How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach[*]

Joachim Freyberger and Bradley J. Larsen[†]

September 6, 2024

Abstract

This study provides a structural analysis of detailed, alternating-offer bargaining data from eBay, deriving bounds on buyers and sellers private value distributions and the gains from trade using a range of assumptions on behavior and the informational environment. These assumptions range from weak (assuming only that acceptance and rejection decisions are rational) to less weak (e.g., assuming that bargaining offers are weakly increasing in players' private values). We estimate the bounds and show what they imply for consumer negotiation behavior and inefficient breakdown. For the median product, bargaining ends in impasse in 37% of negotiations even when the buyer values the good more than the seller.

[*]We thank Carol Lu, Sharon Shiao, Evan Storms, Greg Sun, Caio Waisman, and Wendy Yin for outstanding research assistance, and Brigham Frandsen, Matt Gentzkow, Phil Haile, Lars Lefgren, and Francesca Molinari for helpful comments, as well as seminar and conference participants at BEET 2018, NBER Summer Institute 2024, Workshop on Econometrics and Models of Strategic Interactions 2024, Berkeley Haas, BYU, Cornell, Indiana University, Michigan, Microsoft Research, Minnesota, MIT, Northwestern - Kellogg, Toulouse, Tsinghua, Universitat de les Illes Balears, UCL, University of Melbourne, University of Pennsylvania - Wharton, Washington University in St. Louis - Olin, Wisconsin, and Yale. We acknowledge support from NSF Grants SES-1530632 and SES-1629060. Larsen was a postdoctoral researcher, and then paid contractor, at eBay in the beginning stages of this project. Supporting material is available as an online appendix, see Freyberger and Larsen (2024)

[†]Freyberger: University of Bonn; freyberger@uni-bonn.de. Larsen: Washington University in St. Louis – Olin Business School and NBER; blarsen@wustl.edu.

1 Introduction

Bilateral bargaining is one of the oldest and most common forms of trade. A large theoretical literature and growing structural empirical literature examine the topic, but modeling choices of theorists and empiricists diverge widely, especially in how they consider the issue of *impasse*. Theoretical work (e.g., Myerson and Satterthwaite 1983) allows for the possibility that negotiators fail to agree even when gains from trade exist, whereas the workhorse model for empirical studies — Nash bargaining, in various forms — assumes that inefficient impasse *never* occurs. In these empirical models, negotiating agents know the opposing party's value precisely, and hence agents only negotiate over how to split a pie of known size. In many real-world settings, these strong assumptions are immediately rejected by data. In this paper, we analyze a large, detailed dataset of alternating-offer sequences from consumers negotiating online. We propose an approach to bound buyers' and sellers' values and the degree of inefficient impasse. Unlike Nash bargaining, the approach accommodates the presence of incomplete information. Our quantification exercises are aimed at providing motivation for more realistic theoretical empirical bargaining models and determining how well bargaining performs in practice in consumer markets.

The data comes from eBay's Best Offer platform, with thousands of eBay listings, each corresponding to a particular product identifier (such as an iPhone 6 or X-Box 360). For each listing, the seller posts a list (a Buy-It-Now) price and a buyer begins negotiating with an offer. We observe these prices and all counteroffers between any buyer-seller pair.

We model each bilateral bargaining pair as a buyer with value $B \sim F_B$ negotiating sequentially with a seller of value $S \sim F_S$. The key objects we wish to bound are F_B, F_S, and $P(B \geq S)$, the probability that the buyer values the good more than the seller (the trade probability in a first-best world). Comparing this object to the realized trade probability in the data offers a measure of the inefficient impasse relative to the first-best outcome.

The challenges we face are, first, S and B are not observed in the data, and second, there is no theoretical characterization of equilibria in our game (bilateral negotiations in which both parties potentially have incomplete information and can make offers), and hence no obvious way to identify F_S and F_B from observed bargaining actions.[1] Indeed, unlike *auction* games or *complete-information bargaining* games (e.g. the Rubinstein 1982 model of non-cooperative bargaining with complete information), there is no canonical model of bargaining under *incomplete* information. This dearth is especially pertinent for studying consumer price negotiations, where agents meet and negotiate infrequently and where it is arguably particularly unrealistic to model agents as perfectly informed about the game's structure or opponents' values (as Nash bargaining presumes, for example).

