Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions*

Matthew Backus[†] Thomas Blake[‡] Brad Larsen[§] Steven Tadelis[¶]

July 15, 2019

Abstract

We study patterns of behavior in bilateral bargaining situations using a rich, new dataset describing back-and-forth bargaining occurring in over 25 million listings from eBay's Best Offer platform. We demonstrate that several patterns in the data can be explained by existing theoretical models. These include some interactions ending quickly or ending in agreement after some delay, and stronger bargaining power or better outside options improving a player's outcome. Other robust patterns, however, remain unexplained by existing theories. These include negotiations resulting in delayed disagreement, gradually changing offers that are reciprocal, and most notably, we show that players exhibit a preference for making offers that split the difference between the two most recent offers. These robust patterns call for new explorations in the theory of bargaining.

JEL classifications: C78, D82, D83, M21.

*We are grateful to eBay for allowing us to make the data used herein accessible and publicly available. We thank Luis Cabral, Peter Cramton, Emin Karagozoglu, Saul Lach, and Axel Ockenfels for thoughtful comments. We thank Brenden Eum, Yuyang Jiang, Ziao Ju, Rebecca Li, Carol Lu, Caio Waisman, and Chuan Yu. for outstanding research assistance. Part of this research was supported by NSF Grant SES-1629060.

†Columbia University and NBER, matthew.backus@columbia.edu

‡eBay Research, tomblake@gmail.com

§Stanford University and NBER; bjlarsen@stanford.edu

¶UC Berkeley, NBER, CEPR, and CESifo, stadelis@berkeley.edu

1 Introduction

Bilateral bargaining is one of the oldest and most common forms of trade. Nations negotiate trade deals, arms control, and climate change mitigation; legislators engage in horse-trading to build coalitions and pass legislation; business people haggle over contracts from corporate acquisitions to labor agreements; lawyers wrangle settlements both civil and criminal, and private individuals bargain over wages, real estate, and the allocation of household chores. Bargaining determines the allocation of surplus in these settings, as well as the likelihood of breakdown—the latter with real economic and human costs. Therefore, understanding how people bargain, and the institutions, norms, and practices that affect bargaining outcomes, is a question of first-order importance.

Over the past sixty years, a large literature in economics has examined various aspects of bargaining in theory and in laboratory experiments. The theoretical literature typically assumes a particular information structure and extensive form of the game, which can sometimes be replicated in a controlled laboratory setting. Bargaining in real-world settings, however, tends to be less structured, and little evidence has been presented about how people bargain in the field and how prices actually form in real-world negotiations. Indeed, Fudenberg et al. (1985) explained that the "thorny issue" arising in much of the bargaining literature is that the researcher does not actually know the extensive form of real-world bargaining scenarios. For example, a street vendor bargaining over price might state an offer, watch the facial reaction of the buyer, and immediately state a lower price without waiting for a spoken response by the buyer. It is unclear whether this situation should be modeled with alternating offers, one-sided offers, a concession game, or any number of possible options.

The advent of online marketplaces provides a new opportunity to study negotiations in a real-world setting where the extensive form of the game is similar to those studied in the theoretical and experimental literature, and where the data collected is on a massive scale. In this paper, we utilize data from over 88 million listings on the eBay.com "Best Offer" platform, where sellers offer items at a listed price and invite buyers to engage in alternating, sequential-offer bargaining, very much in the spirit of Rubinstein (1982). In over 25 million of these listings, buyers chose to make an offer, initiating the alternating-offer game. Within this setting, we document a variety of facts on how bargaining proceeds and how prices form, and the forces that play a role in this process.

We find evidence consistent with the most salient predictions of economic theory, and at the same time present robust patterns that suggest that behavioral factors based on reciprocal and equitable norms play a significant role in bargaining outcomes.

While more widely known for its auctions and fixed price listings, eBay has offered sales through alternating-offer bargaining for over a decade. Our data come from eBay's Best Offer platform, through which almost ten percent of eBay transaction volume occurs as buyer-seller pairs engage in bargaining. Given the sheer volume of trade on eBay and the simple extensive form of the game, the Best Offer platform provides a useful setting for studying the determination of agreed-upon prices in sequential bargaining situations. The bargaining in this setting is only over a single dimension (price), making it more straightforward to analyze than many other bargaining settings (such as procurement contracts; Bajari et al. 2009), while still yielding the benefit of being a real-world setting. Furthermore, the data allows us to link buyers and sellers over time. Our dataset is, to our knowledge, the largest offer-level negotiations dataset to be analyzed in the literature.[1]

Section 2 describes background on the Best Offer platform and introduces our dataset. Section 3 then documents how patterns observed in the data relate to a variety of game-theoretic models of bargaining. We provide a breakdown of how bargaining sequences unfold in practice and the frequency with which different responses and outcomes occur. We find that there are often few back-and-forth offers between a given bargaining pair, which is predicted by complete-information, common-priors models of bargaining, such as the classical Rubinstein (1982) model. Bargaining also frequently ends in disagreement early on, consistent with the incomplete information model of Perry (1986). Some interactions involve a delay and end in agreement, consistent with some models of heterogeneous priors (Yildiz 2003) or incomplete information (Rubinstein 1985; Admati and Perry 1987). However, a number of sequences end in *disagreement* after a delay, a feature which is absent in nearly all existing bargaining models, with Cramton (1992) being the only exception we are aware of.

[1] In cooperation with eBay, we have anonymized the dataset and have made it publicly available for research purposes. The data can be accessed at http://www.nber.org/data/bargaining.html or by contacting the authors. We hope that it will further fuel the recent surge of empirical work studying bargaining in economics and stimulate additional work in the area, both empirical and theoretical.

In Section 4 we examine several conventional drivers of bargaining outcomes. Bargaining differs when players bargain over expensive versus inexpensive products in a way that is consistent with fixed costs of bargaining playing a role, as opposed to the more commonly modelled discount-factor approach. We then examine several forms of bargaining strength. We find that buyers who are more patient (as measured by their ex-post choice of shipping speed) tend to obtain lower prices in the bargaining. Buyers who are more experienced in bargaining on this platform (as measured by the number of previous Best Offer negotiations the buyer has participated in) also tend to achieve lower final prices, and experienced sellers achieve higher final prices. These results are consistent with common models of bargaining in which patience or other measures of a player's bargaining power affect outcomes (Rubinstein 1982, 1985; Watson 1998); they are also consistent with laboratory evidence (Rapoport et al. 1995) and survey data (Scott Morton et al. 2011), but, to our knowledge, have not been previously confirmed with data from actual bargaining outcomes.[2]

In Section 5 we document some patterns unexplained by existing bargaining theories, which exhibit flavors of reciprocity. First, we find that bargaining offers tend to change gradually over the course of the bargaining interaction. This is in contrast to a number of existing bargaining models, in which any delay in bargaining is war-of-attrition-like delay: No information is revealed until one party concedes everything, completely revealing her valuation to the opponent. In contrast, the behavior we observe is consistent with a gradual revelation of information and a gradual concession of bargaining positions. We demonstrate furthermore that this gradualism is *reciprocal*: the opponent responds to stubbornness by not conceding, and to concession by reciprocating with more concession.

We then show that a player often makes offers lying halfway between the player's own previous offer and the opponent's current offer. We further demonstrate that such "split-the-difference" offers have a higher likelihood of being accepted—higher even than some offers that would be even more favorable in money terms for the accepting party. Such behavior is not consistent with any existing theory of rational behavior, but it may be consistent with behavioral models; similar split-the-different behavior is

[2]One study documenting causal evidence on delay in bargaining from field data (but not actual offer data) is Ambrus et al. (2018), studying delay induced by travel times of Spanish ransom teams negotiating with North African pirates in the 1600s.

discussed throughout the experimental and theoretical behavioral bargaining literature, in which market participants may care about fairness, and often favor a split-the-difference strategy (Roth and Malouf 1979; Roth 1985; Binmore et al. 1985; Bolton 1991; Bolton and Ockenfels 2000; Charness and Rabin 2002; Andreoni and Bernheim 2009).

Our paper is related to a growing literature studying negotiated price settings. Many such papers have only data on final negotiated prices and only for cases in which trade occurs, such as in Crawford and Yurukoglu (2012). In contrast, we observe all back-and-forth offers, even for bargaining interactions that failed to reach an agreement. Similar datasets, although smaller than ours in size and scope, are analyzed in Keniston (2011),Bagwell et al. (2017), Hernandez-Arenaz and Iriberri (2018), Larsen and Zhang (2018), Larsen (2018), and Backus et al. (2019). Several papers provide tests of implications of bargaining theory, as ours does, such as Scott Morton et al. (2011) and Grennan and Swanson (2016).

The large scale of our data and the variation across several measures of heterogeneity help paint a useful and comprehensive picture of sequential bargaining in the real world that adds great detail to the existing literature. The patterns we uncover confirm some of the most basic insights of bargaining theory, yet at the same time reveal behaviors that are not explained by conventional theoretical approaches. As such, the patterns we uncover suggest that further developments in the theory of bargaining are warranted, especially those that take more seriously aspects of reciprocity and fairness.

2 eBay's Best Offer Mechanism: Facts and Data

eBay is one of the world's largest online marketplace for consumer-to-consumer transactions. It began in 1995 using second-price-like auctions as the sole format for transacting on its platform. The site eventually allowed users the option of selling goods through a single posted fixed-price. In 2005, the site began to allow sellers to sell through an alternating-offer protocol referred to as "Best Offer." This feature can be enabled (at no cost) by the seller at the creation of the listing, and is only available for fixed-price listings—there is no equivalent mechanism for auctions.

The Best Offer platform is currently a fast-growing sales format on eBay. Figure 1 shows the growth of this format relative to auctions and fixed price listings over the

Figure 1: Growth of Best Offer

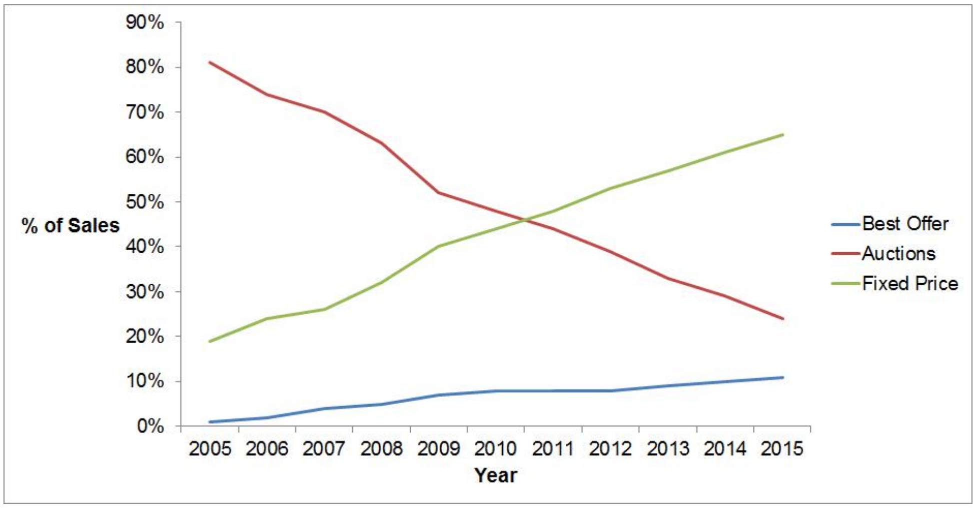

Notes: This figure depicts the percentage of Gross Market Value made up by three mechanisms on the eBay platform—Best Offer, auctions, and fixed price listings—from 2005 to 2015, computed from internal eBay data. The tabulation for fixed price listings includes Best Offer sales. The tabulation for Best Offer includes only listings that were bargained; it does not include Best Offer-enabled listings that sold at the listing price.

past ten years. In 2005, when the format was first rolled out, only a tiny fraction of listings were Best Offer listings, less than 1% of all eBay transactions occurred through a buyer actually placing an offer (rather than accepting the Buy It Now price). By 2012, that fraction had grown to just under 9%.

Goods offered for sale under the Best Offer format are listed as "accepts Best Offer" in eBay search results.[3] Throughout, we refer to these postings as Best Offer listings. A buyer viewing a Best Offer listing sees similar information to a buyer viewing a fixed price listing (referred to as a "Buy It Now" (BIN) listing), including the auction title, seller id and feedback score, at least one picture of the item, and any other information about the item that the seller decides to display.[4] The buyer sees the BIN price, as in a standard fixed price listing, but also sees an additional option, a button labeled "Make Offer", as illustrated in Figure 2. Selecting the Make Offer button allows the buyer to send an offer to the seller. As such, we treat the BIN price (equivalently, "listing price") as the seller's first offer to any buyer who wished to bargain.

[3]Potential buyers may filter search results to display only those listings that accept offers.

[4]Throughout, we will use the term "buyer" to refer to the user interested in potentially buying the item whether or not the transaction actually occurs.

Figure 2: Best Offer User Interface

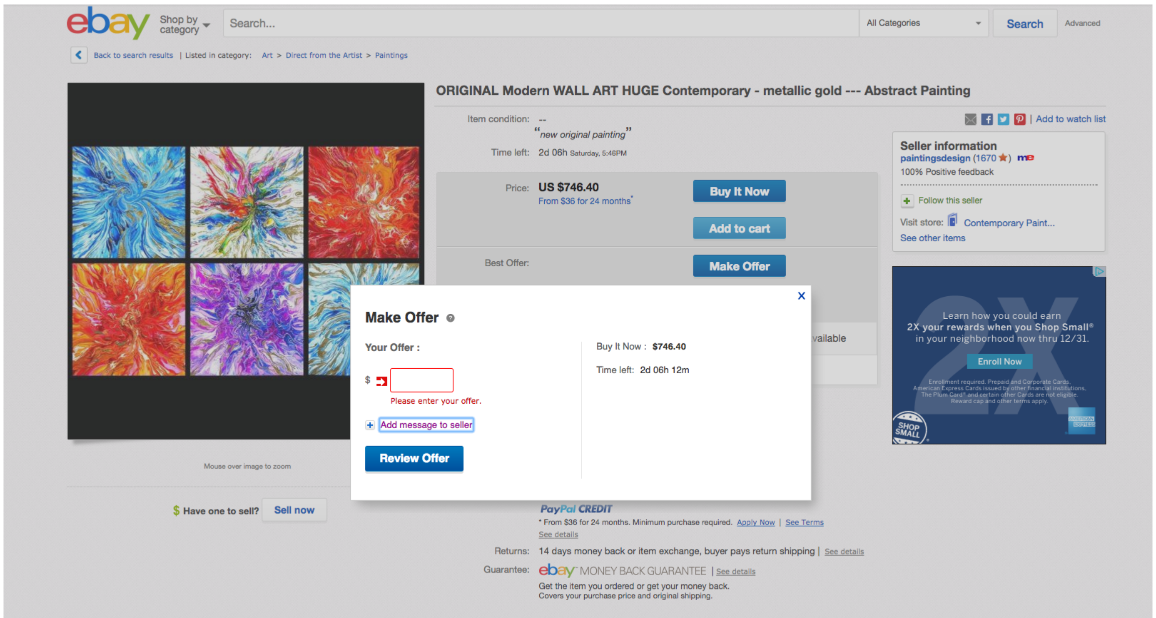

Notes: This figure depicts the "view item" page for a listing with Best Offer enabled. The potential buyer may click on "Buy it Now" to purchase the painting at the listed price of $746.40—or they may click on "Make Offer" and be prompted to propose a price.

