_t, so the buyer can infer $s \leq p_t = \frac{1}{3} + \left(\frac{1}{2}\right)^{t-1} \alpha$. The seller rejected an offer of p_{t-1}, so the buyer can infer $s \geq \frac{1}{3} - \phi(\frac{1}{3} - p_t) = \frac{1}{3} - \phi \left(\frac{1}{2}\right)^{t-2} \alpha$. If both these bounds hold with equality, we have that $s \sim U[\frac{1}{3} - \phi \left(\frac{1}{2}\right)^{t-2} \alpha, \frac{1}{3} + \left(\frac{1}{2}\right)^{t-1} \alpha]$.

Consider the threshold type $\underline{b}$ that is indifferent between accepting and rejecting. This threshold type satisfies $\underline{b} - \frac{1}{3} = \phi(p_t - \frac{1}{3}) = \phi \left(\frac{1}{2}\right)^{t-1} \alpha$. The threshold type's payoff from accepting at t is

$$\underline{b} - p_t = \underline{b} - \left(\frac{1}{3} + \left(\frac{1}{2}\right)^{t-1} \alpha\right) = (\phi - 1) \left(\frac{1}{2}\right)^{t-1} \alpha. \tag{9}$$

The threshold type will make the counteroffer at $t+1$, since $\underline{b} > \frac{1}{3} > p_{t+1}$, and if that offer is rejected, will accept the seller's offer at $t+2$, since $\underline{b} - \frac{1}{3} > \phi(p_{t+2} - \frac{1}{3})$. Given the assigned strategy profile, the expected utility from rejection is

$$\begin{aligned}
&\frac{\frac{3\phi}{4}}{\phi+\frac{1}{2}}\delta\left(\underline{b}-p_{t+1}\right)+\frac{\frac{\phi}{4}+\frac{1}{8}}{\phi+\frac{1}{2}}\delta^2\left(\underline{b}-p_{t+2}\right)\\
&\quad=\frac{\frac{3\phi}{4}}{\phi+\frac{1}{2}}\delta\left(\phi\left(\frac{1}{2}\right)^{t-1}\alpha+\left(\frac{1}{2}\right)^{t}\alpha\right)+\frac{\frac{\phi}{4}+\frac{1}{8}}{\phi+\frac{1}{2}}\delta^2\left(\phi\left(\frac{1}{2}\right)^{t-1}\alpha-\left(\frac{1}{2}\right)^{t+1}\alpha\right)\\
&\qquad=\left(\frac{1}{2}\right)^{t-1}\alpha\left[\frac{\frac{3\phi}{4}}{\phi+\frac{1}{2}}\delta\left(\phi+\frac{1}{2}\right)+\frac{\frac{\phi}{4}+\frac{1}{8}}{\phi+\frac{1}{2}}\delta^2\left(\phi-\frac{1}{4}\right)\right]. \qquad (10)
\end{aligned}$$

Equating (9) and (10) yields

$$\phi-1=\frac{\frac{3\phi}{4}}{\phi+\frac{1}{2}}\delta\left(\phi+\frac{1}{2}\right)+\frac{\frac{\phi}{4}+\frac{1}{8}}{\phi+\frac{1}{2}}\delta^2\left(\phi-\frac{1}{4}\right) \qquad (11)$$

which implies

$$\phi=\frac{4-\delta}{4(1-\delta)}. \qquad (12)$$

To make the stationary argument above work even for $t=1$, we have chosen α so that, conditional on the seller making the offer instead of quitting at $t=0$, we have $s \sim U[\frac{1}{3}-\phi\left(\frac{1}{2}\right)^{t-2}\alpha, \frac{1}{3}+\left(\frac{1}{2}\right)^{t-1}\alpha]$. This is true if $\phi\alpha=\frac{1}{6}$, that is, $\alpha=\frac{2(1-\delta)}{3(4-\delta)}$.

Take any $t \in \{1,3,5,\ldots\}$. For a seller with $s > p_t$, quitting yields a payoff of 0, and making the offer yields no more than 0. For a seller with $s \leq p_t$, quitting yields a payoff of 0, and making the offer yields at least 0. For a buyer receiving an offer of p_t, the payoff from accepting minus the payoff from rejecting is non-decreasing in b; since the threshold type is indifferent, no type can benefit from deviating. A symmetric argument holds for $t \in \{2,4,6,\ldots\}$.

A.1.2 Deviations at Histories Involving Off-path Offers

We now construct beliefs and strategies that yield the path of play in Proposition 1, and prove these are a PBE. It is sufficient to specify, for each t and each offer not equal to $\frac{1}{3}+\left(-\frac{1}{2}\right)^{t-1}\alpha$, beliefs and continuation strategies such that:

1. Those beliefs and continuation strategies form a PBE of the subgame that follows from the deviation.

2. On the path of play, no type of the player called to play at t has a profitable one-stage deviation to a counteroffer of $p_{t+1} \neq \frac{1}{3} + \left(-\frac{1}{2}\right)^t \alpha$.

We will deal with two cases separately: deviations at the first offer p_1 and deviations at any subsequent offer.

Deviations at the first offer

If the seller chooses $p_1 \neq \frac{1}{3} + \alpha$, then we specify the beliefs $s = 0$ and $b \sim U[0, 1]$.[37] This corresponds to the "no-gap" case in models of bargaining with one-sided incomplete information.[38]

There exists a PBE of this bargaining game such that, if the buyer has not deviated previously, the following holds:

1. In even periods t, the lowest-cost seller $s = 0$ offers $p_{t+1} = \frac{\sqrt{1-\delta^2}}{1+\sqrt{1-\delta^2}}\bar{b}_t$, where $\bar{b}_t$ is the highest buyer type that has not yet accepted.

2. In odd periods t, the buyer accepts p_t if $b\sqrt{1-\delta^2} > p_t$ and counteroffers with $p_{t+1} = 0$ otherwise.

The seller's behavior in even periods is constructed as in Example 1 of Gul et al. (1986), which is due to Stokey (1981). For δ close enough to 1, the buyer's behavior in odd periods can be enforced by specifying the belief $b = 1$ for any deviating offer $p_{t+1} \neq 1$ (Gul and Sonnenschein, 1988, p. 610).

We specify that if the seller chooses $p_1 \neq \frac{1}{2} + \alpha$, then continuation play is as in the above equilibrium. Now we prove that, for δ close enough to 1, no type of s profits by deviating in this way.

After a first-offer deviation, no seller strategy results in an accepted price greater than $\sqrt{1-\delta^2}$. Consequently, an upper bound for the seller's payoff after a first-offer deviation is $\max\{\sqrt{1-\delta^2} - s, 0\}$.