To study this setting empirically, we propose a bounds approach based on an *incomplete* model. We first derive bounds on F_B and F_S. We begin with weak rationality assumptions on agents' behavior. We then propose stronger conditions that appear in previous game theoretic bargaining models: monotonicity (an agent's value being weakly increasing in her first offer) and independence (an agent's value being independent of her opponent's first offer). We demonstrate theoretically that these assumptions can be violated by unobserved game-level heterogeneity, and then propose two weaker conditions, stochastic monotonicity (an agent's value being stochastically increasing in her offer) and positive correlation (an agent's value being stochastically increasing in her opponent's value). Building on these assumptions, we derive bounds on the gains from trade and, in turn, the first-best trade probability, leading to bounds on the degree of inefficient breakdown in the eBay data.

The bounds under any given set of assumptions are sharp. We propose nonparametric estimators for the bounds and estimate them separately for each product, limiting to products for which we have at least 200 bargaining sequences. To assess the validity of the underlying assumptions, we look for cases where bounds cross. We find evidence that our strongest

[1] Previous theoretical discussions have emphasized the difficulties of incomplete information in bargaining models. Fudenberg and Tirole (1991) claimed that "the theory of bargaining under incomplete information is currently more a series of examples than a coherent set of results. This is unfortunate because bargaining derives much of its interest from incomplete information." Binmore et al. (1992) observed, "In spite of this progress [in bargaining theory], important challenges are still ahead. The most pressing is that of establishing a properly founded theory of bargaining under incomplete information. A resolution of this difficulty must presumably await a major breakthrough in the general theory of games of incomplete information." These challenges, still unsolved today, arise from the sequential nature of the game: belief updating after off-equilibrium-path actions can sustain an infinite set of on-path behavior. Refinements, such as perfect Bayes equilibrium, do little or nothing to narrow the set of equilibria. See discussion in Gul and Sonnenschein (1988), and see Ausubel et al. (2002) for a survey.

assumptions (weak monotonicity of the seller's first offer or independence of the buyer's value and seller's first offer) can be violated. Bounds based on our weaker assumptions, such as stochastic monotonicity, do not cross.

We also exploit *auto-accept* and *auto-decline* thresholds that sellers can report secretly to the platform; sellers respond to price offers that fall between these thresholds, but eBay automatically rejects or accepts prices lying outside of these thresholds. These secret prices are themselves bounds on the true distribution F_S. We confirm that the estimated bounds are consistent with these auto-accept/decline bounds (which are not explicitly used anywhere in computing our bounds).

Having demonstrated the informativeness of these bounds on the marginal distributions, we estimate bounds on the first-best trade probability separately for each product in our sample. Under the weakest assumptions, bounds on this object are uninformative, with the lower bound corresponding to the sale probability observed in the data and the upper bound being 1. Under our strongest assumptions, the bounds can cross. We propose assumptions of intermediate strength that are reasonable, informative, and do not cross.

The lower bound on the first-best trade probability can be compared to the trade probability in the data to infer a bound on the degree of inefficient impasse. For example, for a popular cell phone product in our sample, agents agree in the real-world negotiations 27.6% of the time. Under our preferred assumptions, the counterfactual first-best trade probability is bounded below by 0.508, suggesting that at least 46% of the time $(1-0.276/0.508=0.46)$, agents fail to reach an agreement even when the buyer truly values the good more than the seller. For the median product, this lower bound on inefficient impasse is 37.3%. The inefficient impasse lower bound ranges from 18.0% to 54.2% across all products.

We explore features of the negotiation or the agents themselves to study when efficiency appears to improve, which we define as a decrease in the inefficient impasse lower bound. We find that this lower bound decreases when agents communicate with one another via messages on the platform or when the seller provides more photos of the item (although these decreases are not statistically significant). The largest improvement comes from negotiations of new products (as opposed to used ones) and from sellers choosing to use eBay's auto-accept/decline feature, both of which are associated with a decrease in the lower bound on inefficient impasse. More seller reviews or additional buyer experience, in contrast, are associated with an increase in the inefficient impasse lower bound.