Upon receiving this offer, the seller may accept the offer, make a counteroffer, or decline the offer (without making a counteroffer in return).[5] If the seller makes a counteroffer, the buyer then can accept, decline, or counter in response. Play continues until either party accepts or until the buyer declines. If the seller declines, the buyer may still respond with a counteroffer or can, at any time, purchase at the BIN price. Each party is limited to three offers (not including the listing price), and each offer expires 48 hours after being placed.[6] We will refer to a sequence of back-and-forth offers —i.e. a given buyer and seller pair bargaining over a given item — as a *thread*.

To form our primary dataset, we obtain internal eBay data from all Best-Offer-enabled, single-unit listings created in May 31, 2012 – June 1, 2013 from the US eBay site.

[5]As a time-saving device for sellers, the platform offers sellers the option to specify an "auto-accept" price which is unobserved to buyers and which, if exceeded by the buyer's offer, will result in the platform accepting the offer on behalf of the seller. Sellers can similarly specify an "auto-decline" price.

[6]In 2017 this maximum limit was changed to five offers. In our sample, approximately 1.1 percent of interactions reach the binding limit, and most of these fail; see Figure 3 below. The limit was extended in an attempt to encourage these negotiators to succeed. The time frame of our data does not permit us to evaluate this policy change.

This consists of over 90 million Best-Offer-enabled listings. This dataset, anonymized to remove all identifiable information, consistutes the dataset that we have arranged to have released publicly for research purposes. For the primary analysis in this paper, we also restrict attention to listings with BIN prices between \$.99 and \$1,000.00, and eliminate listings with apparent data errors (e.g., cases where we could not locate the original offer corresponding to a counteroffer). Details on our sample construction criteria appear in Appendix B. Our final dataset analyzed in this paper contains approximately 88.4 million listings. Of these, 25.4 million received a bargaining offer. Thus, there are 25.4 million bargaining *threads* (defined as a listing-buyer pair). These bargaining threads involve 1.2 million buyers and 4.7 million sellers.

Table 1 presents mean descriptive statistics for our sample. The first column corresponds the full sample. The next seven columns contain the same variable means for separate product categories in the data: collectibles, electronics, fashion, media, toys, business and industrial (B&I), and other. The final column contains a subsample outside of our main sample, corresponding to business and industrial equipment with list prices greater than \$1,000 (such as construction machinery).

In the main sample, the average list price (BIN) is \$95, and the average sale price is 83% of the list price. Sales include BIN choices as well—conditional on bargaining occurring, the average sale prices comes down to 73% of the list price. We note that almost 80% of listings never receive an offer and do not sell. 54.8% of listings are for used goods, and 26.3% of listings have the BIN price revised at some point by the seller during the listing life.

Table 1: Descriptive Statistics

	Full sample	Collectibles	Electronics	Fashion	Media	Toys	Business	Others	B&I expensive
Listing-Level Data									
Listing Price	94.6	75.7	146	122	31.6	73.7	200	126	5588
Used	0.548	0.56	0.639	0.466	0.724	0.503	0.613	0.413	0.763
Revised	0.263	0.236	0.242	0.31	0.296	0.238	0.208	0.269	0.166
Sold	0.215	0.177	0.373	0.222	0.148	0.287	0.213	0.273	0.126
Sold by Best Offer	0.132	0.115	0.194	0.15	0.0714	0.164	0.128	0.151	0.0978
Received an Offer	0.206	0.176	0.345	0.229	0.102	0.268	0.186	0.241	0.23
Sale Price	69.7	54.8	124	65.9	22.6	62.5	143	77.4	2900
Sale Price / List Price	0.832	0.805	0.89	0.814	0.855	0.858	0.827	0.858	0.795
Bargained Price	74.1	60.7	135	68.9	25.8	68.2	145	81.7	2784
Bargained Price / List Price	0.727	0.701	0.791	0.728	0.703	0.753	0.714	0.747	0.736
No. Listings	88,388,220	34,809,958	6,797,133	23,416,772	9,350,399	7,760,166	3,060,107	3,193,685	560,941
Seller-Level Data									
Feedpack Postitive Percent	99.4	99.5	99.3	99.4	99.5	99.5	99.4	99.4	99.2
No. Listings	73.8	87.7	16.1	52	61.9	27	28.4	14.1	15.1
No. Sales	15.9	15.5	6.01	11.5	9.16	7.74	6.05	3.86	1.91
No. Sales by Best Offer	9.72	10.1	3.13	7.78	4.42	4.41	3.64	2.14	1.48
No. Sellers	1,197,419	397,094	421,169	450,347	151,076	287,817	107,832	225,834	37,163
Buyer-Level Data									
No. Bargaining Threads	5.12	8.22	5.99	5.77	7.68	6.57	6.64	7.05	9.42
No. Offers	8.48	13.5	10.1	9.54	12.2	10.9	10.8	11.5	15.8
No. Purchases	3.21	3.6	1.88	2.37	2.06	2.01	2.16	1.4	1.52
No. Bargained Purchases	2.47	2.76	1.38	1.94	1.4	1.52	1.64	1.07	1.33
No. Buyers	4,701,455	1,446,880	954,628	1,802,583	474,710	835,191	239,031	450,873	41,091
Thread-Level Data									
No. Offers	1.66	1.64	1.7	1.66	1.54	1.68	1.61	1.66	1.79
No. Offers conditional on Agreement	1.48	1.46	1.58	1.46	1.37	1.53	1.45	1.48	1.70
Agreement Reached	0.454	0.491	0.317	0.474	0.596	0.423	0.543	0.451	0.26
First Buyer Offer	86.6	67.6	150	79	26.9	77.1	138	92.2	3377
First Buyer Offer / List Price	0.608	0.575	0.66	0.607	0.602	0.624	0.596	0.625	0.545
No. Threads	25,458,516	8,108,054	4,128,702	7,349,990	1,113,664	2,976,836	719,139	1,062,131	207,550

Notes: This table presents summary statistics for the main dataset. Note that indicator "Used" (for used vs. new status of item) is only available for 60,709,655 listings, and feedback variables are only available for 1,145,426 sellers. See text for a discussion of exclusion criteria and, in particular, Appendix A. Columns 2–8 contain summary statistics for specific subsamples. Column 9 contains summary statistics for business and industrial (B&I) items with list prices over $1,000.

The second and third panels of Table 1 include detailed information on market participants. On average, sellers have a 99.4% positive feedback score. While there are many one-time sellers, the market is skewed towards experienced sellers: the average number of listings per seller is 74, sixteen of which sell, ten of which sell through bargaining. Most of the sales in our dataset are made by a relatively small fraction of the sellers. The population of buyers is skewed, but less so: on average, buyers in our sample are observed in 5 bargaining threads, make 8 offers, and purchase 3 items (2.5 of these coming through bargaining). Finally, at the thread level, Table 1 shows that most bargaining threads are short (only 1.6 offers, on average, where the first offer is always made by the buyer), and surprisingly likely to be successful. On average, buyers offer $86.60, which represents 61% of the list price. Bargaining is ultimately successful 45% of the time.[7]

The category-specific statistics in Table 1, columns 2–8, demonstrate that the majority of listings are for items that could reasonably be characterized as being idiosyncratic or one-of-a-kind inventory, such as collectibles or fashion. Categories with more well-defined, frequently sold products, such as media products or electronics, make up a smaller fraction of the data. Interestingly, Table 1 reveals that collectibles are less likely to receive offers or sell through bargaining than are electronics, but when they do sell through bargaining they do so at a smaller fraction of the list price (the ratio of the bargained price to the list price is 0.70 in the former and 0.79 in the latter). The seller-level and buyer-level panels demonstrate that buyers and sellers in the collectibles and fashion categories are more experienced bargainers than in other categories: the average collectibles seller has 10.1 Best Offer sales and the average fashion seller has 7.78, while in the remainder of the categories these numbers range from 2.14–4.42. Similarly, the average buyer in collectibles and fashion purchases 2.76 and 1.94 items (respectively) through Best Offer, while in the remainder of categories these numbers range from 1.07–1.64. These numbers suggest that bargaining may be important particularly for frequent buyers and sellers and particularly in those categories with more idiosyncratic inventory. Still, a number of statistics are remarkably stable across categories, such as the number of offers per thread (1.54–1.7) and the ratio of the first buyer offer to the list price (.575–0.660).

[7] Appendix D contains an anlysis the variance in prices, sales rates, and number of offers, describing how much of this variance is driven by the players (sellers and buyers) vs. the products.

Table 2: Listings Receiving an Offer vs. Not

	Full Sample	Received at least one offer	Never received any offer
Listing-Level Data			
Listing Price	94.6	119	88.3
Used	0.548	0.558	0.545
Revised	0.263	0.263	0.263
Sold	0.215	0.739	0.0794
Sale Price	69.7	78.7	47.9
No. Listings	88,388,220	18,216,576	70,171,644
Seller-Level Data			
Feedpack Postitive Percent	99.4	99.4	99.4
No. Listings	73.8	21.1	71.8
No. Sales	15.9	15.6	5.7
No. Sellers	1,197,419	861,788	976,698
Buyer-Level Data			
No. Purchases	3.21	2.82	2.18
No. Buyers	4,701,455	4,481,430	1,133,818

Notes: This table presents summary statistics for the listings in our data that ever receive an offer compared to those listings that never receive an offer.

The business and industrial category provides an interesting opportunity to study how bargaining differs for big-ticket vs. small-ticket items. Business and industrial contains a number of lower-priced items, such as machine parts, as well as large, expensive construction equipment items such as a backhoe or skid steer loader. The final column of Table 1 displays descriptive statistics for a separate subsample in the business category for expensive items—those with list prices greater than $1,000—which are excluded from our primary sample. The average list price in this expensive sample is significantly higher ($5,588 as opposed to $200) and these items are more likely to be used than in the main sample (0.763 vs. 0.613). The sellers and buyers of the expensive items are less experienced eBay bargainers than those in the main sample, and the bargaining threads are less likely to end in agreement. However, a number of statistics are remarkably similar across high vs. low-price items, including the ratio of the bargained price to the list price, the ratio of the first offer to the list price, and the number of offers made.

Much of our analysis throughout the remainder of the paper will focus on the subsample of listings in which bargaining takes place. These listings, and the buyers and seller involved in these listings, may differ from those in which no bargaining offer is ever observed. In the top panel of Table 2 we compare the full sample and these two subsamples of listings. We find that listings that did receive at least one offer were much more likely to sell, and at a higher price (note that this finding would not have been a foregone conclusion: listings that receive no offers can still sell through the BIN option). We find that listings bargained over and those not bargained over are equally likely to be used items and to have the BIN price revised at some point. In the bottom two panels of Table 2 we compare the buyers and seller involved in listings with no bargaining and those with bargaining. Sellers involved in bargaining post fewer listings on average but sell more listings, but have similar feedback ratings. Buyers involved in bargaining tend to make more purchases than those who are not.[8]

3 Bargaining Theories and Empirical Evidence

In this section, we walk through a number of existing theoretical models of bargaining, examining their empirical implications and studying how much of the observed data each model might be able to to explain. Clearly these models were intended as abstractions from reality, and as such cannot be expected to explain all features of real-world negotiations. Here we use these models simply as a framework for highlighting features of real-world bargaining that existing theory can and cannot explain well. We focus our discussion in this section primarily on three features that have played a prominent role in motivating existing theory: (i) *whether* agreement or disagreement occurs, (ii) *when* agreement or disagreement occurs (immediately or after a delay), and (iii) *how* the final negotiated price is reached (suddenly or gradually).

Points (i) and (ii) relate to how bargaining ends. To visualize how bargaining ends in the actual data, we display a tree in Figure 3 representing the extensive form of the bargaining game. Square boxes represent the identity of the player (B = buyer, S = seller). At the right of each box, we display the number of observations that reach

[8]Note that, unlike the top panel of Table 2, the comparison of sellers involved in listings that never receive an offer and sellers involved in listings that do receive an offer is not a comparison of mutually exclusive samples. This is also true for the buyer comparison in the bottom panel.

the node. Below each node are edges representing the player's decision to make an initial offer (O), accept (A), decline (D), or counter (C). Each edge shows the percent of observations passing through that edge corresponding to a given action being chosen.[9] We will refer to these percentages below as we explore what fraction of the data existing models could potentially explain. Throughout the discussion below, we will attempt to be generous in attributing what fraction of observations can be explained by a given model. In particular, each observation in the data is a sequence of offers corresponding to some game-tree-path from Figure 3; if a given model can generate this path—and in some cases a particular rate of concession of players' counteroffers—we will state that this model could have plausibly generated this observation.

We emphasize here that the discussion in this section is not meant to suggest that any of these theoretical models are "wrong" in any sense—each is entirely accurate if its corresponding assumptions are satisfied. Rather, by pointing out empirical regularities that some models fail to generate, this discussion highlights that the assumptions of these models (such as complete information of one or both parties) may not be satisfied in real-life settings.

3.1 Immediate Agreement: Nash (1950) and Rubinstein (1982)

Perhaps the two models of bargaining that most quickly come to an economist's mind are those of Nash (1950) and Rubinstein (1982). The model of Rubinstein (1982) consists of two players who sequentially alternate their offers. Each player has some cost of bargaining, such as a per-period discount factor or a per-offer cost applied to each time period of the game, and players' have complete information about their opponents' valuations. The unique subgame perfect equilibrium (SPNE) of this game is related to the axiomatic, cooperative solution proposed in Nash (1950), in which players choose a division of surplus that maximizes a weighted, joint payoff of the two bargaining parties, with each party's weight depending on a "bargaining power" parameter. When this bargaining power is given by players' discount factors in Rubinstein's model, the Nash solution corresponds precisely to the unique SPNE of the Rubinstein game.

[9] For the sake of visual clarity, Figure 3 does not display the buyer's option to buy at the BIN price later in the bargaining sequence, which is always an option buyers have available.