Case 1: Suppose $s \leq \frac{1}{6}$. The buyer accepts an offer of $p_1 = \frac{1}{3} + \alpha$ if $b > \frac{1}{3} + \phi\alpha = \frac{1}{2}$. Thus a lower bound for the seller's payoff from this offer is $\frac{1}{2}\left(\frac{1}{3} + \alpha - s\right) \geq \frac{1}{2}(\frac{1}{6} + \alpha)$. An

[37] $b \sim U[0, 1]$ is implied by the usual requirement that the seller's actions cannot be informative about the buyer's type.

[38] This gap vs. no-gap language appears frequently in the incomplete-information bargaining literature. See, for example, Ausubel et al. (2002).

upper bound for the seller's payoff from deviating is $\sqrt{1-\delta^2}$.

$$\lim_{\delta\to 1}\frac{1}{2}(\frac{1}{6}+\alpha)=\frac{1}{12}>0=\lim_{\delta\to 1}\sqrt{1-\delta^2}. \tag{13}$$

Thus, for δ close enough to 1, the seller cannot gain by deviating at the first offer.

Case 2: Suppose $s>\frac{1}{6}$. The seller's payoff in the constrained game is at least 0. The seller's payoff from the deviation is no more than $\max\{\sqrt{1-\delta^2}-\frac{1}{6},0\}$, and for δ close enough to 1, $\sqrt{1-\delta^2}-\frac{1}{6}<0$, so the seller cannot gain by deviating at the first offer.

Deviations at later offers by the seller

Suppose that at $t\in\{2,4,6,\ldots\}$, the seller facing offer $p_t=\frac{1}{3}+\left(-\frac{1}{2}\right)^{t-1}\alpha$ makes an off-path counteroffer $p_{t+1}\neq\frac{1}{3}+\left(-\frac{1}{2}\right)^{t}\alpha$.

Upon this deviation, we specify the optimistic belief $s=0$. At t, we have that $b\sim U\left[\frac{1}{3}-\frac{1}{2^{t-1}}\alpha,\frac{1}{3}+\frac{1}{2^{t-2}}\phi\alpha\right]$. This is the "gap" case.

Seller punishment construction

We now construct a PBE of the alternating-offer bargaining game starting from a deviating offer by the seller, such that as δ goes to 1, the transaction price converges to strictly less than the lower bound on b. We denote that lower bound as $\underline{b}\equiv\frac{1}{3}-\frac{1}{2^{t-1}}\alpha$. The strategies are as follows:

1. If all the buyer's previous offers (since the deviation) are equal to $p_B\equiv\underline{b}\frac{\delta}{1+\delta}$, then the buyer of type b accepts an offer of p if

$$b-p\geq\delta(b-p_B) \tag{14}$$

   and counteroffers with p_B otherwise.

2. If the seller of type 0 receives an offer of p_B, then he accepts.

3. If the seller of type 0 receives an offer not equal to p_B, then he believes that $b=1$, and types $b=1$ and $s=0$ proceed to play as in the full-information alternating-offer bargaining game.

We now verify that this is an equilibrium via the one-stage deviation principle. By inspection, the assigned strategies are optimal if the buyer makes an offer not equal to p_B.

We now check histories at which the buyer has only made offers equal to p_B. Suppose the buyer receives an offer of p. Accepting p yields payoff $b-p$ and making a counter offer of p_B yields $\delta(b-p_B)$. Thus, he prefers accepting p to making a counteroffer of p_B if and only if $b-p \geq \delta(b-p_B)$. If he makes any counteroffer other than p_B, then the lowest price the seller will later accept is $\frac{\delta}{1+\delta}$, so his payoff is at most $\max\{\delta(b-\frac{\delta}{1+\delta}), 0\}$, which is strictly less than his payoff from counteroffering p_B.

Suppose the seller receives an offer of p_B at t. We now check that it is not profitable to deviate to any counter offer p. A one-stage deviation to p is accepted by the buyer at $t+1$ if equation (14) is satisfied, and otherwise the buyer counteroffers with p_B and is accepted at $t+2$.

Given the buyer's strategy, the seller believes that b is uniformly distributed between $\underline{b}$ and $\bar{b}$, for some $\bar{b} \leq \frac{1}{3} + \frac{1}{2^{t-2}}\phi\alpha$. Thus, the buyer's acceptance threshold is uniformly distributed between $\underline{\tau}$ and $\overline{\tau}$, where $\underline{\tau} = \underline{b}(1-\delta+\frac{\delta^2}{1+\delta}) = \underline{b}\frac{1}{1+\delta}$ and $\overline{\tau} = \bar{b}(1-\delta) + \underline{b}\frac{\delta^2}{1+\delta}$

Thus, making an offer of $p \in [\underline{\tau}, \overline{\tau}]$ yields payoff

$$\frac{\overline{\tau}-p}{\overline{\tau}-\underline{\tau}}\delta p + \frac{p-\underline{\tau}}{\overline{\tau}-\underline{\tau}}\underline{b}\frac{\delta^3}{1+\delta} \tag{15}$$

Note that $\bar{b} \leq 2\underline{b}$, since

$$\bar{b} \leq \frac{1}{3} + \frac{1}{2^{t-2}}\frac{1}{6} \leq \frac{1}{2} \leq 2\left(\frac{1}{3} - \frac{1}{2^{t-1}}\frac{2(1-\delta)}{3(4-\delta)}\right) = 2\underline{b} \tag{16}$$

$\bar{b} \leq 2\underline{b}$ implies that $\overline{\tau} - 2\underline{\tau} + \delta^2\underline{\tau} \leq 0$, which implies that equation (15) is maximized at $p = \underline{\tau}$. Thus, the seller's maximum payoff from any counteroffer is $\delta\underline{\tau} = \underline{b}\frac{\delta}{1+\delta} = p_B$, so the seller cannot profit by rejecting an offer of p_B.

No profitable deviations by the seller

Suppose that the history of the game thus far is such that every offer has been consistent with equation (8). The seller faces an offer of $p_t = \frac{1}{3} - \frac{1}{2^{t-1}}\alpha$. Given the above punishment strategies, any one-stage deviation to a counteroffer $p_{t+1} \neq \frac{1}{3} + \frac{1}{2^t}\alpha$ is accepted only if $p_{t+1} \leq (1-\delta) + p_t\frac{\delta}{1+\delta}$, and otherwise leads to a transaction at price $p_t\frac{\delta}{1+\delta}$.

Since $p_t = \frac{1}{3} - \frac{1}{2^{t-1}}\alpha$ is bounded away from 0 for all $t \in \{2,4,6,\ldots\}$, we can pick δ close enough to 1 so that for all $t \in \{2,4,6,\ldots\}$, $\delta > \sqrt{1-p_t}$, which implies that $(1-\delta) + p_t\frac{\delta}{1+\delta} < p_t$, and the seller cannot profitably deviate to an off-path offer.