Our study contributes to the theoretical and empirical bargaining literature. Theory work studies incomplete information bargaining either by explicitly modeling the extensive form of the game or applying mechanism design tools. Even our strongest assumptions (monotonicity and independence) are satisfied in the environments and equilibria of extensive-form games in the literature (e.g. Perry 1986, Grossman and Perry 1986, and Cramton 1992). We demonstrate, however, that in the presence of unobserved game-level heterogeneity (i.e., features of the negotiation that shift or scale the values of both agents in a given instance of the game, but that are unobservable to the econometrician), monotonicity assumptions can fail. This is not an indication that these theoretical equilibria cannot possibly describe real-world bargaining games well, but rather that data limitations (unobserved heterogeneity) can invalidate any attempt to use these existing theoretical results to analyze bargaining, even if the researcher is confident that she knows which of many equilibria generates the data. We show that our milder assumptions, such as stochastic monotonicity, can still be satisfied under unobserved heterogeneity.

In the mechanism design literature, our study is related to Myerson and Satterthwaite (1983), who demonstrated that when agents have independent values and face uncertainty about whether gains from trade exist, no incentive-compatible, individually rational mechanism will realize the first-best surplus without running a deficit. Our study quantifies how close real-world negotiators in a consumer market get to the first-best trade probability.

Our work relates to a small but recently growing literature estimating structural models of incomplete-information bargaining. The most closely related studies are Keniston (2011), who studied bargaining for auto-rickshaw rides in India, and Larsen (2021), who analyzed bargaining between used-car businesses. Our study is distinct in several dimensions. First, we study a setting where both agents may be inexperienced negotiators (unlike the drivers in Keniston 2011 or used-car businesses in Larsen 2021). The importance of this distinction is that previous studies assumed more about agents' behavior (such as optimality) or knowledge of game outcomes that, while plausible for the frequent market participants and professionals in those studies, are unlikely to hold when applied to consumers in a marketplace like eBay. Our study develops a new, incomplete-model approach that relies on a series of intuitive (and falsifiable) assumptions, and takes these bounds to real-world consumer negotiation data to estimate private values and the degree of inefficient impasse. Our study is also distinct methodologically. Keniston (2011) relied on inequality bounds generated from a two-step dynamic game method. Larsen (2021) relied on auction (in addition to bargaining) data, and applied a special case of one of the bounds we propose herein (seller independence).[2] In contrast, the methodology we develop does not rely on auction data, only sequential-offer bargaining data, and extends beyond the independence case.

Several structural empirical studies have focused on take-it-or-leave-it-offer bargaining (e.g. Silveira 2017, studying judicial settings) or sequential bargaining with all offers by one party (Ambrus et al. 2018, studying ransom negotiations for Spaniards taken captive by North African pirates in the seventeenth century). Li and Liu (2022) studied incomplete-information bargaining in the form of a k double auction, where each party simultaneously makes a single offer. In our setting, multiple offers from both parties can and frequently do occur, and hence the frameworks of these previous papers do not apply.

Our work also relates to a literature exploiting eBay as a laboratory for studying fundamental questions of price discovery and efficiency. The structural literature examining efficiency of eBay trading mechanisms has largely focused on auctions (e.g. Hendricks et al. 2021; Bodoh-Creed et al. 2021). Backus et al. (2020), Keniston et al. (2024), and Green and Plunkett (2022) documented a number of patterns in eBay bargaining data consistent with the existence of incomplete information and cognitive limitations, underscoring the benefit of our flexible approach bounding agents' values without assuming a complete

[2] Two structural empirical studies that also examined the used-car setting are Larsen and Zhang (2021) and Larsen et al. (2022). The latter paper relied on the methodology of Larsen (2021) and examined the impact of intermediaries in bargaining, while the former studied bargaining power using incentive compatibility and optimality assumptions that may be too strong for the consumer negotiations we study.

model of fully rational equilibrium behavior.