Figure 3: Bargaining Sequence Patterns

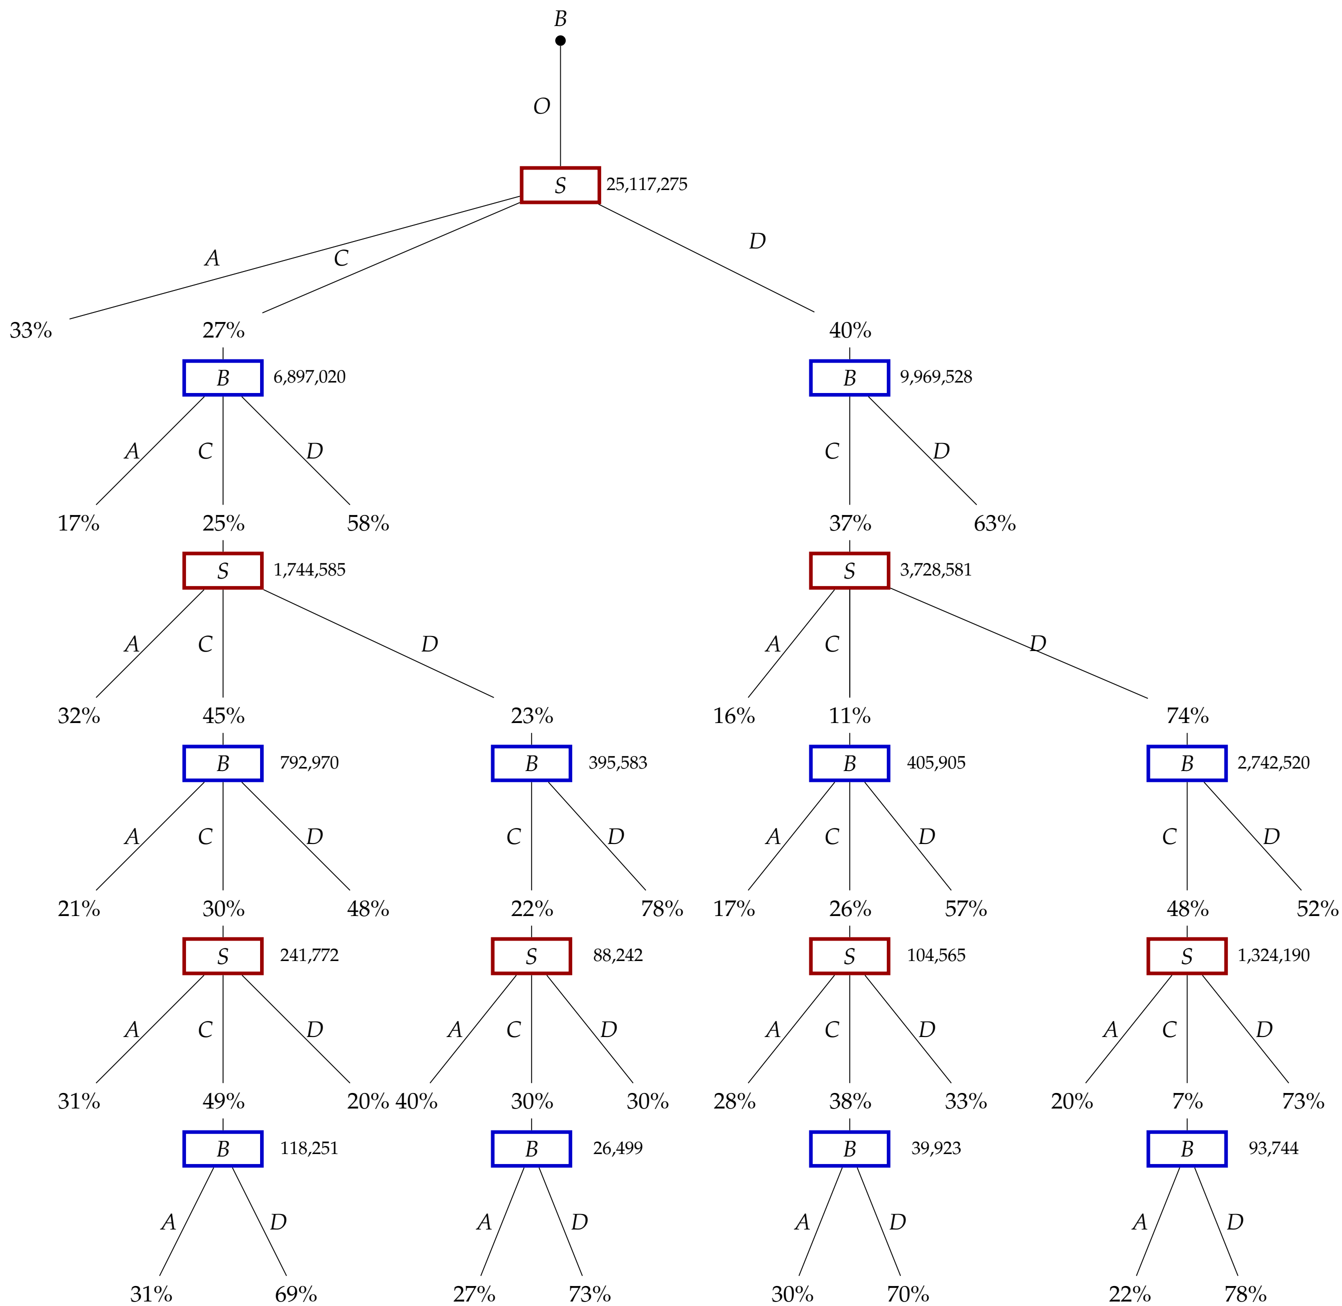

Notes: This figure summarizes the offer-level data in terms of the “game tree” of bargaining. See text for detailed discussion.

These two models have several immediate empirical restrictions. First, all bargaining interactions should end in agreement. Second, the player who moves first makes an offer that is immediately accepted by the second mover.[10] These two restrictions imply that, in order for an observed interaction to be rationalized by the Rubinstein model, it must be the case that the game ends with the first offer being accepted immediately by the opponent. As shown in Figure 3, immediate agreement occurs in 33% of our observations. The remaining 67% of observations are inconsistent with Rubinstein's strategic model or Nash's cooperative model. Below we discuss a number of theoretical approaches that could explain some of these outcomes.

3.2 Immediate Disagreement: Perry (1986)

A salient feature of real-world bargaining is that some negotiations end in impasse, and this cannot be explained by the complete-information models of Nash (1950) and Rubinstein (1982). A number of studies explain such impasse by relaxing the complete information assumption and incorporating incomplete information into the sequential bargaining game or a mechanism design framework. For example, incomplete information is the key to the seminal theorem of Myerson and Satterthwaite (1983). An important contribution to modeling an extensive-form, alternating-offer bargaining game with incomplete information is that of Perry (1986), in which both the buyer and seller have *private* valuations that are not commonly known to the parties. Buyers face a cost (common to all buyers) of making each offer and the seller faces a per-offer cost common to all sellers. The unique sequential equilibrium of this game is that the side with the lowest bargaining cost makes an offer and the other party accepts or rejects, but never makes a counteroffer. Figure 3 implies that this model can rationalize 58% of the observed bargaining sequences; this includes a large percentage that cannot be rationalized by the Rubinstein (1982) or Nash (1950) models.[11] The Perry (1986) model cannot, however, rationalize cases in which multiple offers or delays occur in

[10] A third implication of the Rubinstein and Nash bargaining models is that a player who has more "bargaining power" (or lower bargaining costs) should obtain a better deal. Testing this implication is not simple, as bargaining power and bargaining costs are not tangible or well-defined objects, and are difficult to measure in data. We provide one approach to such an analysis in Section 4.2.

[11] This 58% includes the 33% of immediate-agreement observations as well as an additional 25% of observations—cases where the seller immediately declines (40%) followed by the buyer declining (63%): that is, 40% $\times$ 60% = 25%.

equilibrium. A number of other bargaining models of incomplete information share this feature as well, such as the "k-double auction" of Chatterjee and Samuelson (1983): bargaining can end in agreement or in disagreement, but it always ends immediately (in this case because the game is assumed to be static).

3.3 Delayed Agreement: Gul and Sonnenschein (1988) and Others

A large branch of the theoretical bargaining literature includes models in which parties may delay in reaching an agreement, but nonetheless the parties always agree. Models with this delayed-agreement feature differ widely in how delay is generated. Rubinstein (1985) and Grossman and Perry (1986) study alternating-offer cases in which there may be two offers in equilibrium before agreement occurs: If the offer of the player moving first is not accepted, the player who rejected that offer then makes an offer that is accepted immediately. Gul and Sonnenschein (1988) provide a model in which an informed buyer (with a private valuation) alternates offers with an uninformed seller (with a known valuation). In the equilibria they study, delay can occur before parties agree, but the authors demonstrate that a Coase Conjecture result (Gul et al. 1986) holds: as the time between offers decreases, agreement takes place immediately.[12] Figure 3 suggests that these delayed-agreement models can explain 37.6% of the observed bargaining sequences.[13]

Two separate models that can also generate delayed agreement in alternating-offer bargaining are Cramton (1992) and Abreu and Gul (2000). We will discuss these below, as they also provide testable predictions of how offers themselves change over the course of the bargaining.

[12]The authors also demonstrate that the equilibria they study have the property that all unaccepted offers made by buyers in a given period of the game must be the same for all buyers. Furthermore, the class of equilibria they study nests that of Grossman and Perry (1986). There are a number of other theories of delayed agreement; too many to treat here. For example, Feinberg and Skrzypacz (2005) obtains delayed agreement in a model in which one party has private information about her valuation and the other has private information about his *beliefs* about the other party's valuation.

[13]This quantity corresponds to the 33% of immediate agreement observations and the 27% $\times$ 17% of observations in which the seller countered and the buyer immediately accepted.

3.4 Delayed Disagreement: Cramton (1992)

The models discussed above cannot generate *delayed disagreement*—cases where a buyer and seller exchange multiple offers and then walk away without trading. This is another salient feature of the data shown in Figure 3. The one model of alternating-offers of which we are aware that allows for such delayed disagreement is Cramton (1992), in which both parties have private information about their valuations and may discover at some point that there are no gains from trade, and will then walk away.[14] In the model, parties effectively play a war of attrition game: once a party makes an offer, the offer itself and the timing of the offer completely reveal the offering party's valuation. This model can potentially rationalize any cases in which, after an offer is made, the opposing party (i) immediately accepts or declines it or (ii) makes a counteroffer that is immediately accepted.

The empirical prevalence of these Cramton (1992) cases can be computed in a number of ways. First, if the BIN price is viewed as not fully revealing of the seller's valuation, and the buyer's first offer (and its timing) is viewed as arising from the Cramton separating equilibrium, then the seller should respond to this first buyer offer by immediately accepting or declining or by making an offer that is guaranteed to be accepted. Such cases comprise 62.8% of bargaining sequences in Figure 3.[15] If instead the first bargaining offer is viewed as uninformative to the seller (and simply as a signal that the buyer wishes to enter negotiations), then the first offer with the potential to fully reveal a player's valuation is the *seller's* first counteroffer. The buyer should respond to this by accepting or rejecting, or by countering at a price acceptable to the seller. This behavior is consistent with 83% of the observations in which the seller makes a counteroffer in response to the buyer's first offer (which occurs 27% of the time).[16] In either scenario, a significant fraction of bargaining sequences in the data are too long to be explained fully by the Cramton (1992) equilibrium. Moreover, the war-of-attrition nature of the Cramton model does not match the actual protocol used in the data; the

[14]The model of Cramton (1992) is a two-sided-offer, continuous-types extension of the one-sided offer game of Admati and Perry (1987).

[15]This quantity comes from 33% immediate agreement, 40% $\times$ 63% immediate disagreement, as in the Perry (1986) model, plus an additional 27% $\times$ 17% of cases where the seller counters at a price that the buyer then accepts.

[16]This 83% figure can be seen in Figure 3, following the seller's first counteroffer, as the sum of 17% + 58% + 25% $\times$ 32%.

protocol dictates whose turn is next, which in and of itself poses some difficulty in using this model to explain patterns of behavior in our game.

3.5 What Theory Struggles to Explain: Gradual Offers

A salient feature of the data that existing theories fail to replicate is that counteroffers tend to gradually be more and more favorable to the opposing party with each subsequent offer. This is illustrated in Figure 4. In each panel, the vertical axis shows the *average* amount of the offer and, on the horizontal axis, the period of the game in which the offer is made, with the $t = 0$ offer representing the list (BIN) price. Panels on the left include sequences ending in agreement and panels on the right include those ending in disagreement. We analyze separately those sequences that ended in period 6, where the seller declined (and the buyer not taking any further action) or the seller accepting, and those that ended in periods 7, with the buyer accepting or declining. Each panel also displays the average change in offer price the seller makes from one offer to the next, averaged across all periods of the game, and similarly for the buyer.[17]

This gradualism of offers seems quite intuitive to anyone who has engaged in bargaining in practice—it is exactly how one would expect bargaining offers/counteroffers to evolve over the course of negotiations. However, theoretical models generating this type of pattern are almost non-existent. As highlighted above, most models can rationalize only immediate disagreement or immediate agreement, or, in the cases of Grossman and Perry (1986) or Gul and Sonnenschein (1988), one offer at most by each party. The model of Cramton (1992) can accommodate many offers by each party in equilibrium, but only if all offers (other than the last offer made by each party) are completely uninformative, non-serious offers; offers cannot gradually become more favorable to the opposing party.[18] Indeed, in the model of Cramton (1992), when a player concedes to the opposing player, the concession must be a precipitous jump.

[17] In creating these figures, we treat periods in which a seller declines a buyer offer and the buyer follows up with an additional counteroffer as periods in which the seller did not decline but rather *countered* at her most recent offer (or at the list price in the case where the list price is the most recent price stated by the seller).

[18] As highlighted above, in the equilibrium of Cramton (1992), the timing of offers is endogenous. This feature can be viewed as literally allowing a player to signal her valuation when she first make an offer, or instead allowing the player to delay through making non-serious offers whenever it is her turn, and then finally making a serious offer and signaling her valuation through the cumulative total time passed since the game began until the serious offer is made.