Deviations at later offers by the buyer

Suppose that at $t \in \{1,3,5,\ldots\}$, the buyer facing offer $p_t = \frac{1}{3} + \left(-\frac{1}{2}\right)^{t-1} \alpha$ makes an off-path counteroffer $p_{t+1} \neq \frac{1}{3} + \left(-\frac{1}{2}\right)^{t} \alpha$. Upon this deviation, we specify the optimistic belief $b = 1$. At t, we have that $s \sim U\left[\frac{1}{3} - \frac{1}{2^{t-2}}\phi\alpha, \frac{1}{3} + \frac{1}{2^{t-1}}\alpha\right]$.

Buyer punishment construction

We now construct a PBE of the alternating-offer bargaining game starting from a deviating offer by the buyer. The construction is essentially symmetric. We denote $\overline{s} \equiv \frac{1}{3} + \frac{1}{2^{t-1}}\alpha$. The strategies are as follows:

1. If all the seller's previous offers (since the deviation) are equal to $p_S = \overline{s} + (1-\overline{s})\frac{1}{1+\delta}$, then the seller of type s accepts an offer of p if

$$p - s \geq \delta(p_S - s) \tag{17}$$

   and counteroffers with p_S otherwise.

2. If the buyer of type 1 receives an offer of p_S, then he accepts.

3. If the buyer of type 1 receives an offer not equal to p_S, then he believes that $s = 0$, and types $b = 1$ and $s = 0$ proceed to play as in the full-information alternating-offer bargaining game.

The argument proceeds as before. The only non-trivial part is checking one-stage deviations by the buyer of type 1 after receiving an offer of p_S.

As before, the seller's type is uniformly distributed between some $\underline{s} \geq \frac{1}{3} - \frac{1}{2^{t-2}}\phi\alpha$ and $\overline{s}$, so the seller's acceptance threshold for a counteroffer p is distributed uniformly between $\underline{\tau} \equiv (1-\delta)\underline{s} + \delta p_S$ and $\overline{\tau} \equiv (1-\delta)\overline{s} + \delta p_S = \overline{s} + (1-\overline{s})\frac{\delta}{1+\delta}$. The buyer's payoff from an offer of $p \in [\underline{\tau}, \overline{\tau}]$ is

$$\frac{p-\underline{\tau}}{\overline{\tau}-\underline{\tau}}\delta(1-p) + \frac{\overline{\tau}-p}{\overline{\tau}-\underline{\tau}}\delta^2(1-p_S) \tag{18}$$

Equation 18 is maximized at $p = \overline{\tau}$ if

$$1 - 2\overline{\tau} + \underline{\tau} - \delta(1-p_S) \geq 0. \tag{19}$$

Some algebra reduces this to

$$2(1-\overline{s}) \geq 1 - \underline{s} \tag{20}$$

which holds since $2(1-\overline{s}) \geq 1 \geq 1-\underline{s}$.

Substituting into equation (18), the buyer of type 1 has a payoff of no more than $(1-\overline{s})\frac{\delta}{1+\delta} = 1 - p_S$ from making a counteroffer, so he cannot profit by rejecting an offer of p_S.

No profitable deviations by the buyer Suppose that, so far, every offer has been consistent with equation (8). The buyer faces an offer of $p_t = \frac{1}{3} + \frac{1}{2^{t-1}}\alpha$. Given the above punishment strategies, any one-stage deviation to a counteroffer $p_{t+1} \neq \frac{1}{3} - \frac{1}{2^t}\alpha$ is accepted only if $p_{t+1} \geq \delta p_S = \delta(p_t + (1-p_t)\frac{1}{1+\delta})$, and otherwise leads to a transaction at price $p_t + (1-p_t)\frac{1}{1+\delta}$. Since p_t is bounded away from 1 for all $t \in \{1, 3, 5, \ldots\}$, we can pick δ close enough to 1 so that for all $t \in \{1, 3, 5, \ldots\}$, $p_t < \delta(p_t + (1-p_t)\frac{1}{1+\delta})$, so the buyer cannot profitably deviate to an off-path offer. $\square$

A.2 Proof of Proposition 2

Suppose $b < p_T^*$. Any σ_B that accepts p_T^* at h yields negative utility conditional on h, whereas rejecting at h and at all subsequent histories yields 0 utility. Hence, any σ_B that accepts p_T^* at h is not a sequential best reply to any conditional probability system, which proves Clause 1.

Suppose $b > p_T^*$. Let σ_B^* be the buyer strategy such that:

1. If all previous offers have been consistent with the sequence $\{p_t^*\}_{t=1}^T$ and equal to p_T^* for $t > T$, then the buyer makes the next offer in the sequence if proposing, and accepts any offer weakly less than p_T^* if receiving.

2. Else, if the first inconsistent offer was by the seller, then the buyer plays the full-information subgame-perfect equilibrium with buyer value p_T^* and seller cost 0.

3. Else, the buyer offers p_T^* and accepts an offer if and only if it is no more than p_T^*.

Symmetrically, let σ_S^* be the seller strategy such that:

1. If all previous offers have been consistent with the sequence $\{p_t^*\}_{t=1}^T$ and equal to p_T^* for $t > T$, then the seller makes the next offer in the sequence if proposing, and accepts any offer weakly more than p_T^* if receiving.

2. Else, if the first inconsistent offer was buy the buyer, then the seller plays the full-information subgame-perfect equilibrium with buyer value 1 and seller cost p_T^*.

3. Else, the seller offers p_T^* and accepts an offer if and only if it exceeds p_T^*.

Observe that $\sigma_B^* \in \Sigma_B(h)$ and $\sigma_S^* \in \Sigma_S(h)$.

We now specify beliefs for both buyer and seller: on the path of play of (σ_B^*, σ_S^*), the buyer believes that the seller's strategy is σ_S^* and her cost $s = p_T^*$, and the seller believes that the buyer's strategy is σ_B^* and her value is $b = p_T^*$. Following a deviation by the seller, the buyer believes that the seller's cost is 0 and that she will henceforth play the full-information subgame-perfect equilibrium with buyer value p_T^* and seller cost 0. Symmetrically, following a deviation by the buyer, the seller believes that the buyer's value is 1 and that he will play the full information SPE with buyer value 1 and seller cost p_T^*.