Finally, our work connects to a large literature using partially identified models to study objects of interest, such as distributions of wages or treatment effects. Canonical papers in this literature include Manski (1989, 1990) and Manski and Pepper (2000), among others. Examples that, like ours, begin with weak, uncontroversial assumptions and add stronger assumptions to improve bound informativeness, are Blundell et al. (2007), studying wages, and Frandsen and Lefgren (2021), studying treatment effects. See Ho and Rosen (2017) for additional discussion of this literature.

2 eBay's Best Offer Platform

eBay's *Best Offer* option, introduced in 2006, permits a posted-price seller to "allow offers"; a prospective buyer will then see the Buy-It-Now price (which we will refer to as the *list* price) as well as a *Make Offer* button, illustrated for an iPhone 8 in Figure 1.[3] Clicking this button lets the buyer propose an offer. The seller responds by selecting a button to decline, accept, or counter. If she counters, it is then the buyer's turn to accept, decline, or counter. If the seller declines, the buyer may still choose to make a counteroffer. Each party is limited to three offers, and the buyer can purchase at the list price at any time.[4] If any agent delays responding more than 48 hours, the offer expires (effectively declining).

Our data comes from a mix of business-to-consumer and consumer-to-consumer price negotiations: buyers are typically retail consumers while sellers may be businesses or individuals. Consumers thus play an important role in this market, in contrast to the professional negotiators studied in Keniston (2011), Ambrus et al. (2018), or Larsen (2021). As some consumers may be particularly inexperienced with the eBay game, our study adopts a robust bounds approach that does not require a complete model of equilibrium behavior.

[3] This option has, for many years, been the default, and sellers must opt out to disable it.

[4] This three-offer limit stood during our time period; in more recent years, eBay moved to a five-offer limit. Sellers can specify *auto-accept* and *auto-decline* prices; buyer offers above the auto-accept or below the auto-decline price are automatically accepted or declined by the platform. In our analysis we take advantage of these secretly reported thresholds to examine the validity of our bounds.