Figure 4: Price Convergence Over the Duration of Bargaining ($t = 6, 7$)

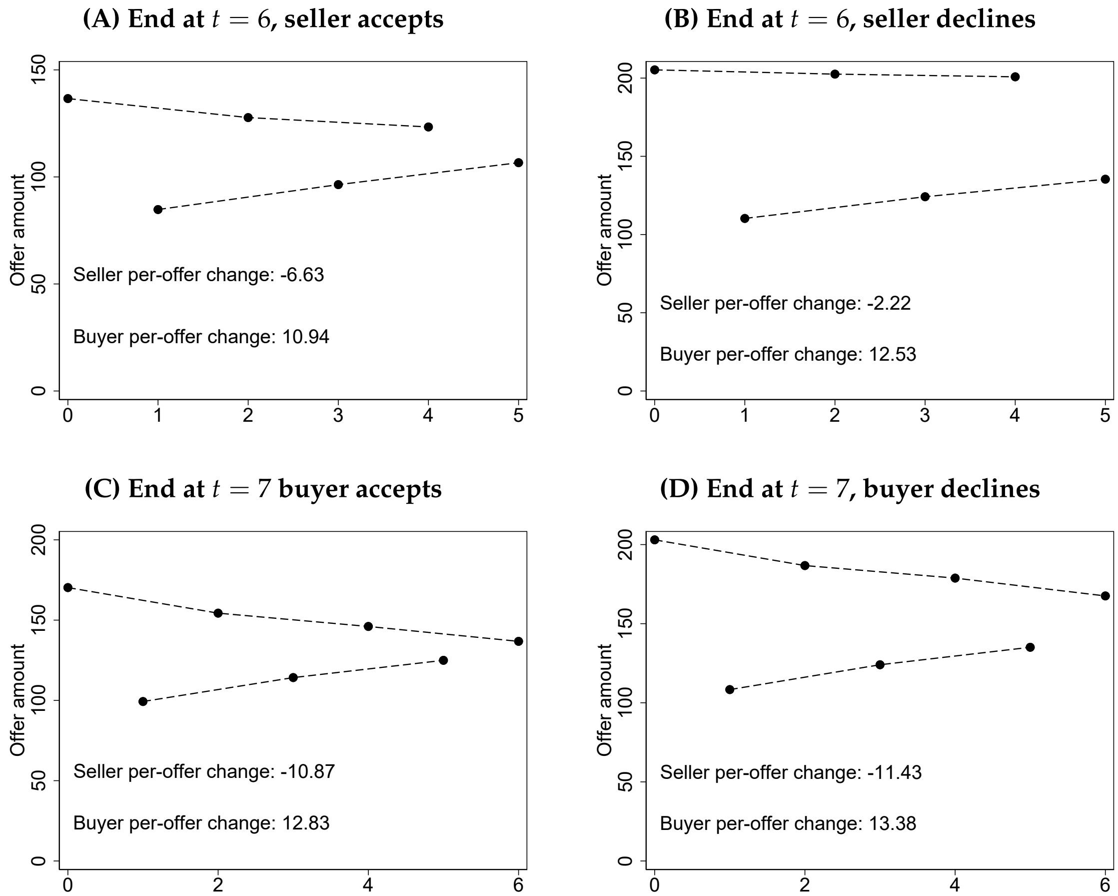

Notes: Figure displays the average first offer (BIN price), average second offer, etc. for bargaining sequences that ended in six (Panels A and B) or seven (Panels C and D) periods. Panels on the left ended in acceptance and panels on the right ended in decline.

Other models have a similar war-of-attrition flavor to the Cramton (1992) model and can thus explain sudden, but not gradual, offer changes. One example includes Abreu and Gul (2000). Their model applies reputational game theory concepts (Kreps et al. 1982) to bargaining, as suggested by Myerson (1991). In their model, there is some probability that an opponent is a "crazy" or "obstinate" type who will never budge from her initial offer. Rational types find it profitable to mimic obstinate types until the game ends, at which point they will concede. This model yields sudden offer changes when a rational player does concede, but no gradual changes along the way.[19] We are

[19] A gradual change in offers is the equilibrium outcome of several models of *one-sided* bargaining, where only one party gets to make offers to the opposing party, and the opposing party only gets to

aware of only two previous model generating gradually changing offers on both sides: Compte and Jehiel (2004) and Abreu and Pearce (2003). In the first paper, the authors model a setting in which both parties have complete information, but, by making offers, parties can influence the outside-option payoff to the opposing party. In the second paper, the authors behavioral types who, for exogenous reasons, concede differently than rational agents. Both models share an unrealistic feature of many of the above models in that they only allows for equilibrium *agreement*, not disagreement. We explore the gradualism behavior that we observe in more detail in Section 5.1.

4 Conventional Drivers of Bargaining Outcomes

In this section we analyze three features that play important roles and drive outcomes in a number of bargaining models: costs of bargaining, bargaining power, and players' outside options. Each of these features can be related to the others and are not necessarily distinct in their empirical implications. We explore a fourth dimension of analysis and examine the impact of reducing adverse selection on bargaining outcomes.

4.1 Bargaining Costs

Rubinstein (1982) proposed two models of bargaining costs: one in which the surplus at stake is discounted exponentially, as if the primary cost of bargaining were delayed consumption, and a second, in which there is a fixed cost of making offers. In the first case bargaining costs scale up with the value of the transaction, while in the latter they are fixed. Many subsequent bargaining models have also adopted one or the other (or both) of these types of costs (e.g. Cramton 1991). Because these types of bargaining costs differ in how they relate to the value of the transaction, in this section we search for evidence of such costs by examining how bargaining outcomes differ at different levels of the listing price.

accept or reject these offers. For example, both Fudenberg et al. (1985) and Gul et al. (1986) allow for a seller (with no private valuation for the good) to screen a buyer (who does have a privately known valuation) by making successively decreasing offers. This generates gradual offers, but only through the constraint that only one party is allowed to make offers, which is perhaps not a realistic constraint for many bargaining settings (and clearly not for ours). Other papers with similar assumptions and gradualism include Deneckere and Liang (2006), Fuchs and Skrzypacz (2013), and Gerardi et al. (2014).

Figures 5 and 6 present smoothed plots of expected outcomes against the listing price for our sample. To construct these plots, we employed a stratified subsampling approach discussed in Appendix C. The distribution of listing prices is presented in Panel A of Figure 5, where we see that the vast majority of listings fall in the \$.99 to \$100 range. While average first offers are decreasing throughout the range (Panel B), bargained prices are initially rising and then fall (Panel C), and the slope of the expected sale prices flips from negative to positive and back again (Panel D). Figure 6 provides some insight into this pattern. For very cheap items, more buyers exercise the BIN option and forego bargaining (Panels B and C). Moreover, sellers who do receive offers on cheaper items tend to accept them immediately (Panel D).

We interpret these outcomes as informative about the costs of bargaining. Assuming higher listing prices correspond to settings with a larger surplus on the table, our data is consistent with the existence of fixed costs of bargaining: when the listing price is greater and the amount of surplus to be negotiated is large, parties are more willing to engage in the back and forth of negotiation; and when the listing price is low and there is little surplus on the table, bargaining power tends to sit with whomever is making the current offer. This model of costs is also consistent with casual empiricism: bargaining in street markets is less frequent in developed economies with higher incomes—it is in some sense an inferior good—but bargaining remains prevalent among high-value transactions, e.g. salary negotiations, plea bargaining, terms of a merger, and trade deals; or even big-ticket consumer transactions, such as cars, large appliances, or homes. Fixed costs of bargaining are not the only possible explanation for the findings; another possible explanation would be that the types of players or equilibrium of the game differs markedly between high- and low-listing price sequences.

These findings are especially important insofar as most testing of theoretical models of bargaining has been primarily done in the lab. Experimental work focuses, for reasons of feasibility, on low-stakes bargaining. If players behave differently when the stakes are high, whether because there are fixed costs of bargaining or because of some other reason, then this implies an important caveat to the external validity of those findings. It also highlights the importance of complementing this experimental testing in the lab with evidence from the field.

Figure 5: Bargaining Outcomes by Listing Price

Notes: Panel A depicts a histogram of the listing prices for the full sample of listings. The remaining panels depict LOWESS plots of the outcome variables in terms of the listing price. In Panel B the variable of interest is the mean first offer of bargaining threads; in Panel C it is the bargained price, conditional on sale *and* the buyer not executing the BIN option; and in Panel D we are interested in the sale price, conditional on sale.

4.2 Bargaining Power

An additional feature of many theoretical models of bargaining is that players with more bargaining power obtain a greater share of the surplus. This "bargaining power" is captured differently in different contexts. In some models, such as Rubinstein (1982) and Rubinstein (1985), bargaining power is explicitly represented by a player's patience (discount factor). In other bargaining models, in particular many recent models applied in empirical research in bargaining settings (e.g. Crawford and Yurukoglu 2012; Grennan 2013), bargaining power is instead a reduced-form feature of the model rather than an underlying primitive, with a direct correspondence to the share of the surplus the player would receive in a static Nash bargaining game, where both players agree to maximize the total surplus weighted by the bargaining power weights (see Binmore et al. 1986). In these models, bargaining power can represent concepts such as a bargaining party's negotiation skill or experience.

Here we use a simple, yet novel approach to identify buyers who may have a greater degree of patience than others. In particular, we identify patient buyers as those who, ex-post (after the bargaining ended), chose the slowest shipping option when multiple options were available. Namely, at checkout, a buyer often can choose between several shipping options, where faster shipping costs more that slower shipping. Hence, by revealed preference, buyers who choose a slower shipping method reveal that they are willing to wait rather than spend more money, and are thus more patient than buyers who opt for faster shipping at a higher price. We also construct a measure of experience for buyers and sellers using their accumulated number of previous bargaining threads participated in (summary statistics for this measure are shown in Table 1).

Table 3 shows the results of regressing the bargained outcome on our measures of buyer patience and both parties' experience. For this analysis, we rely on a subsample of the data for which we can compute a reference price for each good. (See Appendix B for a discussion of, and summary statistics for, that sample.) In this table, we treat new products (columns 1–3) separately from used products (column 4–6). Our dependent variable in each of the regressions in Table 3 is the final price from a bargaining transaction in which agreement occurred, divided by the reference price for that item.

We find that, for used products, more patient buyers tend to have lower final prices in bargaining, consistent with theoretical models. Column 4 demonstrates that buyers who selected the slowest shipping option obtained prices that were lower by 7.2 percentage points (of the reference price). Column 6 controls for both experience and patience, and finds a similar premium for patient buyers, who obtain prices that are 5.7 percentage points lower than than less-patient buyers. For new goods, the relationship between patience and prices is not statistically significant.

It is important to note that, while this evidence is consistent with a role for patience, these regressions may also simply be capturing an effect of willingness or ability to pay: buyers who are willing/able to pay less for the item may also be willing/able to pay less for fast shipping, and hence the better deal obtained by buyers whom

Table 3: Bargaining Power and Prices

	(1) Norm. Price	(2) Norm. Price	(3) Norm. Price	(4) Norm. Price	(5) Norm. Price	(6) Norm. Price
Slowest	-0.0238		0.0318	-0.0722***		-0.0566***
Shipping	(0.0322)		(0.0411)	(0.0128)		(0.0147)
Log Seller		0.0228***	0.0643***		0.0137***	-0.00227
Experience		(0.00214)	(0.00915)		(0.00105)	(0.00279)
Log Buyer		-0.00234	-0.00423		-0.0184***	-0.0242***
Experience		(0.00304)	(0.0137)		(0.00166)	(0.00418)
Constant	1.215***	0.893***	0.858***	1.224***	1.081***	1.260***
	(0.0226)	(0.0112)	(0.0543)	(0.00977)	(0.00748)	(0.0234)
Condition	New	New	New	Used	Used	Used
R^2	0.0000264	0.00154	0.00423	0.000515	0.00123	0.00128
N	20,200	79,176	12,079	58,661	211,286	37,020

Notes: This table presents results from regressions where the dependent variable is the normalized price (see text for a discussion of the construction of reference prices) and the regressors are buyer and seller attributes. Robust standard errors are presented in parentheses. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

we label as "patient" may actually be due to those buyers' low willingness/ability to pay. This highlights an important point: a player's patience, willingness/ability to pay, bargaining costs (from the previous section), or outside options (discussed in the following section) can all affect the strength of a player's bargaining position and the surplus the player obtains.

Table 3 also demonstrates that, for both used and new items, more experienced buyers tend to obtain lower prices, although the point estimates are only statistically significant (and are larger in magnitude) in the used goods sample. Column 5 suggests and that 1% change in the buyer experience measure is associated with a decrease in prices corresponding to 0.018 percent of the reference price. When controlling for patience as well (column 6), the estimate increases in magnitude to 0.024 percentage points. For seller experience in both new and used good markets we find a similarly intuitive result: sellers with more experience tend to obtain higher prices, with the price increase ranging from 1.4 to 6.4 percent of the reference price, although this is result disappears in column 6 when controlling for patience. We show additional evidence in Appendix E (Table E1) that buyers' and sellers' counteroffer behavior follows a similar pattern: players concede more in nearly every round of the game when facing a more experience opponent, and concede less when they are more experienced themselves.

4.3 Outside Options:

Another feature that can affect outcomes is the outside option of a negotiator. In our setting, the outside option of a seller is to exit and search for another player with whom to negotiate (or to leave the platform). In some cases, a player may even be engaged in bargaining simultaneously with multiple parties at the same time. In our data we see 14.3% of bargaining threads in which the seller has offers from multiple buyers that overlap in time; that is, the window of time in which the seller bargained with one buyer overlaps the window in which she bargains with another distinct buyer. One the buyer side, we see 5.1% of threads in which the buyer is bargaining with more than one seller of the same cataloged product at the same time. When a buyer fails to reach agreement with a seller, in 2.8% of threads the buyer trades with another seller of the same product within a day's time.[20]

The option of a player to end negotiations with one player and engage with another can affect prices and other bargaining outcomes. To examine this we run regressions of the final price (normalized by the reference price, as in Table 3) on several measures of competition and outside options, including the log of the number of buyers with whom the current seller is bargaining simultaneously for this item (labeled "competing buyers" in Table 4), the log of the number of sellers (selling the same cataloged product) with whom the current buyer is bargaining simultaneously (labeled "competing sellers"), and the log of the number of other listings available at the same time as the current listing offering the same product as the current listings (labeled "competing products"). In these regressions, we also include leaf-category fixed effects.

The results are displayed in Table 4. The first four columns display results for new items and the last four for used items. For new goods we do not detect any significant effects of competition and outside options on the final price except for a significant and positive coefficient on the number of competing products, suggesting that prices are actually higher when more listings of a particular product are available; this likely reflects an equilibrium response of supply. Among used items, we instead see a positive

[20]A related interesting statistic is the fraction of repeat interactions by the same buyer and seller pair. We find that 9% of buyer-seller pairs meet in at least two separate bargaining threads. Together, these repeat buyer-seller pairs constitute 23.5% of the interactions in the data. This feature is another interesting aspect that could be exploited in future research with our public dataset to study, for example, reputation-building and learning in bargaining.

Table 4: Competition, Outside Options, and Prices

	(1) Norm. Price	(2) Norm. Price	(3) Norm. Price	(4) Norm. Price	(5) Norm. Price	(6) Norm. Price	(7) Norm. Price	(8) Norm. Price
Log competing buyers	0.00168			-0.00250	0.0228***			0.0243***
	(0.00782)			(0.00816)	(0.00485)			(0.00492)
Log competing sellers		0.0216		-0.0125		-0.0600***		-0.0558***
		(0.0650)		(0.0628)		(0.00617)		(0.00639)
Log competing products			0.0391***	0.0395***			-0.00668***	-0.00523**
			(0.00758)	(0.00688)			(0.00179)	(0.00188)
Constant	1.022***	1.022***	0.990***	0.991***	1.127***	1.134***	1.140***	1.135***
	(0.00429)	(0.00388)	(0.00611)	(0.00582)	(0.00232)	(0.00237)	(0.00324)	(0.00316)
Condition	NEW	NEW	NEW	NEW	USED	USED	USED	USED
R^2	0.0266	0.0266	0.0270	0.0270	0.0445	0.0446	0.0445	0.0446
N	125359	125359	125359	125359	328033	328033	328033	328033

Notes: This table presents results from regressions where the dependent variable is the normalized price (see text for a discussion of the construction of reference prices) and the regressors are the (log of) the number of overlapping buyers competing on the same thread for a give product, the number of sellers offering a given product, and the number of listings of the same product live on the site at the same time as a given listing. Robust standard errors are presented in parentheses. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

and significant increase in price (0.02% of the reference price) when the number of competing buyers increases by 1% for a given listing. Conversely, we see a negative and significant drop of 0.06% when the number of competing sellers increases by 1%, and a drop of 0.006% when the number of competing products increases by 1%. These results for used goods are consistent with the role one would expect for market thickness and outside options of one side of the market or the other.