For δ close enough to 1, σ_B^* is a sequential best reply to the specified beliefs for a buyer with value $b > p_T^*$. For any history at which the seller made the first deviating offer, σ_B^* is a sequential best reply by construction, since it specifies that the buyer plays his part in the full-information SPE with buyer value p_T^* and seller cost 0. For any history consistent with (σ_B^*, σ_S^*), playing according to σ_B^* yields utility $\delta^{T-1}(b - p_T^*)$. By contrast, if the buyer deviates to an off-path offer, then the seller will henceforth only offer $\frac{\delta p_T^* + 1}{1+\delta}$ and will only accept offers that exceed $\frac{p_T^* + \delta}{1+\delta}$. Hence the buyer's utility following a deviating offer is upper bounded by $b - \frac{p_T^* + \delta}{1+\delta}$. Similarly, the buyer's utility from deviating to accept an earlier offer is upper bounded by $b - \min_{t \in \mathbb{T}_S} p_t^*$, where $\mathbb{T}_S$ denotes the periods strictly before T in which the seller made offers. Hence the buyer's gain from deviating first is upper bounded by the expression

$$\max\left\{b - \frac{p_T^* + \delta}{1+\delta}, b - \min_{t \in \mathbb{T}_S} p_t^*\right\} - \delta^{T-1}(b - p_T^*) \tag{21}$$

which, as $\delta \to 1$, converges (uniformly in b) to

$$\max\left\{\frac{p_T^* - 1}{2}, p_T^* - \min_{t \in \mathbb{T}_S} p_t^*\right\} < 0. \tag{22}$$

where the inequality follows since the offer sequence $\{p_t^*\}_{t=1}^T$ is monotone. This argument yields Clause 2 of Proposition 2. A symmetric argument applies to the seller. $\square$

B Additional Cleaning Steps of Datasets

In this section, we describe additional details about our cleaning procedure for each dataset.[39]

Before describing each dataset in more detail, we first show the number of observations that are dropped due to our final restrictions described at the beginning of Section 2. First, we drop any threads in which an agent's offer is an exact repeat of the opponent's previous offer (which logically should have led to the game ending in agreement), but additional offers are recorded afterward (Restriction 1 in Appendix Table A1). Second, we drop any threads in which the seller makes an offer that is strictly below a buyer's offer (Restriction 2). Third, we drop any threads in which a buyer makes an offer that is strictly below her own previous offer or a seller makes an offer that is strictly above her own previous offer (Restriction 3).

Appendix Table A1 shows that fewer than 2% of observations are dropped in most settings. In the settlement, TV show, and housing cases, 19%, 24%, and 16%, respectively, are dropped. In these settings, changes to the bargaining environment during the game (such as the arrival of new information), or simply misrecorded offers, may be more prevalent.

Table A1: Observations Dropped

	(1) Cars	(2) Settlement	(3) TV Show	(4) Rides	(5) Housing	(6) Trade	(7) eBay
# Threads $\geq$ 3 rounds	22,134	91,617	268	2,058	210	44,893	7,057,219
Restriction 1	22,069	87,390	264	2,058	198	44,732	7,048,997
Restriction 2	21,955	80,576	258	2,058	198	44,732	7,029110
Restriction 3	21,734	74,356	204	2,058	176	44,048	6,976,776
% Dropped	1.81%	18.84%	23.89%	0%	16.19%	1.88%	1.14%

Notes: Table shows the number of sequences/threads dropped due to the restrictions described at the beginning of Section 2.

B.1 Used Car Bargaining in the U.S.

In this data setting, when the auction price falls short of the seller's secret reserve price, the typical next step is for the seller and the highest bidder to engage in bargaining. In some cases, this negotiation may end quickly, with the seller deciding on the spot to not accept

[39]We do not describe any additional details for the eBay dataset, as we use the same cleaning steps as Backus et al. (2020).

the auction price and to not negotiate further. In these cases, a buyer who is not the highest bidder may, on occasion, approach the auction house salesperson and make an offer on the car, asking the auction house to call the seller to notify her. We include such sequences in our main analysis. Another possibility in cases when the auction price falls below the reserve price is that the seller and the highest bidder may engage in bargaining and, at the same time, a buyer who is not the highest bidder may approach the auction house and make a "backup offer," which the auction house can turn to if the negotiation between the seller and the highest bidder breaks down. As these backup offers are more complicated to interpret in our alternating-offer framework, we drop any sequences that contain backup offers. We also drop any inexplicable sequences, such as those that end with a counteroffer or end with a party accepting an offer after an opponent had already supposedly ended the negotiations. Neither of these restrictions excludes a large fraction of sequences.

B.2 Pre-trial Settlement Bargaining from Insurance Claims in the U.S.

The dataset contains multiple proposed *offers* (an offer here means a proposal from the insurer, i.e., the defendant) or *demands* (meaning a proposal from the claimant, i.e., the plaintiff) on the same day with no exact timestamp. Sometimes consecutive proposals from the same party are recorded, a pattern that does not satisfy the alternating-offer feature of the bargaining we are analyzing. To construct the bargaining sequences for analysis, we rearrange all bargaining sequences following the rules described below.

1. In each claim, we order all the offers and demands by date.
2. If there are multiple offers on the same day, we order them in the *increasing* amount.
3. If there are multiple demands on the same day, we order them in the *decreasing* amount.
4. If there are both offers and demands on the same day, we assume within the same day, demands and offers should be alternating.
5. If there are both offers and demands on the same day, whether this day starts with a demand or an offer follows the following rule:
    (a) If the date is the first date in a claim, we assume *offers* come first.

(b) If the date is not the first date in a claim

    i. If the most recent date before this day ends with an offer, we assume this day starts with a demand.

    ii. If the most recent date before this day ends with a demand, we assume this day starts with an offer.

6. After we arrange all offers and demands in this way, if there are consecutive offers or demands, we keep the last offer or demand and drop others.

Our main analysis pools together sequences that begin with an offer and those that begin with a demand. When we examine these separately, we find a similar mass point at 0.5. We also find that split offers are more likely to be accepted regardless of whether the sequence begins with a proposal from the insurer or the plaintiff. These results are shown in the first two columns of Appendix Table A2. Note that some sequences end at round 10, and we have no data on whether round 10 offers were accepted or not. This leads to the total number of observations in Table 1 being slightly larger than that in the analysis that examines acceptance, such as Table 2 and Appendix Table A2.

B.3 Street Bargaining from a TV Game Show in Spain

The bargaining sequences in the dataset are not necessarily alternating. Sometimes the proposer or the respondent can make consecutive offers. In such cases, we only keep the last offer in consecutive offers made by the same party. We also drop a small number of sequences that start from the respondent, so all remaining sequences start from the proposer.

B.4 Auto Rickshaw Rides Bargaining in India

As described in the paper, there are two broad types of bargaining: "real" bargaining and "scripted" bargaining. We keep all real bargaining sequences. For scripted sequences, we exclude offers and acceptance decisions from surveyors, as these are not actual decisions made by negotiating agents. This point leads to the total number of observations for the rides setting differing in analyses where we examine offers made (such as Table 1 and Figure 1) vs. analyses where we examine offers accepted (such as Table 2 or Appendix Table A2), because, in acceptance analysis, some scripted bargaining acceptances must be dropped, and yet some scripted bargaining offers can be included.