Figure 1: Illustration of Best Offer Listing

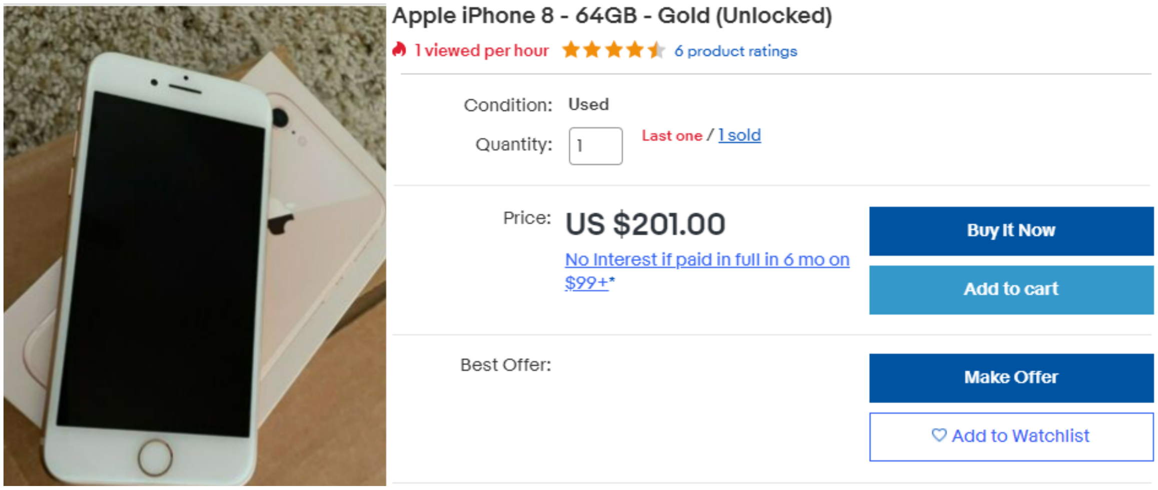

We use data created for the descriptive analysis of Backus et al. (2020), which studied bargaining on the U.S. eBay site from June 2012 to May 2013. We arrange the data with an observation being a given bargaining sequence, containing the list price and all offers.[5] For this paper, we focus on a subset of items with well-labeled product identifiers (bar codes), such as "Apple iPhone 8 64 GB" or "Xbox 360". We observe each item's condition (used/new), and we consider a *product* to be a combination of the condition type and product identifier. The data includes a *reference price*, the average of all non-Best-Offer posted-price sales of that same product during the sample period.[6] We restrict the sample to products with reference prices constructed from at least ten sales. We limit the sample to listings to which a buyer makes an offer. This means we do not analyze cases where a buyer arrives at a listing page but immediately buys at the Buy-It-Now price or leaves without making an offer. Our motivation for focusing on these listings is that we wish to analyze the degree of inefficient impasse conditional on both the buyer and seller indicating an interest in negotiating.

We refer to this sample as our original data. We impose several restrictions to remove

[5] We use the terms *negotiation* (or *negotiation sequence*) and *bargaining sequence* (or *sequence*) interchangeably.

[6] We use the reference price as a normalization to put each product on a similar scale by dividing prices/offers by its reference price.

incomplete sequences or cases where offers are extreme outliers. Section 3.2 discusses two additional sample restrictions that are best understood after introducing the model. To leave sufficient data to estimate our bounds, we also limit to products for which we observe at least 200 negotiation sequences after imposing all other restrictions. In the end, we are left with 12,012 bargaining sequences corresponding to 36 products. Table A1 in the online appendix describes each sample restriction.

Table 1: Descriptive Statistics: Highest-Selling Product per Category and Full Sample

Category	Reference Price ($)	*n*	P(sale)	Final Price Over List (if trade)	Buyer Price Over List (if no trade)	Seller Price Over List (if no trade)
Consumer Electronics	51.44	401	0.44	0.77	0.53	0.97
Video Games/Consoles	80.78	330	0.43	0.84	0.61	0.96
Cell Phones/Accessories	224.82	1,147	0.27	0.87	0.67	0.99
Computers/Tablets/Networking	87.46	318	0.37	0.84	0.62	0.98
All Products	216.64	12,012	0.30	0.84	0.64	0.98

Notes: First four rows show statistics for top-selling product in each category. Final row shows same statistics for estimation sample of 36 products (12,012 observations). *n* represents number of observations. Final three columns are averages across sequences within a product.

The data does not specify product titles (only anonymous product identifiers) but does specify product categories. Table 1 displays statistics for the top-selling product in each of the four categories in our final sample. These four products are all used products (only 3 of the 36 products in our sample are new). The products in Table 1 vary in reference price, ranging from $51 for the consumer electronics product to $224.82 for the cell phone product. When trade occurs, the final price as a fraction of the list price ranges from 0.77 to 0.87 on average. When trade fails, the highest price offered by the buyer as a fraction of the list price ranges from 0.53 to 0.67, whereas the lowest price offered by the seller in these disagreement cases ranges from 0.96 to 0.99. A key object of interest in this study is the fourth column, the sale probability ($P(sale)$), which varies from 0.27 for the cell phone product (which was the subject of 1,147 negotiations) to 0.44 for the consumer electronics product (401 negotiations). Our empirical approach allows us to quantify inefficiency by constructing bounds on the *counterfactual first-best* trade probability.

3 Bounds on Values in a Bargaining Game

In this section we present bounds on buyers' and sellers' values in an alternating-offer bargaining game, the protocol used on eBay.[7] We begin with bounds under minimal assumptions, using revealed preference arguments only. We then introduce assumptions on strategic behavior and the dependence of buyer and seller values to tighten bounds.

3.1. Bargaining Game Setup and Notation. A seller has value $S \sim F_S$ and a buyer has value $B \sim F_B$.[8] If the buyer and seller agree on a price P, the buyer's payoff is $B - P$ and the seller's is P. If they break up the bargaining (i.e., some agent quits), the seller gets S and the buyer gets 0. Throughout the paper, we maintain the assumption that, in a given instance of the game, S and B are known to the seller and buyer, respectively, before any actions take place and are held fixed throughout the negotiation.[9] To match the structure of eBay's protocol, we treat the first bargaining offer as coming from the seller (the list price). The buyer makes the second offer. The seller can then choose to accept, counter, or quit. We refer to each turn as a *period*, beginning with the seller at $t = 1$.