4.4 Adverse Selection

One aspect of bargaining that has not be modeled extensively in theory literature is adverse selection. We do not attempt to provide such a model here, but instead we present several facts that may help guide future theoretical and empirical work. First, in previous work studying online markets (eBay Motors), Lewis (2011) demonstrated that the number of photos included in a listing serves as a useful measure of the amount of information revealed about an item's quality. Here we examine how bargaining outcomes differ with the number of photos.

Figure 7: Time to First Offer by Number of Photos

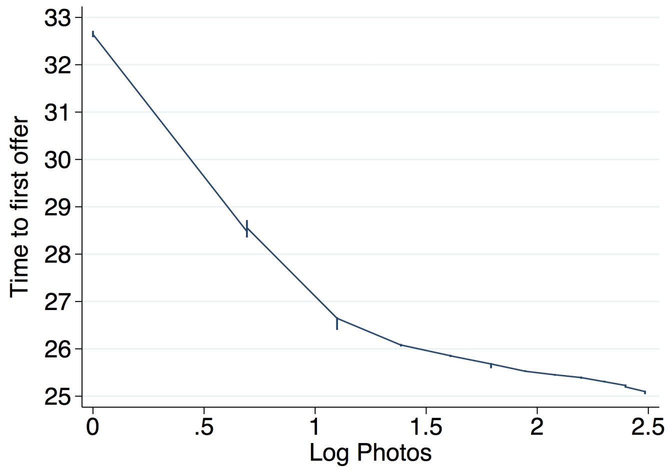

Notes: Figure displays the LOWESS estimates of the expected amount of time in days until the first offer arrives (after the listing is posted) conditional on the log number of photos.

First, in Figure 7, we examine how long it takes days for the first bargaining offer to arrive as a function of the log of the number of photos. We find that the larger the photo count the quicker the first offer arrives. This finding may be due to buyers being more willing to engage in bargaining when the amount of potential adverse selection is lower.

Table 5 presents regressions of bargaining outcomes on the log of the number of photos included with the listing. These regressions use our reference price sample, as in Tables 3 and 4. We find that the normalized price is higher and the likelihood of sale is smaller in both new and used goods when the number of photos is larger. This may suggest that sellers can extract more surplus in lower-adverse-selection settings. We also see that the number of bargaining offers per sequence is higher in higher-photo-count settings. In additional results (not shown) we also see that higher-photo-count listings are more likely to have received an offer and to have sold through Best Offer. These results are consistent with the idea that players may be more willing to negotiate when the potential for adverse selection is reduced.

Table 5: Bargaining Outcome and Number of Photos

	(1)	(2)	(3)	(4)	(5)	(6)
	Norm. Price	Norm. Price	Sold	Sold	No. offers	No. offers
Log photos	0.201***	0.0650***	-0.0125***	-0.0411***	0.0540***	0.0645***
	(0.00786)	(0.00304)	(0.000279)	(0.000232)	(0.000560)	(0.000435)
Constant	0.877***	1.056***	0.428***	0.523***	1.651***	1.537***
	(0.00388)	(0.00366)	(0.000364)	(0.000362)	(0.000714)	(0.000660)
Condition	NEW	USED	NEW	USED	NEW	USED
R^2	0.00922	0.00128	0.000231	0.00292	0.00111	0.00210
N	125429	328119	8589646	10614542	8589646	10614542

Notes: This table presents results from regressions of the normalized price (see text for a discussion of the construction of reference prices) on the log of the number of photos of the listing. Robust standard errors are presented in parentheses. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

5 Unexplained Behavioral Patterns

We now turn to players' choices of counteroffers in order to further explore the gradual concessions we discussed in subsection 3.5 and, in the process, uncover another interesting behavioral pattern that is unexplained by conventional bargaining theories. In general, we are interested in exploring further patterns about how the offer in period t relates to the offers in periods $s < t$.

Let $\gamma_1 = p_1/p_0$, and, for $t = 2, 3, ..., 6$, let $\gamma_t \in [0, 1]$ be the weight such that $p_t = \gamma_t p_{t-1} + (1 - \gamma_t)p_{t-2}$. Therefore, γ_t represents the weight that player t places upon the opponent's previous offer, and $1 - \gamma_t$ represents the weight the player places on his or her own previous offer. Note that by definition, $p_1 = \gamma_1 p_0 + (1 - \gamma_1)0$, so we can think of the buyer's "previous" offer when he makes his first offer as his bliss point of paying nothing for the good. It is useful to think of γ_t as how much a player concedes to her opponent when she is making a counter offer, or her concession weight. This notion will prove useful to explore concession behavior in more detail as we show below.

5.1 Reciprocal Gradualism

As we explained in Section 3.5, existing theories have trouble explaining the appearance of gradual changes in counter offers, and as Figure 4 showed, the path of average counteroffers exhibits very strong gradual concessions. In fact, gradual changes in

Table 6: Concessions as a Function of Previous Concessions by One's Opponent

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	γ_3	γ_4	γ_5	γ_6	γ_3	γ_4	γ_5	γ_6
γ_{t-1}	0.116***	0.219***	0.191***	0.148***	0.120***	0.220***	0.191***	0.0864***
	(0.00149)	(0.00296)	(0.00638)	(0.00873)	(0.00190)	(0.00386)	(0.00850)	(0.0127)
$1\{\gamma_{t-1} = 0\}$	-0.0434***	-0.0279***	0.0368***	-0.0237***	-0.0469***	-0.0221***	0.0285***	-0.0423***
	(0.00316)	(0.00325)	(0.00309)	(0.00515)	(0.00360)	(0.00377)	(0.00407)	(0.00680)
Constant	0.347***	0.164***	0.274***	0.167***	0.349***	0.169***	0.268***	0.188***
	(0.000677)	(0.00106)	(0.00190)	(0.00266)	(0.000851)	(0.00138)	(0.00256)	(0.00390)
Condition	NEW	NEW	NEW	NEW	USED	USED	USED	USED
R^2	0.0513	0.106	0.0985	0.139	0.0565	0.0949	0.0989	0.130
N	514564	201767	56752	26427	336699	123139	32882	14269

Notes: This table presents results from regressions of γ_t on γ_{t-1} for $t = 3, 4, 5$, and 6. $1\{\gamma_{t-1} = 0\}$ is an indicator that a player's opponent did not budge in his or her last round of offers. Robust standard errors are presented in parentheses. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

offers are vastly more common in the data than are the sudden changes predicted by most theoretical analyses. In the data, among observations in which the seller makes at least two offers (beyond the BIN), we observe that only 0.78% of observations involve a seller standing firm at the Buy-it-Now price for several periods and conceding. In contrast, in 98.4% of the observed sequences the seller's first counteroffer already concedes a bit relative to the Buy-it-Now price. Examining analogous numbers for buyers, we find that only 0.42% involve concession only suddenly after holding firm at the previous offer for at least one period, whereas 95.9% involve a buyer conceding gradually.

To explore individual behavior rather than average behavior across bargaining threads, Table 6 presents results from regressing a player's concession weight, γ_t, on the opponent's previous concession weight, γ_{t-1}, where the unit of observation is one movement along each bargaining thread starting from the buyer's first counteroffer in period $t = 3$. A larger γ_{t-1} means that the player's opponent conceded more in his previous round, and a positive coefficient means that the player in period t is herself conceding more in response to a larger concession. This is the pattern we see for all periods $t = 3, 4, 5$, and 6, for both new and used goods: the more a player concedes in period $t - 1$, the more will his opponent concede in period t. Hence, gradual behavior is *reciprocal*, in that deeper concessions of one player are "rewarded" by deeper concessions of the following player.

To see what happens with extreme behavior when players do not budge at all from their earlier offers, we denote by $1\{\gamma_{t-1} = 0\}$ an indicator that the player in period $t-1$ did not budge from his previous offer (in $t-3$). As Table 6 shows, this behavior causes the opponent in period t to be *less* generous compared to what she would do if the player in period $t-1$ offered some additional concession beyond his offer from $t-3$. This further strengthens the narrative of reciprocal gradualism.

Together with the discussion in Section 3.5, the data strongly suggest that a major gap in the bargaining literature is a theory that can generate—from both parties in equilibrium—agreement *and* disagreement, delay *and* immediate termination, *and* the very robust pattern of reciprocal gradualism.

5.2 "Split-the-Difference" Offers

In defining the concession weight γ_t, it is insightful to observe the distributions of γ_t in different rounds of the game. Figure 8 displays histograms of these concession weights for the bargaining threads observed in the data. For simplicity, for this figure, we limit to threads with back-and-forth sequences corresponding to the left side of the game tree displayed in Figure 3; that is, we focus on threads with a series of offers and counteroffers, while ignoring threads in which the seller declines but the buyer continues to make additional offers. Panel A plots a histogram of γ_1, Panel B plots a histogram of γ_2, limiting to those threads in the data in which a period-2 offer was made, and so on.

Several interesting patterns are evident in this analysis. We note first that offers typically make nonzero concessions that are closer to the sender's own prior offer than the other party's (concession weights below 0.5), with the exception of the first offer, which is often close to the BIN price. Second, some common mass points emerge, and of particular notice are counteroffers that are halfway between the previous two offers, or "split-the-difference" counteroffers. The pattern even holds for buyers' first offers, where the modal initial offer is half of the BIN. The mid-point offer is also the modal offer for the first seller counter. In subsequent seller counters, the modal offer gives zero or nearly zero weight to the opponent's most recent offer, and second only to this choice is again the split-the-difference point.[21]

This pattern is consistent with previously documented laboratory evidence and behavioral economic theory (Roth and Malouf 1979; Roth 1985; Binmore et al. 1985; Bolton 1991; Bolton and Ockenfels 2000; Charness and Rabin 2002; Andreoni and Bernheim 2009), in which market participants may care about notions of fairness and may favor a split-the-difference strategy in negotiations. Interestingly, however, the split-the-difference pattern we observe is not a pattern of splitting *surplus* between the two parties, as the surplus is not necessarily known to the players given the potential presence of incomplete information about opponent valuations. Rather, the split-the-difference phenomenon we observe in our data regards to splitting the two most recent *offers*, regardless of how those offers relate to surplus.

We now explore how a player's choice of offer, as measured by the weight, γ_t, relates to later outcomes in the bargaining game. We create a measure for whether the offer is a 'split' offer by creating an indicator that is equal to one if γ_t is equal to 0.5 (after being rounded to the nearest hundredth) for each $t \in 1, 2, 3, 4, 5, 6$. We find that about 7 percent of offers are split offers by this definition.[22] We estimate a local linear regression of an indicator for whether each offer is accepted on both this split indicator and the underlying γ_t.[23] Results are shown in Table 7.

Table 7 demonstrates that, as would be expected, the coefficient on the concession rate (γ) is positive: the more a player concedes relative to previous offers, the more likely it is that the opposing player accepts the offer. The key result of Table 7, however, is that an offer in bargaining is more likely to be accepted if it is a *split* offer than if it is not, and this effect is both statistically significant and surprisingly large in magnitude, as well as being curiously stable across periods of the bargaining, lying in a range of 5–10% independent of what point in the bargaining game the split offer occurs. We supplement this approach with a more flexible fit of γ_t and plot fractional polynomial

[21]Appendix Figure E1 demonstrates that this same pattern emerges for very expensive items—those with list prices greater than \$1,000—which are not in our main sample.

[22]Broader definitions of split, by rounding γ_t to the nearest five hundredths or nearest tenth yield 9 percent and 14 percent split rates, respectively.

[23]We follow Fan and Gijbels (1992) in the construction of the optimal variable bandwidth for estimation of the effect at 0.5 using a rectangular kernel. See their paper for details.

Table 7: Probability of Split Offer Accepted

	(1)	(2)	(3)	(4)	(5)	(6)
	Pr(Accept)	Pr(Accept)	Pr(Accept)	Pr(Accept)	Pr(Accept)	Pr(Accept)
Split	0.0512***	0.0683***	0.0879***	0.0993***	0.0982***	0.0980***
	(0.000358)	(0.000559)	(0.00111)	(0.00216)	(0.00333)	(0.00645)
γ_i	0.894***	0.623***	1.199***	0.807***	1.018***	0.769***
	(0.00174)	(0.00269)	(0.00481)	(0.00643)	(0.0103)	(0.0117)
Constant	0.187***	0.181***	0.416***	0.343***	0.462***	0.489***
	(0.000136)	(0.000249)	(0.000592)	(0.00117)	(0.00177)	(0.00333)
Observations	9573704	3170486	993246	353115	135332	61656

Notes: This table displays the results from a linear regression of the probability of an offer being accepted regressed on the offer weight, γ_t, and on an indicator for whether γ_t is approximately equal to 0.5. Columns 1 through 6 display results for γ_t for $t = 1, ..., 6$. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

fits of acceptance and γ_t in Figure 9. As can be seen, the underlying relationship between γ_t and acceptance is positive and split offers are substantially more likely to be accepted.

This pattern of behavior is particularly surprising because, taken seriously, it implies a non-monotonicity in the likelihood of acceptance—that is, a player is more likely to accept a split-the-difference offer than an even slightly *more favorable* offer. This kind of discreet and non-monotonic behavior does not seem to be consistent with social preferences such as altruism or inequity aversion, or for that matter, any standard preferences for which bargaining can be modeled as a Bayesian game. Namely, if one considers the probability of trade as an equilibrium outcome, then a general mechanism design approach as in Myerson and Satterthwaite (1983) should imply that the probability of trade should be monotonic along some notion of private types if standard notions of interim incentive compatibility hold. Hence, the observed behavior presents a challenge to standard equilibrium theory, and suggests that there is some kind of "numerosity" effect at the split-the-difference points that seems to go beyond preferences, possibly some kind of norm, or some notion of salience.