In columns 3–5 of Appendix Table A2, we repeat our analysis of the likelihood that split offers are accepted, doing so separately for real bargaining data, the scripted bargaining sequences that begin with a driver offer, and the scripted bargaining sequences that begin with a surveyor offer. As highlighted in Section 4.1, these scripted bargaining offers are the most interesting for this analysis, as these give something closer to a causal estimate of the effect of split offer (because surveyor's offers are assigned by the experiment designer rather than arising endogenously).[40] In this subset of the data, shown in columns 4–5, we find positive point estimates, and a particularly large and marginally significant positive point estimate in those sequences that begin with a surveyor offer. When we examine histograms of concession weights separately for these three subsamples, we also find a large mass point at 0.5 and sparse data at other points, as in the main sample.

B.5 Bargaining Over Housing

In this dataset, we observe a seller identifier (which is simply the address of the home), but we do not observe the buyer identifier. This means that when we observe multiple offers for the same house, these could come from the same buyer or different buyers. We are therefore required to make some assumptions to identify a distinct bargaining thread (i.e., a negotiation between a given seller and given buyer). For each observation, we can observe the agent commission type (call it AgentType), which is either a fixed amount (e.g., $5,000) or a percentage (e.g., 2%). We assume if two observations have different AgentType, they must be different buyers. If two observations have the same AgentType, they are not necessarily the same buyer. For all offers in each house-AgentType pair, we sort them by their submitted timestamp.

In the main sample, we only keep those house-AgentType pairs where all offers always weakly increase over time. Among these pairs, if the last offer is accepted or no offer is accepted, we assume all these offers belong to the same buyer. If some middle offer is accepted, we assume all offers up to the accepted offer belong to the same buyer and all offers after the accepted offer belong to another buyer.

For pairs where not all offers weakly increase over time, we follow the rules below to identify distinct buyers:

[40] Even though the offer sequence of the surveyor is assigned exogenously, the event that an offer is a "split" offer is not entirely exogenous, as this depends on the offer of the driver as well.

- When an offer is higher than the previous offer and the previous offer is not accepted, then the two offers come from the same buyer.
- When an offer is higher than the previous offer and the previous offer is accepted, then the second offer is made by a new buyer.
- When an offer is lower than the previous offer, then the second offer is made by a new buyer.

In doing so, we assume that the seller only bargains with one buyer at the same time. We are less confident in this assumption, and thus in the main sample, we exclude these observations. When these observations are included, we find a similar mass point at 0.5 in the concession weights and a positive (but insignificant) point estimate of the effect of a split offer on the probability of acceptance. The latter result is shown in column 6 of Appendix Table A2.

B.6 International Trade Tariff Bargaining

This dataset is publicly available on the journal website. We define one bargaining round as a combination of Proposer-Target-Stage-Date. A stage can be "request," "offer," "final offer," or "modification." An example of one bargaining round is the following: Australia makes requests to India on 10/16/1950. To create product-level concordances across negotiations, Bagwell et al. (2020) connect product-level descriptions to HS 1988 6-digit (HS6) codes. A product-level description can involve multiple tariff items. We refer to a combination of HS6 and tariff item as one product. If multiple observations exist for one product in one bargaining round, we use their average tariffs. A bargaining thread contains two counties negotiating over the tariffs for a certain product in a certain direction.

In the raw dataset, there are two types of tariffs: "Specific" means the request/offer is in dollars and "Ad Valorem" means the tariff is quoted as a fraction of prices. In most threads, only one type of tariff is used. In some rare cases, both types can exist. For each thread, if all observations have Ad Valorem tariff terms, we use this variable as the price variable. If not all observations involve Ad Valorem tariff terms, but all observations involve Specific tariff terms, we use this variable as the price variable. We drop threads with no consistent tariff types. We also drop observations with missing stage variables and with inconsistent country names.

Within each thread, we sort all observations by date. In a few cases, there are multiple observations on the same day, which come from modifications of offers/requests/final offers. We treat these modifications as having come after their corresponding proposals. For each thread, if there is a "final offer" or "modification of final offer" stage, we assume the price in this round is the final price. If there is no such stage, we assume this bargaining thread does not reach an agreement. For each round, if the price is equal to the final price, we assume this offer or request is accepted. Otherwise, it is rejected/countered. By construction, the price in the "final offer" or "modification of final offer" stage is accepted.

In all threads, 66.5% start with a request, 14.1% start with an offer, and the rest start a final offer. We drop threads that start with a final offer. We further restrict to threads with the following patterns: threads that start with a request (which includes items labeled "request," "request-offer," "request-final offer," and "request-offer-final offer") and threads that start with an offer (which includes items labeled "offer" and "offer-final offer"). These threads account for more than 80% of total threads.

For purposes of examining split-the-difference behavior, we consider two benchmarks as though they are default bargaining offers at the beginning of any given thread. These are a zero tariff, which can be seen as the initial request from any proposer, and the status quo tariff before the negotiations, which can be seen as the initial offer from the target. For threads that start with an offer and for the last round in threads with the pattern "request-offer-final offer," we replace $\gamma_{j,t}$ with $1 - \gamma_{j,t}$, so that $\gamma_{j,t}$ still measures the extent of concession in two consecutive offers from the target. We exclude the last round in threads with the pattern "offer-final offer," as this implies three consecutive offers and we cannot define concession in this case.

In our main analysis, we pool together sequences that begin with a request and those that begin with an offer. When we examine these two subsets of the data separately, we find a similar mass point at 0.5 in the concession weights, and a similarly strong positive and significant effect of split offers on the acceptance probability. These latter results are shown in the last two columns of Appendix Table A2. The effect of a split offer is especially large for sequences that begin with an offer.

C Additional Placebo Concession Analysis

C.1 Placebo Based on Secret Reserve Prices in Used Car Bargaining

In the used car bargaining data, we can observe the reserve price the seller reports to the auction house. This is a *secret* reserve price, in that it is not announced to the buyer. If the auction price is above this secret reserve price, the highest bidder is awarded the car. Otherwise, the seller and the buyer can bargain. The bargaining starts with the auction price from the buyer, $p_{j,1}$, and then alternates between the seller and buyer.

Below are two placebo concession weights we construct that rely on sellers' secret reserve prices:

- In round 3, the true concession of the buyer is $\gamma_{j,3} = \dfrac{p_{j,3} - p_{j,1}}{p_{j,2} - p_{j,1}}$. The placebo concession $\gamma_{j,3}^{pl}$ replaces the offer of the seller, $p_{j,2}$, with the reserve price.
- In round 4, the true concession of the seller is $\gamma_{j,4} = \dfrac{p_{j,4} - p_{j,2}}{p_{j,3} - p_{j,2}}$. The placebo concession $\gamma_{j,4}^{pl}$ replaces the offer of the seller, $p_{j,2}$, with the reserve price.