Denote the offers of the seller in odd periods t by P_t^S and the offer by the buyer in even periods t by P_t^B. Denote the decision of the buyer in even t by $D_t^B \in \{a, c, q\}$ (representing "accept", "counter", and "quit"). Similarly, let D_t^S be the decision of the seller in odd t. Either player accepting or quitting ends the game. For a given instance of the game, the data consists of the sequence of offers and any decisions to accept or quit.[10]

We define the following random variables at the *sequence*- (rather than period-) level: Let $X_{AC}^S = \min\{\{P_t^B : D_{t+1}^S = a\}, \min_t\{P_t^S : D_t^S = c\}\}$ be the smallest offer the seller makes or accepts in a given sequence (where "AC" stands for "accepting or countering"). Note

[7]The bounds we derive can be modified to allow for protocol other than alternating offers.

[8]Throughout, we use uppercase letters to denote random variables and lowercase to denote realizations.

[9]We allow for the buyer and seller to be learning about their opponents' values during the game, but not for agents to learn any additional information about *their own* values (such as learning about the quality of the good), which Desai and Jindal (2024) show in laboratory experiments is certainly a possibility.

[10]If the buyer has not reached the three-offer limit, she can make an additional offer following a seller's choice to decline; we reclassify such sequences as consisting of the seller having not declined but rather having countered at her previous offer. Thus, sequences ending with a seller *quitting* only occur if a seller declines an offer and there is no further buyer action. Appendix D considers alternative seller quit definitions.

that X^S_{AC} is always defined because the seller always makes the first offer (so $D^S_1 = c$). Also, let $X^S_Q = \max\{\{P^B_t : D^S_{t+1} = q\}, 0\}$ be the offer at which a seller quits, if she does quit, and 0 otherwise. There can be at most one offer in a sequence at which an agent quits. The definition ensures that X^S_Q is well defined even if $D^S_t \neq q$ for all t. Let $X^B_{AC} = \max\{\{P^S_t : D^B_{t+1} = a\}, \max_t\{P^B_t : D^B_t = c\}\}$ be the largest price a buyer accepts or offers in a given sequence. This is always defined because we focus on cases where buyers make offers P^B_2 (so $D^B_2 = c$). Let $X^B_Q = \min\{\{P^S_t : D^B_{t+1} = q\}, \infty\}$ be the offer at which a buyer quits, if indeed the buyer quits, and ∞ if $D^B_t \neq q$ for all t.[11]

A number of the arguments we derive below rely on the following identities, which are representations for F_S and F_B applying the law of iterated expectations:

$$P(S \leq x) = \int P(S \leq x \mid P^S_1 = y) dF_{P^S_1}(y) \tag{1}$$

$$P(B \leq x) = \int P(B \leq x \mid P^S_1 = y, P^B_2 = z) dF_{P^S_1, P^B_2}(y, z) \tag{2}$$

where $F_{P^S_1}$ is the CDF of P^S_1 and $F_{P^S_1, P^B_2}$ is the joint distribution of P^S_1, P^B_2. When written as a function, $P(\cdot)$ represents the probability of a given event.

As a final piece of notation, for a given random variable Y, a random vector $X \in \mathbb{R}^{d_x}$, and a set $A \subset \mathbb{R}^{d_x}$, let $\overline{\text{supp}}(Y \mid X \in A)$ and $\underline{\text{supp}}(Y \mid X \in A)$ be the supremum and infimum of the support of Y given $X \in A$, respectively. Define $X^{S*}_{AC}(y) \equiv \underline{\text{supp}}(X^S_{AC} \mid P^S_1 \geq y)$ and $X^{S*}_Q(y) \equiv \overline{\text{supp}}(X^S_Q \mid P^S_1 \leq y)$. Thus, $X^{S*}_{AC}(y)$ is the smallest accept/counter price of sellers conditional on the first offer of sellers being at least y. $X^{S*}_Q(y)$ has a similar interpretation. For all (y, z) on the support of (P^S_1, P^B_2), define $X^{B*}_{AC}(y, z) \equiv \overline{\text{supp}}(X^B_{AC} \mid P^B_2 \leq z, P^S_1 = y)$ and $X^{B*}_Q(y, z) \equiv \underline{\text{supp}}(X^B_Q \mid P^B_2 \geq z, P^S_1 = y)$. Conditional on P^S_1, $X^{B*}_{AC}(y, z)$ is the largest accept/counter price of buyers conditional on events where the second offer of buyers is at most z. $X^{B*}_Q(y, z)$ has a similar interpretation.