Figure 9: Probability of Split Offer Accepted

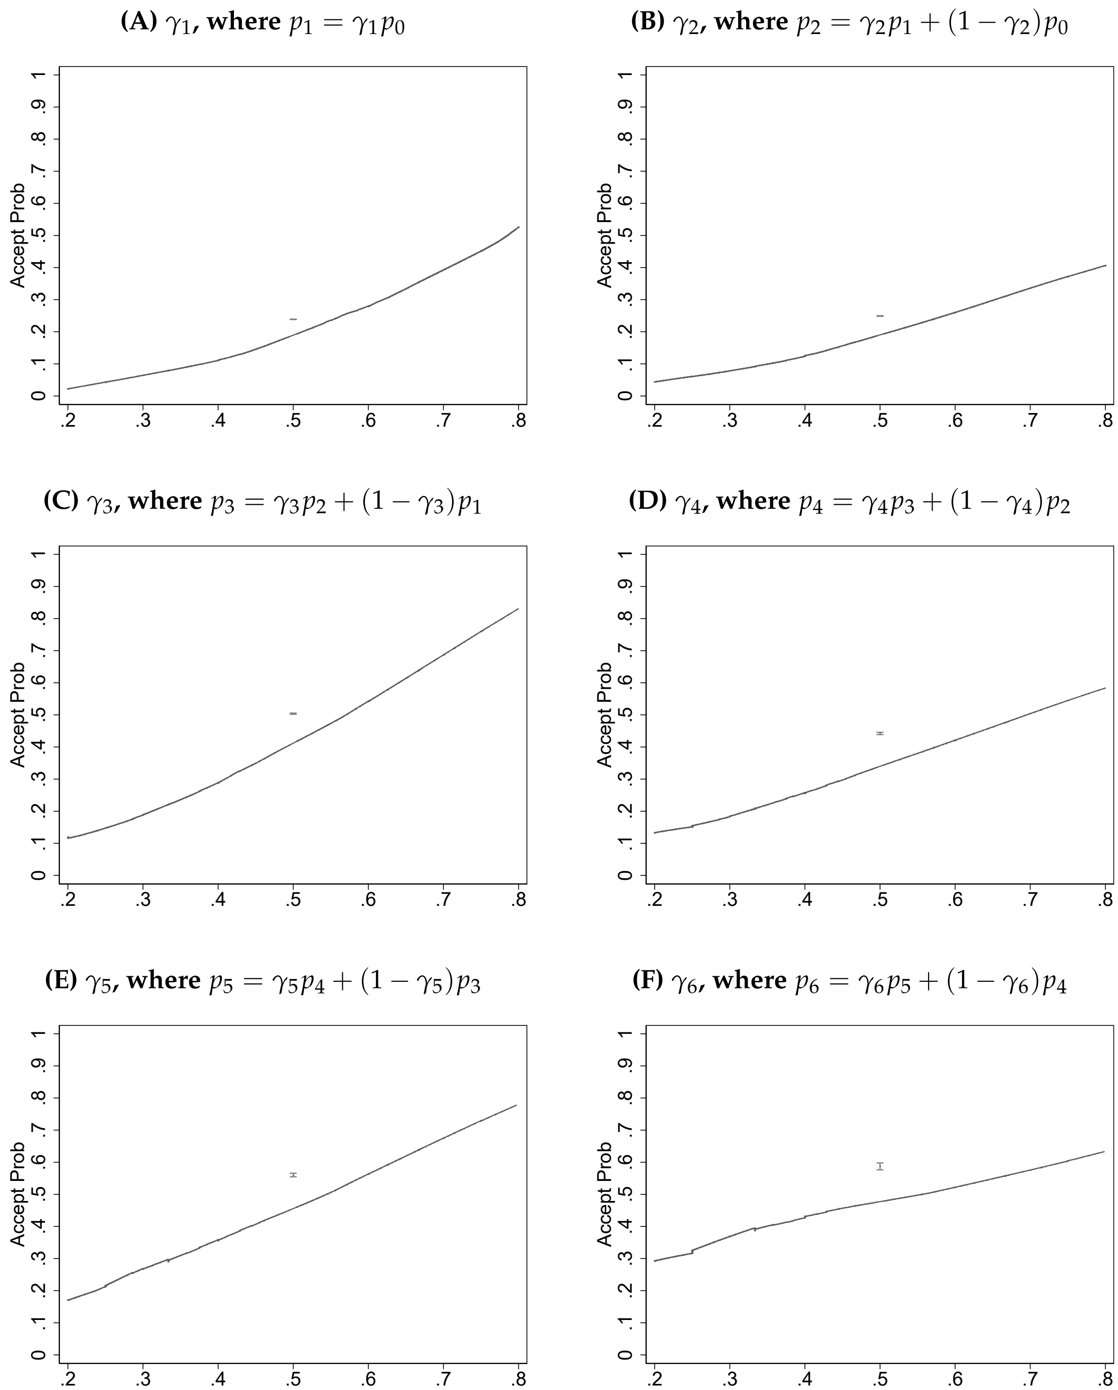

Notes: This figure displays a local polynomial fit of the probability of an offer being accepted regressed on the offer weight, γ_t, and on an indicator for whether γ_t is approximately equal to 0.5. From left to right, top to bottom, the panels display results for γ_t, where t ranges from 1 to 6.

6 Conclusion

In this paper we analyzed a novel dataset of bilateral bargaining used by millions of users in a live ecosystem. We documented a number of facts consistent with rational theories of bargaining behavior. In particular, we found that existing theories can explain a non-trivial fraction of the data in terms of how the bargaining ends (immediately vs. with delay, and in agreement vs. disagreement). We also found evidence favoring fixed over proportional costs of bargaining. This is important for understanding mechanism selection but also for thinking about the external validity of experimental studies of low-stakes bargaining: if our conclusion is correct, behavioral patterns in high- and low-stakes bargaining will be very difference. We also found that more patient or more experienced players obtain better deal, as do players facing less competition (or having a better outside option).

We then documented several features of the data that are clearly not consistent with any existing models of bargaining. First, we showed that bargaining often ends in disagreement after several back and forth offers. Second, we showed that buyer and seller offers tend to change gradually, not suddenly, over the course of the bargaining sequence, and that this gradualism is not one-sided: more concession by one player is associated more concession by her opponent as well. These two data patterns stand in stark contrast with existing bargaining theory models.

Finally, we offered evidence of "splitting-the-difference" behavior, a result that supports the incorporation of behavioral elements to understanding bargaining dynamics. We documented the surprising fact that counteroffers lying halfway between the two preceding offers are significantly more likely to be accepted by the opposing party than are offers which are even slightly more favorable to the opposing party.

We believe that the rich data we used herein, which we have made publicly available, offers opportunities to explore the ways in which people bargain, and can help shed light on what determines bargaining outcomes in the real world.

References

Abreu, D. and Gul, F. (2000). Bargaining and reputation. *Econometrica*, 68(1):85–117.

Abreu, D. and Pearce, D. (2003). A behaviorial model of bargaining with endogenous types. Working Paper, Yale University.

Admati, A. and Perry, M. (1987). Strategic delay in bargaining. *Review of Economic Studies*, 54(3):345.

Ambrus, A., Chaney, E., and Salitsky, I. (2018). Pirates of the mediterranean: An empirical investigation of bargaining with transaction costs. *Quantitative Economics*, 9(1):217–46.

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.

Backus, M., Blake, T., and Tadelis, S. (2019). On the empirical content of cheap-talk signaling: An application to bargaining. *Journal of Political Economy*, 127(4).

Bagwell, K., Staiger, R. W., and Yurukoglu, A. (2017). Multilateral trade bargaining: A first peek at the GATT bargaining records. Working paper, Stanford University.

Bajari, P., McMillan, R., and Tadelis, S. (2009). Auctions versus negotiations in procurement: An empirical analysis. *Journal of Law, Economics, and Organization*, 25(2):372–399.

Binmore, K., Rubinstein, A., and Wolinsky, A. (1986). The Nash bargaining solution in economic modelling. *RAND Journal of Economics*, pages 176–188.

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

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

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

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

Chatterjee, K. and Samuelson, W. (1983). Bargaining under incomplete information. *Operations Research*, pages 835–851.

Compte, O. and Jehiel, P. (2004). Gradualism in bargaining and contribution games. *The Review of Economic Studies*, 71(4):975–1000.

Cramton, P. (1992). Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59:205–225.

Cramton, P. C. (1991). Dynamic bargaining with transaction costs. *Management Science*, 37(10):1221–1233.

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

Deneckere, R. and Liang, M.-Y. (2006). Bargaining with interdependent values. *Econometrica*, 74(5):1309–1364.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics*, 7(2):215–247.

Fan, J. and Gijbels, I. (1992). Variable bandwidth and local linear regression smoothers. *The Annals of Statistics*, 20(4):2008–2036.

Feinberg, Y. and Skrzypacz, A. (2005). Uncertainty about uncertainty and delay in bargaining. *Econometrica*, 73(1):69–91.

Fuchs, W. and Skrzypacz, A. (2013). Bridging the gap: Bargaining with interdependent values. *Journal of Economic Theory*, 148(3):1226–1236.

Fudenberg, D., Levine, D., and Tirole, J. (1985). Infinite-horizon models of bargaining with one-sided incomplete information. In Roth, A., editor, *Game Theoretic Models of Bargaining*, pages 73–98. Cambridge: Cambridge University Press.

Gerardi, D., Hörner, J., and Maestric, L. (2014). The role of commitment in bilateral trade. *Journal of Economic Theory*, 154:578–603.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–177.

Grennan, M. and Swanson, A. (2016). Transparency and negotiated prices: The value of information in hospital-supplier bargaining. Working Paper, Wharton.

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

Gul, F. and Sonnenschein, H. (1988). On delay in bargaining with one-sided uncertainty. *Econometrica*, pages 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.

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.

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

Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2):245–252.

Larsen, B. (2018). The efficieny of real-world bargaining: Evidence from wholesale used-auto auctions. NBER Working paper 20431.

Larsen, B. and Zhang, A. (2018). A mechanism design approach to identification and estimation. NBER Working Paper 24837.

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of ebay motors. *American Economic Review*, 101(4):1535–46.

Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard Univ Press.

Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18(2):155–162.

Perry, M. (1986). An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica*, pages 313–321.

Rapoport, A., Erev, I., and Zwick, R. (1995). Bargaining behavior in a bilateral monopoly with one-sided incomplete information. *Management Science*, 41:377–394.

Roth, A. E. (1985). Toward a focal point theory of bargaining. volume 259, pages 265–67. Cambridge University Press Cambridge, UK.

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

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

Rubinstein, A. (1985). A bargaining model with incomplete information about time preferences. *Econometrica*, pages 1151–1172.

Scott Morton, F., Silva-Risso, J., and Zettelmeyer, F. (2011). What matters in a price negotiation: Evidence from the US auto retailing industry. *Quantitative Marketing and Economics*, pages 1–38.

Watson, J. (1998). Alternating-offer bargaining with two-sided incomplete information. *Review of Economic Studies*, 65(3):573–594.

Yildiz, M. (2003). Bargaining without a common prior—an immediate agreement theorem. *Econometrica*, 71(3):793–811.

Appendix

A Sample Construction

Starting with our original sample of approximately 92 million Best Offer listings, we imposed a handful of sample restrictions. Note that all sample restrictions are imposed at the listing level. That is, even if the sample restriction pertains to characteristics of an offer, we removed all listing that are associated with *any* offer that violates the restriction. In total, this leaves us with a sample of 82 million listings.

The sample restrictions are as follows:

Restrictions on Listing Attributes

(L1) *Listing price at or below $1000.*

(L2) *In the event of a sale, the sale price is at or below the listing price.*

Restrictions on Thread Attributes

(T1) *All offers are at or below the listing price.*

(T2) *Neither the buyer nor the seller makes more than three offers.*

(T3) *For all offers with a status of "countered," a counter-offer exists in the dataset.*

(T4) *For all offers accepted, there are no subsequent offers in the thread.*

The quantitative significance of these sample restrictions is described in Table A1. Restriction (L1), the largest, is an arbitrary restriction to simplify the analysis and the graphics and excludes approximately 10% of our sample. The second listing-level restriction, (L2), binds rarely, for only 42 thousand listings. It is possible that it happens because the sellers have agreed to bundle other products or services with the sale, however this is abnormal and inconsistent with eBay guidelines for communication in Best Offer bargaining.

Among the thread-level restrictions, (T1) is the most significant, affecting approximately 335 thousand listings. We suspect that it happens when an offer is made and

Table A1: Summary of Sample Restrictions

	No. Violations	Fraction of Listings
L1	9,547,987	0.0971
L2	42,524	0.000433
T1	386,096	0.00393
T2 - buyer	3,529	0.0000359
T2 - seller	0	0
T3	1,453	0.0000148
T4	1,111	0.0000113
No. Listings Before	98,307,281	
No. Listings After	88,388,220	

Notes: This table summarizes the incidence of violations of sample restrictions L1-L2, and T1-T4, to document the exclusion of approximately 10m listings from our main estimation sample.

the seller subsequently revises her price downward. In that case we only observe the final listing price, not the standing listing price that the buyer saw when making the offer. Indeed, we see from Table 1 that about 26% of listings have their prices adjusted at some point.

In addition to imposing this restriction on our sample, we check to make sure that all of our results are robust to excluding listings that were ever revised (results available on request). Restrictions (T2) - (T4) are rarely binding, and we do not have an explanation for them besides data processing errors.

B Reference Price Sample

Here we describe our reference price subsample—a portion of the data for which we observe catalog product identifiers, where each product identifier represents a distinct product SKU that can be linked to third-party catalogs to which eBay subscribes. These products are narrowly defined, matching a product available at retail stores, such as: "Microsoft Xbox One, 500 GB Black Console," "Chanel No.5 3.4oz, Women's Eau de Parfum," and "The Sopranos - The Complete Series (DVD, 2009)." We also construct a flag for the condition of the item as being new or used. For each product-by-condition cell, we compute what we refer to as a *reference price* (as in Einav et al. (2015)), which is the average price of sold fixed price listings of the same product and condition over the time frame of our sample that did not have the Best Offer option enabled, limiting to product-by-condition cells with at least 20 such transactions. These reference prices are computed entirely outside of our main sample, as our sample consists only of Best-Offer-enabled listings. For each thread in our bargaining data, we then compute a normalized price by dividing the final sales price (when a sale occurred) by the reference price for that product.

Table B1 shows an analogue of the descriptive statistics from Table 1 using this subsample of listings for which we are able to construct reference prices. The advantage of this subsample is that we have some evidence on the expected outcome price, so we can think about whether the buyer got a "deal." However, this comes at the cost of a somewhat opaque sample construction. By ruling out one-of-a-kind listings, for which no reference price will exist, it biases the sample towards the kinds of listings for which we expect relatively less bargaining. Consistent with this intuition, we see that these listings are substantially more likely to sell (47% as compared to 22%), but that there is less "room to bargain." Sale prices are substantially closer to list prices (91% as compared to 83%), and bargaining threads are substantially less likely to succeed (25% as compared to 45%). However, we still find that just over half of listings that sell in this sample are bargained (.227/.435 yields 52%).