In many cases, the placebo concession is out of the range $[0, 1]$. Appendix Figure A1 plots the distribution of these placebo concession weights in rounds 3 and 4, where we limit to cases where the weight lies in $[0, 1]$. In the right panel of Appendix Figure A1, we observe a spike at 0.5, suggesting that some sellers do propose offers that split the difference between the buyer's most recent offer and the seller's secret reserve price. This tendency to splitting the difference is much weaker here, however, than in the histogram of the main sample shown in Figure 1. The spike at 0.5, for example, is similar to those at other levels of $\gamma_{j,4}^{pl}$ (such as those around 0.7 or 0.8), suggesting that a stronger norm for splitting the difference between the two most recent offers than between an offer and a privately known quantity.

Figure A1: Distribution of Placebo Concession, Used Car Bargaining

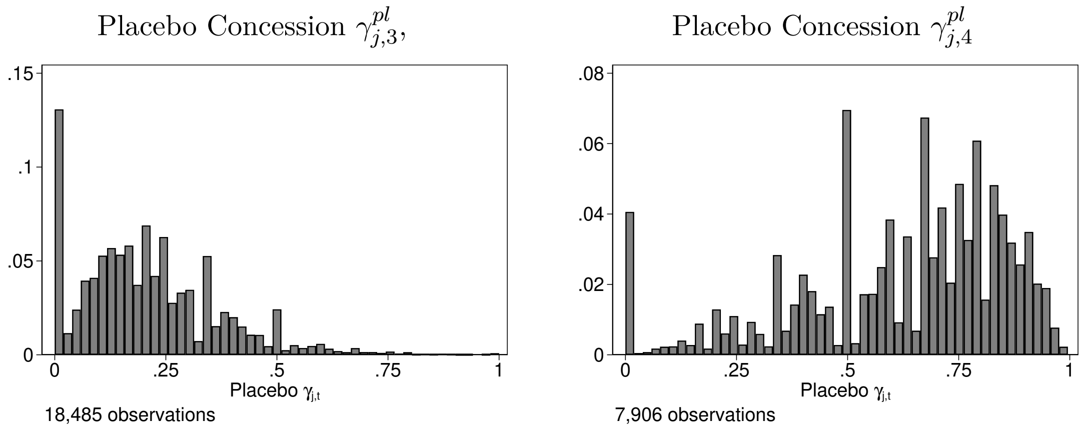

Notes: Each panel shows a histogram of the placebo concession weights in the used-car bargaining data where the seller's round 2 offer is replaced with the seller's secret reserve price. The left panel uses the round 3 placebo concession and the right panel uses the round 4 placebo concession.

C.2 Placebo Based on Private Reserve Estimate in Pre-trial Settlement Bargaining

In the pre-trial settlement data, we can observe the insurer's "reserve price," which is an estimate known only to the insurer of how much the insurer expects the case to cost the company. There are two types of bargaining sequences: those that start with a demand and those that start with an offer.

- If the sequence starts with a demand, in round 3, the true concession of the plaintiff is $\gamma_{j,3} = \frac{p_{j,3} - p_{j,1}}{p_{j,2} - p_{j,1}}$. The placebo concession replaces the offer of the insurer in round 2, $p_{j,2}$, with the reserve price.

- If the sequence starts with an offer, in round 3, the true concession of the insurer is $\gamma_{j,3} = \frac{p_{j,3} - p_{j,1}}{p_{j,2} - p_{j,1}}$. The placebo concession replaces the offer of the insurer in round 1, $p_{j,1}$, with the reserve price.

- If the sequence starts with a demand, in round 4, the true concession of the insurer is $\gamma_{j,4} = \frac{p_{j,4} - p_{j,2}}{p_{j,3} - p_{j,2}}$. The placebo concession replaces the offer of the insurer in round 2, $p_{j,2}$, with the reserve price.

- If the sequence starts with an offer, in round 4, the true concession of the plaintiff is $\gamma_{j,4} = \frac{p_{j,4} - p_{j,2}}{p_{j,3} - p_{j,2}}$. The placebo concession replaces the offer of the insurer in round 3, $p_{j,3}$, with the reserve price.

In many cases, the placebo concession is out of the range $[0, 1]$. Appendix Figure A2 plots the distribution of placebo concession in rounds 3 and 4 for sequences that start with a demand and start with an offer separately, limiting to those with concession weights in $[0, 1]$. We do not observe strong support for agents favoring offers that split the difference between the privately known reserve and the most recent public offer.

Figure A2: Distribution of Placebo Concession, Pre-trial Settlement Bargaining

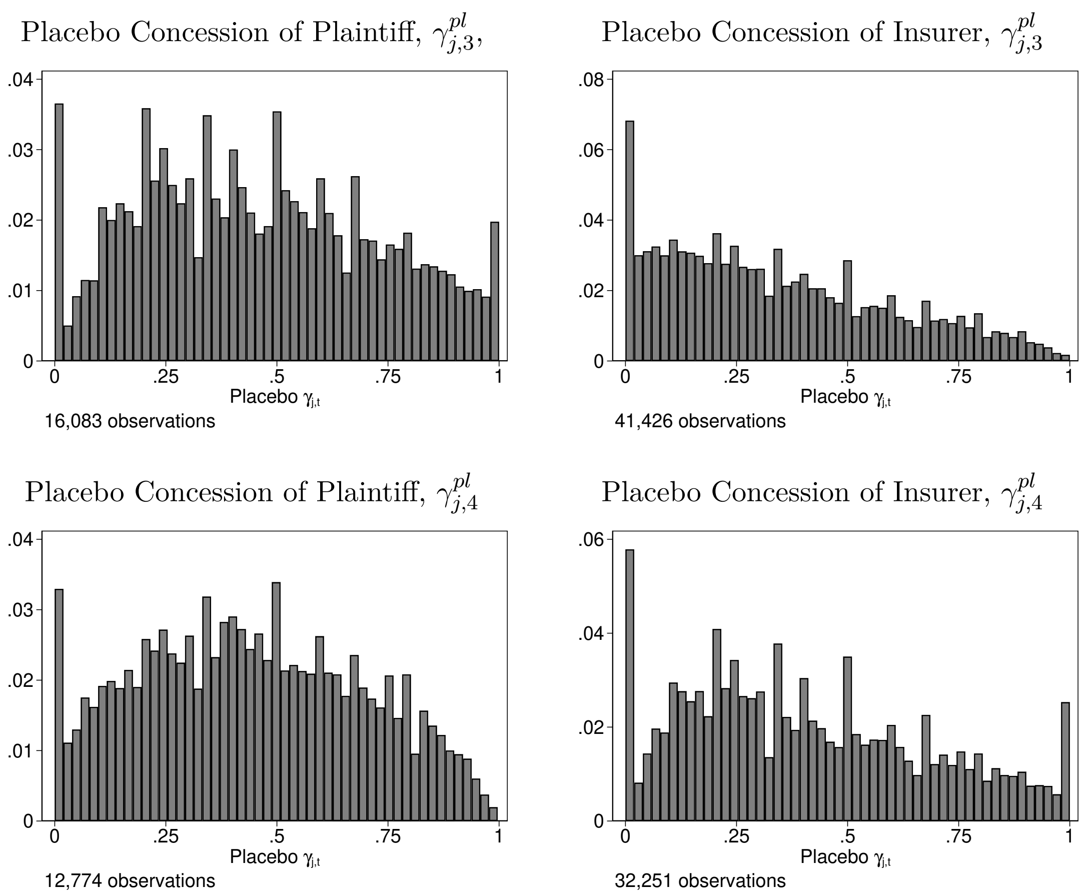

Notes: Each panel shows a histogram of the placebo concession weights in the settlement bargaining where an insurer offer is replaced with the insurer's privately known reserve amount.

C.3 Placebo Based on Secret Auto-Accept/Decline Prices in eBay Best Offer Bargaining

In the eBay data, we can observe a seller's *auto-accept* and *auto-decline* prices. Reporting these price thresholds is optional for a seller. When reported, these prices serve a similar role to proxy bids in an eBay auction. If a buyer makes an offer above the auto-accept price, the platform automatically accepts the offer on the buyer's behalf. If a buyer makes an offer below the auto-decline price, the platform declines. These prices are known only to the seller.

In the eBay setting, a bargaining sequence starts with the seller's list price and alternates between the buyer and seller. The list price is $p_{j,1}$, the initial offer from the buyer is $p_{j,2}$, and the first offer from the seller is $p_{j,3}$. In round 3, the true concession of the seller is $\gamma_{j,3} = \dfrac{p_{j,3} - p_{j,1}}{p_{j,2} - p_{j,1}}$. The placebo concession replaces the list price $p_{j,1}$ with the auto-accept or auto-decline price. In many cases, the placebo concession is out of the range $[0, 1]$. Appendix Figure A3 plots the distribution of placebo concession using auto-decline and auto-decline prices separately, limiting to those cases with concession weights in $[0, 1]$. We observe some mass at 0.5, but far less than in the main sample in Figure 1, suggesting again a stronger norm for splitting the difference between the two most recent offers than between an offer and a quantity known only to one party.

Figure A3: Distribution of Placebo Concession, eBay Best Offer Bargaining

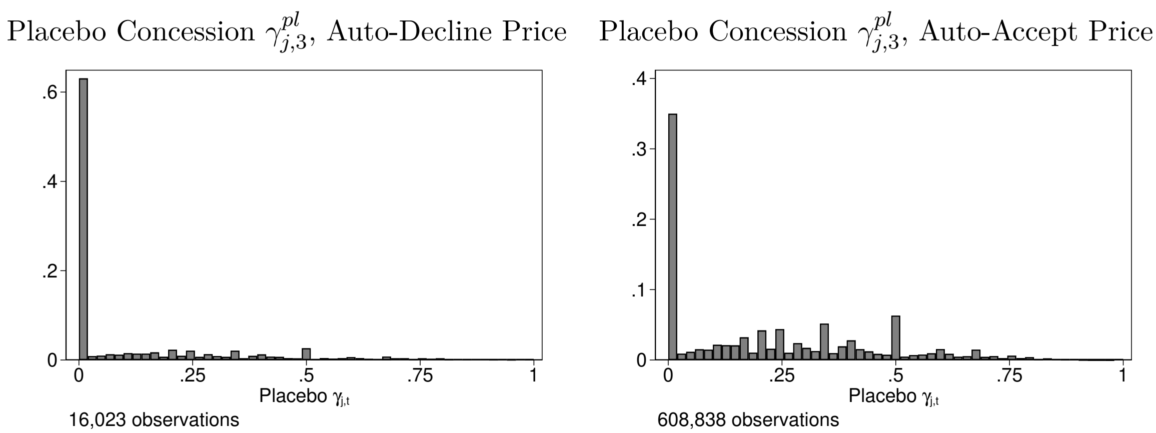

Notes: Each panel shows a histogram of the placebo concession weights in the eBay bargaining. In each case, the list price of the seller ($p_{j,1}$) is replaced with either the auto-decline (left panel) or auto-accept (right panel) price of the seller.

Table A2: Probability of a Split Offer Being Accepted

	Pre-trial Settlement Bargaining		Auto Rickshaw Rides Bargaining			Housing	Trade Tariff Bargaining	
	Insurer First	Plaintiff First	Real	Driver	Surveyor	Full	Request First	Offer First
Split	0.170***	0.165***	-0.0485	0.0333	0.165*	0.0983	0.0224***	0.341***
	(0.00602)	(0.00991)	(0.0458)	(0.0567)	(0.0944)	(0.114)	(0.00235)	(0.0218)
N	148173	55968	1224	1050	736	338	41473	5512
Order of $\gamma_{j,t}$	3	3	3	3	3	3	3	3
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Accept rate	0.37	0.28	0.30	0.13	0.16	0.60	0.05	0.25
Split rate	0.04	0.04	0.31	0.13	0.09	0.07	0.17	0.17
R^2	0.128	0.183	0.264	0.108	0.0932	0.0163	0.349	0.0814

Notes: Table shows the estimated coefficient on the split indicator from the regression described by equation (3), as in Table 2, using different subsamples in several of the data settings. Columns 1 and 2 correspond to settlement bargaining sequences that begin with a plaintiff or insurer proposing, respectively. Columns 3–5 correspond to the auto rickshaw rides data, with the real bargaining only in column 3, scripted bargaining beginning with a driver moving first in column 4, and scripted bargaining with a surveyor moving first in column 5. Column 6 uses the housing data, without excluding less trustworthy observations, as described in Appendix B.5. Columns 7–8 use the trade data, with sequences beginning with a *request* in column 7 and those beginning with an *offer* in column 8. The accept rate is the mean of the dependent variable and the split rate is the mean of the split indicator. Standard errors are shown in parentheses. $*: p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$

Table A3: Probability of a Split Offer Being Followed by Opponent Exit

	Auto Rickshaw Rides Bargaining			Trade Tariff Bargaining	
	Real	Driver	Surveyor	Request First	Offer First
Split	0.0271	-0.0240	-0.157*	-0.00385	-0.359***
	(0.0331)	(0.0652)	(0.0829)	(0.00408)	(0.0235)
N	1224	1050	736	41473	5512
Order of $\gamma_{j,t}$	3	3	3	3	3
Round FE	Yes	Yes	Yes	Yes	Yes
Exit rate	0.14	0.68	0.47	0.88	0.70
Split rate	0.31	0.13	0.09	0.17	0.17
R^2	0.0471	0.0659	0.221	0.151	0.0610