C Stratified Subsampling for Figures 5 and 6

Figures 5 and 6 were constructed as LOWESS plots. To make this feasible, given the large sample size of the dataset, we subsampled. As Panel A of Figure 5 makes clear, however, stratification is required to obtain an adequate sample of listings at prices above $100. Therefore we sampled 70,000 listings from 20 bins, each $50 in length (inclusive on the upper extreme). Figure C1 presents two histograms to document the subsampling: Panel A is a histogram with a bin width that reflects the stratification strategy, and Panel B has 100 bins of length $10. The regular peaks in Panel B reflect the prevalence of round numbers (since our bins are constructed to be inclusive on the upper extreme). The increasing use of round numbers at higher prices may be affecting sale price outcomes from Figure 5, consistent with Backus et al. (2019).

Table B1: Descriptive Statistics for Reference Price Sample

	Mean	Std. Dev.	Min	Max
Listing-Level Data				
Listing Price	100	162	.99	1,000
Used	.719	.45	0	1
Revised	.233	.423	0	1
Sold	.467	.499	0	1
Sold by Best Offer	.224	.417	0	1
Received an Offer	.427	.495	0	1
Sale Price	111	158	.99	1,000
Sale Price / List Price	.91	.127	.00099	1
Bargained Price	122	161	.99	1,000
Bargained Price / List Price	.813	.124	.00099	1
No. Listings	2,047,079			
Seller-Level Data				
Feedpack Postitive Percent	99.3	4.84	0	100
No. Listings	9.05	120	1	26,132
No. Sales	4.23	42.6	0	9,415
No. Sales by Best Offer	2.02	18	0	4,078
No. Sellers	226,237			
Buyer-Level Data				
No. Bargaining Threads	7.96	29.5	1	4,734
No. Offers	13.4	48	1	7,229
No. Purchases	1.49	4.1	1	1,059
No. Bargained Purchases	1.07	2.46	0	408
No. Buyers	427,935			
Thread-Level Data				
No. Offers	1.7	.959	1	6
Agreement Reached	.25	.433	0	1
First Buyer Offer	152	166	0	1,000
First Buyer Offer / List Price	.681	.188	0	1
No. Threads	1,815,601			

Notes: This table presents summary statistics for the subsample of our data for which we have reference prices. Note that indicator "Used" (for used vs. new status of item) is only available for 2,044,419 listings, and feedback variables are only available for 218,739 sellers. See text for a discussion of exclusion criteria and, in particular, Appendix A.

D Heterogeneity: Players or Products?

Here we analyze the question of whether variance in bargaining outcomes is more a feature of *who* is bargaining or of *what* is being bargained over. The outcomes we examine are whether or not the bargaining pair comes to an agreement, how many periods the bargaining takes, and what price the players agree on when they do agree. The bargaining literature provides a number of possible explanations for why player heterogeneity may matter in explaining these outcomes: players may differ in their levels of patience, experience, or other measures of bargaining power/ability, or may differ in their valuations for the good. We also see the literature as establishing a role for heterogeneity in the items being bargained over, as markets for different items may be characterized by varying degrees of asymmetric information, for example. We explore these issues by regressing outcomes on buyer fixed effects, seller fixed effects, and product fixed effects in three separate regressions and reporting the R-squared and adjusted R-squared. For this exercise, we limit to our reference price sample.

The results are displayed in Table D1, separately for new items (columns 1–3) and used items (columns 4–6). For each of the three outcomes—normalized price, a dummy for whether the bargaining thread ended in agreement, and the number of offers—we find that buyer fixed effects explain more of the variation in the outcome than do seller fixed effects or product fixed effects. Specifically, over 83% of the variance in prices is explained by buyer fixed effects alone, while only 45–60% is explained by seller or product fixed effects. Variation in sales and the number of offers is also explained approximately twice as well by buyer fixed effects than by seller or product fixed effects. This finding may be unsurprising, as there are many more buyer fixed effects.

When we look at the adjusted R^2 values, the story is rather more subtle. For new products, product identity is substantially more relevant for predicting prices, but not for used ones—this is intuitive, as used products introduce heterogeneity poorly captured by product identifiers. In terms of the probability that a bargaining thread is successful or the number of offers, the buyer fixed effects are also favored, and the seller effects appear to explain more than the product fixed effects. We take this as evidence that, except in the case of new products where there is a clear outside option for both parties, buyer characteristics are of first-order importance for understanding bargaining outcomes. Seller characteristics also appear to be more important than the product identity in explaining variance in outcomes, suggesting that the traits of players play an important role in bargaining. We will return to this finding in Section 4.2 when we study buyer and seller characteristics influencing bargaining outcomes.

Table D1: Explaining Heterogeneity in Bargaining Outcomes

	Dependent Variable: Normalized Price					
	(1)	(2)	(3)	(4)	(5)	(6)
R^2	.8948	.598	.5938	.8389	.5203	.4553
Adj. R^2	.3714	.3719	.3652	.4563	.3551	.3187
No. FE	104,437	45,154	45,181	230,908	84,046	65,763
N	125,429	125,429	125,429	328,119	328,119	328,119
Condition	New	New	New	Used	Used	Used
Fixed Effects	Buyer	Seller	Product	Buyer	Seller	Product

	Dependent Variable: Sold					
	(1)	(2)	(3)	(4)	(5)	(6)
R^2	.5878	.3398	.2895	.5039	.2604	.1885
Adj. R^2	.2856	.2177	.1759	.2357	.171	.1254
No. FE	160,709	59,278	52,348	349,321	107,363	71,867
N	379,896	379,896	379,896	995,490	995,490	995,490
Condition	New	New	New	Used	Used	Used
Fixed Effects	Buyer	Seller	Product	Buyer	Seller	Product

	Dependent Variable: No. Offers					
	(1)	(2)	(3)	(4)	(5)	(6)
R^2	.5113	.2672	.1702	.44	.2191	.09192
Adj. R^2	.1529	.1317	.03753	.1372	.1247	.02126
No. FE	160,709	59,278	52,348	349,321	107,363	71,867
N	379,896	379,896	379,896	995,490	995,490	995,490
Condition	New	New	New	Used	Used	Used
Fixed Effects	Buyer	Seller	Product	Buyer	Seller	Product

Notes: This table presents R^2 and adjusted R^2 coefficients from regressions of three dependent variables—normalized prices conditional on sale (see text for a discussion of the construction of reference prices), a dummy for whether the thread ends in a sale, and the number of offers—where we vary both the condition of the item and the inclusion of buyer, seller, and product fixed effects.

We now perform this same analysis category-by-category. Results are displayed in Table D2. Each cell in the table represents the adjusted R^2 from the corresponding fixed effects regression. We find substantial variation across categories and across item condition (new vs. used) in the explanatory power of buyer, seller and product fixed effects. For example, in collectibles, buyer effects explain little or none of the variation in prices of new goods, but explain nearly all of the variation (an adjusted R^2 of 0.894) in used-good prices. Seller fixed effects alone explain much of the variation in new- and used-good prices (values of 0.821 and 0.790, respectively). Product fixed effects explain much of the variation in new good prices (0.491) but not for used goods (0.127). In the media category, on the other hand, buyer and product fixed effects explain more of the price variation for new goods than for used goods.

E Concession Weights and Experience

In Table E1, we explore how these concession weights relate to bargaining power as measured by buyer and seller experience. We run regressions of each stage's concession weight on buyer and seller experience (measured as in Table 3). The coefficients of the linear-log regressions can therefore be interpreted as the effect of experience on the concession weight (γ). Recall that odd t's represent buyer turns. In column 3, for example, the significant positive coefficient on seller experience indicates that a one log point increase in the seller experience measure is associated with a 0.002 increase in the concession weight placed by the buyer on the seller's previous offer, and the significant negative coefficient on buyer experience in column 3 indicates that a one log point change in the buyer experience measure is associated with a 0.012 decrease in the buyer's concession weight. Thus, in their period 3 offers, buyers tend to concede more to more experienced sellers and concede less if they themselves are more experienced. This same pattern is observed in the period 5 offers in column 5.

Results for seller offers weights can be seen in even columns. In columns 2 and 4 we find the same patterns for sellers that we observed for buyers: sellers tend to concede more to more experienced buyers and concede less if they themselves are more experienced. For example, we find that, in seller's period 2 offers, a one log point change in the buyer experience measure is associated with a 0.003 increase in

Table D2: Heterogeneity in Bargaining Outcomes by Category

	Dependent Variable: Normalized Price					
Collectibles	0.000	0.821	0.491	0.894	0.790	0.127
Electronics	-0.0687	0.358	0.318	0.039	0.264	0.134
Fashion	0.151	0.304	0.180	0.106	0.586	0.209
Media	0.685	0.397	0.294	0.627	0.531	0.068
Toys	0.290	0.780	0.643	0.383	0.245	0.116
Business	0.592	-0.459	-0.531	0.163	0.827	0.132
Others	-1.204	0.313	0.198	0.625	0.487	0.134
	Dependent Variable: Sold					
Collectibles	0.193	0.173	0.00406	0.0267	0.0661	0.0846
Electronics	0.255	0.187	0.130	0.213	0.144	0.102
Fashion	0.0973	0.139	0.0986	0.0361	0.0668	0.0417
Media	0.199	0.191	0.113	0.178	0.161	0.0328
Toys	0.156	0.140	0.105	0.0434	0.0986	0.0682
Business	0.0449	0.0124	0.0401	-0.0161	0.0140	0.0281
Others	0.144	0.141	0.0785	0.0620	0.100	0.0494
	Dependent Variable: No. Offers					
Collectibles	0.215	0.161	0.0444	0.238	0.205	0.0205
Electronics	0.170	0.137	0.0406	0.145	0.125	0.0258
Fashion	0.0549	0.0931	0.0501	0.157	0.221	0.0758
Media	0.116	0.133	0.0325	0.131	0.151	0.0137
Toys	0.094	0.0666	0.0274	0.0851	0.0861	0.0132
Business	0.189	0.109	0.0126	0.228	0.135	0.0387
Others	0.157	0.124	0.0598	0.0945	0.120	0.0471
Condition	New	New	New	Used	Used	Used
Fixed Effects	Buyer	Seller	Product	Buyer	Seller	Product

Notes: This table presents the adjusted R^2 coefficients from regressions of three dependent variables—normalized prices conditional on sale (see text for a discussion of the construction of reference prices), a dummy for whether the thread ends in a sale, and the number of offers—where we vary both the condition of the item and the inclusion of buyer, seller, and product fixed effects. We perform each regression separately for each of the listing categories displayed in the table.

the amount conceded by the seller, and a one log point change in the seller experience measure is associated with a 0.004 decrease in the seller's concession.

This pattern is remarkably robust across time periods and across both buyers and sellers, with two minor exceptions. First, in column 1, we see that an increase in the seller experience measure is associated with a decrease in the buyer's concession. We expect that this is an artifact of selection due to the seller being the player to set the initial price (if more experienced sellers tend to set higher initial prices, buyer offers on these listings will appear to be conceding less to the seller, all else equal). Second, in column 6, we see that more experienced sellers appear to concede more in their final

Table E1: Offer Weights and Experience

	(1)	(2)	(3)	(4)	(5)	(6)
	γ_1	γ_2	γ_3	γ_4	γ_5	γ_6
Log of Buyer Experience	-0.0104***	0.00282***	-0.0121***	0.00252***	-0.00292***	0.00340***
	(0.0000225)	(0.0000455)	(0.0000869)	(0.000139)	(0.000257)	(0.000368)
Log of Seller Experience	-0.00353***	-0.00445***	0.00188***	-0.00363***	0.00204***	0.00109***
	(0.0000156)	(0.0000322)	(0.0000632)	(0.000101)	(0.000190)	(0.000276)
Constant	0.667***	0.446***	0.424***	0.257***	0.325***	0.182***
	(0.000123)	(0.000250)	(0.000503)	(0.000807)	(0.00155)	(0.00220)
Observations	24,695,728	6,741,903	1,679,447	731,893	217,156	101,771

Notes: This table presents results from regressions where the dependent variable is γ_t (see text for a discussion of the construction of this variable) and the independent variables are measures of buyer and seller experience. Note that buyers make offers when t is odd, and sellers make offers when t is even. The number of observations changes across columns (becomes smaller) because fewer observations reached later periods of bargaining. *: $\alpha = 0.10$, **: $\alpha = 0.05$, and ***: $\alpha = 0.01$.

offer weight (γ_6). We suspect this is positive because experienced sellers are more likely to be aware that this is the final offer, and are therefore more generous.