Notes: Table shows the estimated coefficient on the split indicator from the regression described by equation (4), as in Table 2, using different subsamples in several of the data settings. Columns 1–3 correspond to the auto rickshaw rides data, with the real bargaining only in column 1, scripted bargaining beginning with a driver moving first in column 2, and scripted bargaining with a surveyor moving first in column 3. Columns 4–5 use the trade data, with sequences beginning with a *request* in column 4 and those beginning with an *offer* in column 5. The exit rate is the mean of the dependent variable and the split rate is the mean of the split indicator. Standard errors are shown in parentheses. $*: p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$

D Splitting When There is Only One Previous Offer

We present a further test of the idea that negotiating agents' behavior is consistent with the potential surplus-split argument we propose in Section 5.3. Here, we consider settings in which the seller makes the first offer. When it is the buyer's turn to make the first counteroffer, there is only one previous offer to which agents can apply an optimistic-inference potential-surplus view. The most optimistic lower bound on the seller's value in such cases is therefore zero. For example, in the eBay setting, suppose a seller posts a list price of $100. If the buyer rejects this offer, the history of rejected offers from which the players can make inferences will consist only of the rejected list price. What, then, should be the notion of fairness toward which agents gravitate? In the spirit of our sequential rationality argument, the potential surplus at this stage of the game is the range from $0 to $100, an equal split of which is an offer of $50.

This same argument can be applied to any of our settings in which the seller moves first.[41] Among our data settings, these include the subset of settlement bargaining where the plaintiff starts, the subset of rides bargaining where the driver starts, and eBay bargaining. We also examine the trade bargaining case here, where, as we explain in Section 2.6, we set

[41]Note that cases in which a buyer moves first do not offer a zero lower bound: suppose the buyer first offers $200, and then it is the seller's turn. If the seller counters, she will clearly counter at a price above $200, but there is no natural upper bound that the agent would assume about the buyer's value.

zero as the first bargaining offer throughout the paper. In Figure A4, we examine where the buyer's offer in these settings lies relative to zero and relative to the seller's first offer. In the rides, trade, and eBay settings, we observe a mass point at 0.5 (although the mass point at 0.5 in the rides case is similar in size to that at about 0.6 or 0.7). These mass points are consistent with agents focusing on the halfway point of the potential surplus even at this early stage of the game when that potential surplus is defined by zero and the first seller offer.

Figure A4: Splitting the Difference Between Zero and the First Proposal

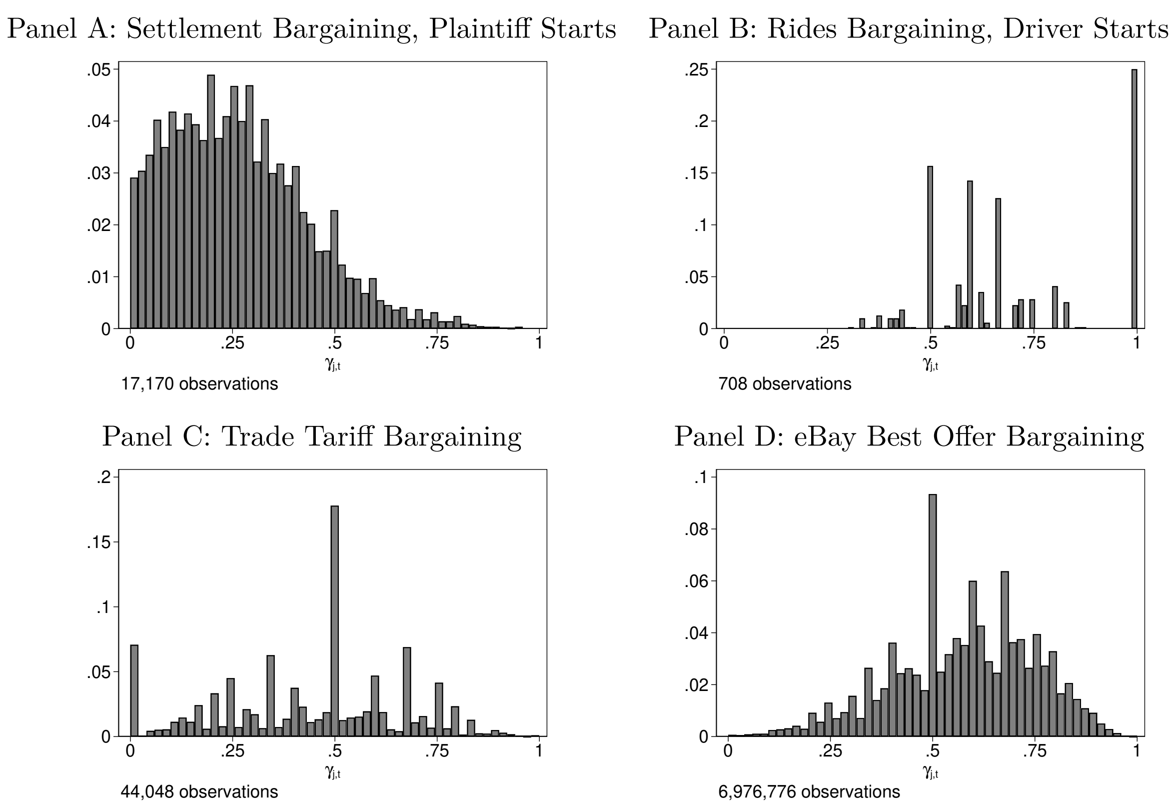

Notes: Histograms of concession weights defined as where the second offer in a thread lies relative to the first offer and zero, where the first offer is from a seller-like agent, meaning the agent who wants the price higher. Thus, panel A uses only settlement threads starting with the plaintiff in panel A and rides threads starting with the driver in panel B. In panel D—eBay bargaining—all threads start with a seller. In panel C—trade bargaining—we already treat zero as the first offer throughout the paper.

Pre-trial settlement bargaining in panel A is the one setting in Figure A4 in which the first offer is *not* frequently a 50-50 split between 0 and the first offer. There we observe a slight uptick at 0.5 relative to surrounding points, but the contrast is much smaller than in the other panels. This finding gives some insight into the limits of split-the-difference behavior. In this setting, the plaintiff's initial offer is often exorbitant relative to where the bargaining eventually ends; thus, a halfway offer on the part of the buyer (i.e., the insurer) would typically be overly generous (and might often be higher than the insurer's valuation, i.e., the expected court ruling). Housing bargaining would exhibit a similar result: it would be unheard of for a buyer to counter at a price that is 50% of the seller's list price. Such an offer would undoubtedly give surplus to the buyer, but it would surely not be seen as "fair" by the seller, and panel E of Figure 5 demonstrates that indeed buyers do not make such low-ball offers. These potential constraints on split offers may drive the broader patterns observed in Section 4.5.