Figure 6: More Bargaining Outcomes by Listing Price

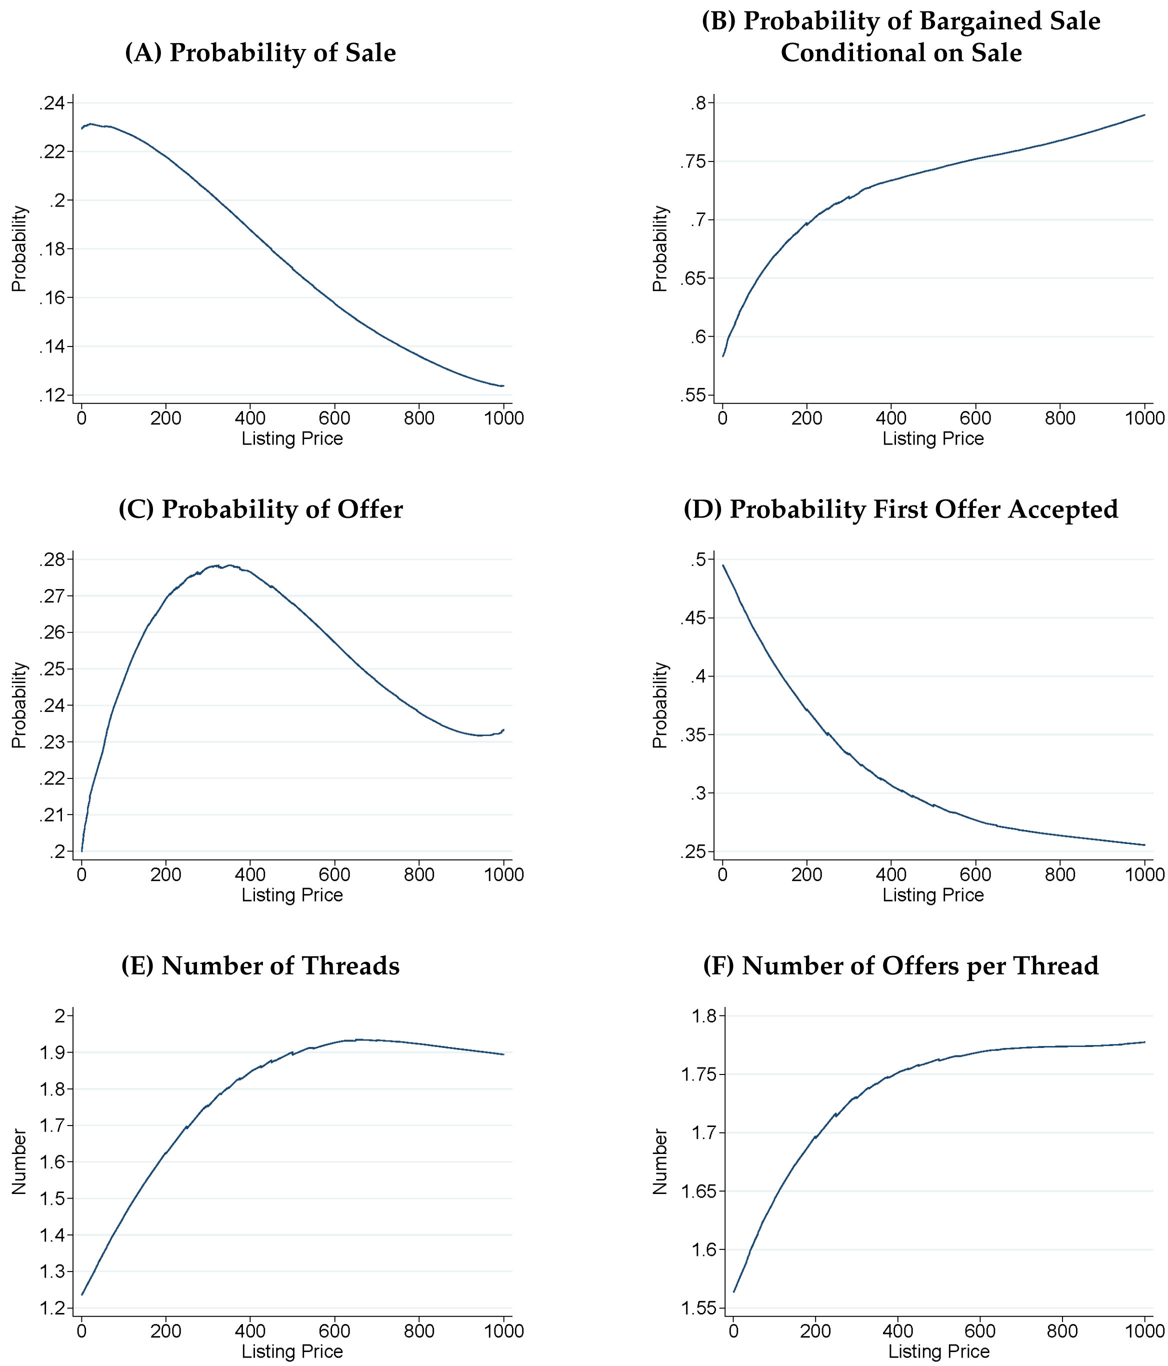

Notes: These panels depict LOWESS plots of bargaining outcomes in terms of the listing price. Panel A concerns the probability of sale for all listings; Panel B restricts attention to successful listings, and plots the likelihood that the price was bargained (as opposed to a buyer executing the Buy-it-Now option); Panel C concerns the empirical likelihood of receiving *any* offer; Panel D concerns the likelihood that, conditional on such an offer arriving, it is immediately accepted; Panel E concerns the number of bargaining threads per listing, and finally Panel F measures the number of offers associated with each thread, not including the listing price as an offer.

Figure 8: Where Current Offer Lies Relative to Previous Offers

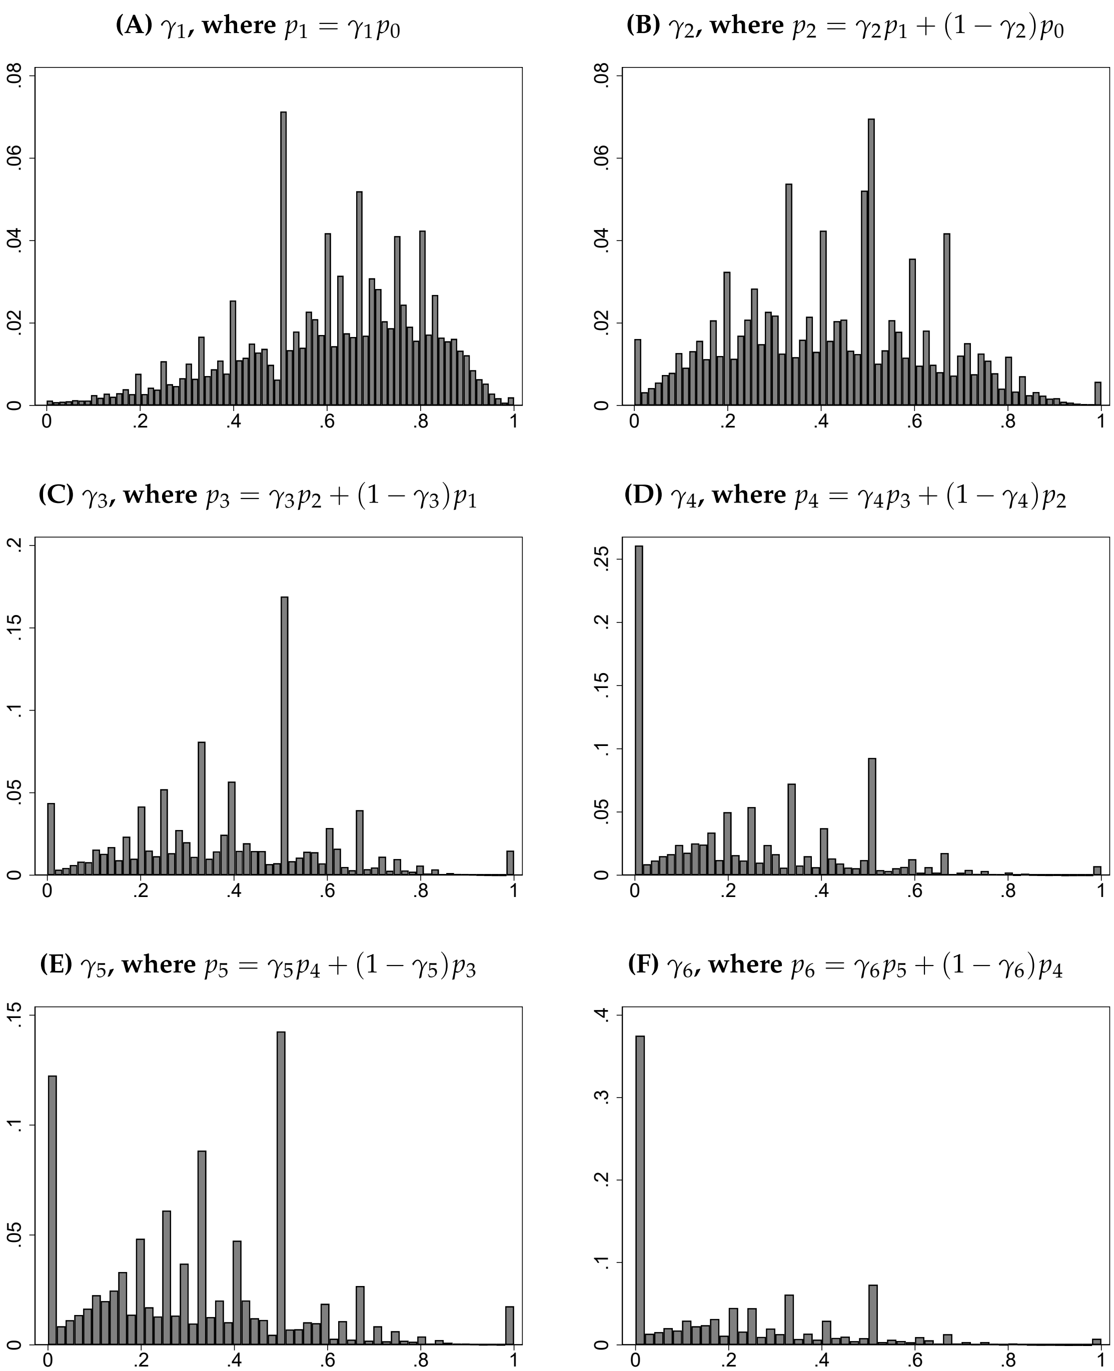

Notes: Each panel displays a histogram of offer weights defining how the current offer relates to the previous offers, where $\gamma_1 = p_1/p_0$, and, for $t = 2, 3, ..., 6$, γ_t is such that $p_t = \gamma_t p_{t-1} + (1-\gamma_t) p_{t-2}$.

Figure C1: Subsample Histograms

(A) Histogram with 20 Bins

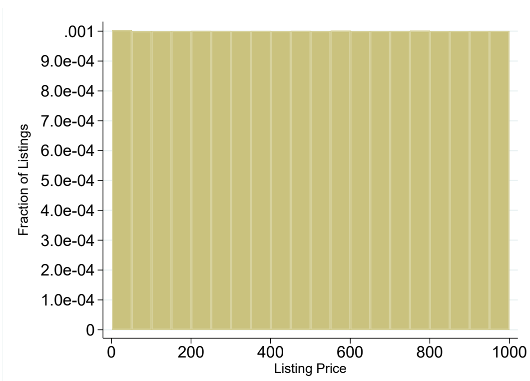

(B) Histogram with 100 Bins

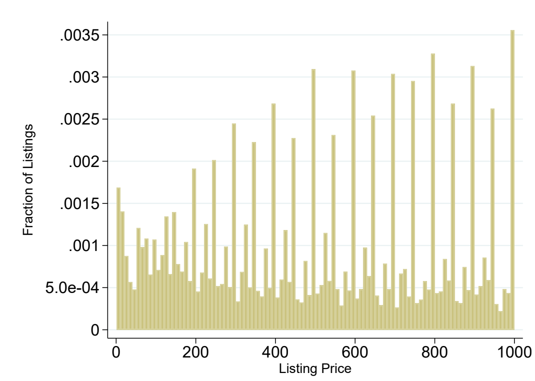

Notes: This figure presents two histograms of the stratified subsample with bandwidth 50 (panel A) and 10 (panel B) in order to document the clustering at round prices, even in the stratified subsample.

Figure E1: Where Current Offer Lies Relative to Previous Offers – Using Obs with BIN>$1,000

(A) γ_1**, where** $p_1 = \gamma_1 p_0$

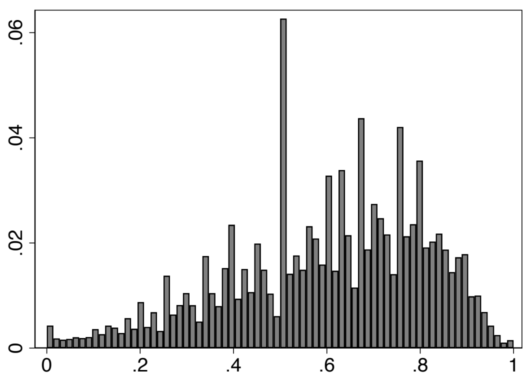

(B) γ_2**, where** $p_2 = \gamma_2 p_1 + (1 - \gamma_2) p_0$

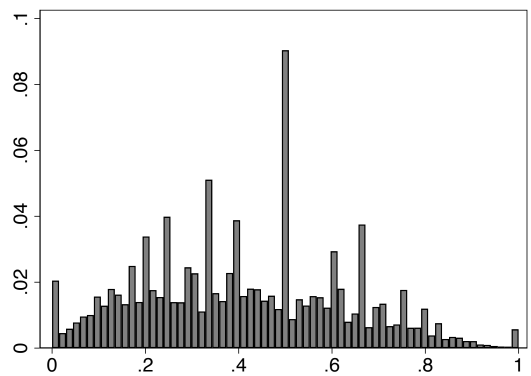

(C) γ_3**, where** $p_3 = \gamma_3 p_2 + (1 - \gamma_3) p_1$

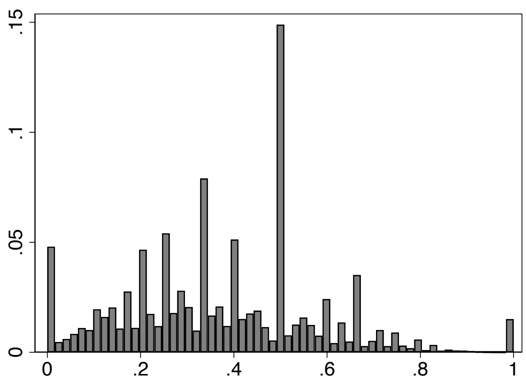

(D) γ_4**, where** $p_4 = \gamma_4 p_3 + (1 - \gamma_4) p_2$

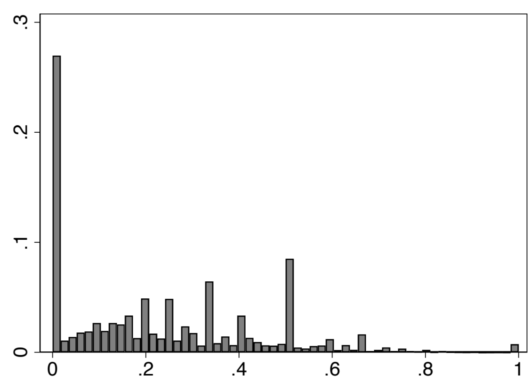

(E) γ_5**, where** $p_5 = \gamma_5 p_4 + (1 - \gamma_5) p_3$

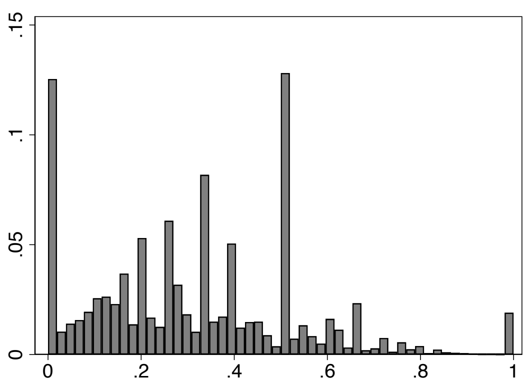

(F) γ_6**, where** $p_6 = \gamma_6 p_5 + (1 - \gamma_6) p_4$

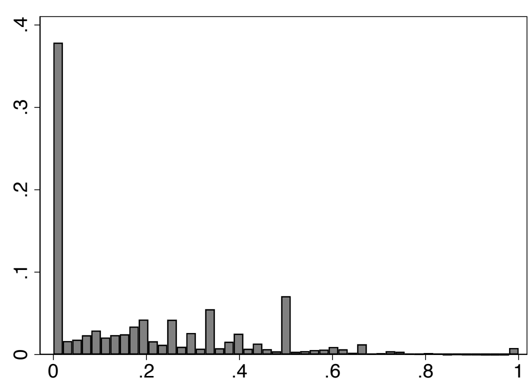

Notes: This figure uses the sample of listings with BIN prices over $1,000. Each panel displays a histogram of offer weights defining how the current offer relates to the previous offers, where $\gamma_1 = p_1 / p_0$, and, for $t = 2, 3, ..., 6$, γ_t is such that $p_t = \gamma_t p_{t-1} + (1 - \gamma_t) p_{t-2}$.