3.2. Model Discussion Before deriving bounds, we discuss dynamics *across* negotiations, an issue not explicitly modeled but allowed for in our framework. While we refer to B and S as "values" for brevity, more precise terms would be "net values," or "willingness to pay"

[11] The support points ∞ and 0 are conservative and can easily be replaced with other assumptions.

and "willingness to sell." Throughout the paper, we focus on one negotiating pair at time, but our framework allows for the possibility that, in a given negotiation, the buyer may receive gross utility V from trading and a nonzero outside option μ from not trading, where μ is a continuation value in a broader game where the buyer returns to the platform or looks for the item elsewhere. The buyer's willingness to pay in this negotiation is $V-\mu$, and we refer to this quantity as the buyer's value, $B \equiv V-\mu$. For our purposes, when studying buyers, we do not need to separately identify bounds on V and μ (nor can we using our method); we only seek bounds on F_B.

Similarly, the seller's willingness to sell S is the utility the seller receives if trade does not occur, which, in a broader continuation game, represents the value she receives from re-entering the eBay market, attempting to sell the item elsewhere, or keeping the good. Thus, in this broader game, the objects S and B are not primitives. They are nonetheless the initial objects we are interested in bounding because doing so will allow us to study several properties of bargaining offers and a notion of efficiency that does not require separately identifying V and μ or unpacking S into more primitive objects.[12]

This notion of B representing a buyer's willingness to pay suggests that, in the eBay context, an upper bound on B for a given item will be P_1^S (the list/Buy-It-Now price), because an eBay buyer can choose to purchase at that price at any time, even after a breakdown in negotiations. We can incorporate this information into our notation above by modifying the definition of X_Q^B (which, as we will show, provides an upper bound on B) to be $\min\{\{P_t^S : D_{t+1}^B = q\}, P_1^S, \infty\}$. In our analysis, we will say that we use the *list-price-recall condition* when taking advantage of this alternative definition of X_Q^B (see Section 5).

As highlighted above, we treat B and S as fixed during a given negotiation and do not consider the possibility of these changing (such as through the buyer's outside option changing). Such changes are certainly possible if a buyer negotiates with multiple sellers at once or vice versa. To minimize this possibility, we limit our sample to negotiations in which neither party negotiates with multiple opponents in an overlapping time window.[13] This re-

[12] Appendix D shows how several forms of bargaining costs would also fit into this framework.

[13] We define this time window for a given negotiation as the time from the buyer's first offer to the last

striction drops 54.3% of negotiations from the original data. While this ensures that agents in our sample do not engage in multiple simultaneous negotiations on eBay, it is in principle still possible that agents' values could change during a given negotiation in ways that we cannot account for. For example, buyers may engage with auctions or non-Best-Offer posted-price listings on eBay (which are not in our data) or search for better prices on other platforms, any of which may change their willingness to pay in the current negotiation.[14]

Among the remaining, non-overlapping sequences, we limit to the first seller of a given product with whom a given buyer negotiates, and, among these, the first buyer with whom a given seller negotiates. This drops 13% of negotiations (see Table A1 in the Appendix) and yields a unique set of negotiating pairs: no seller appears twice for the same listing and no buyer appears twice for the same product. It is tempting to retain data on cases where a given seller negotiates with multiple (non-overlapping) buyers over time (for example, one buyer today and another buyer a week later), under the additional assumption that S stays constant *across* negotiations, not just within a negotiation. If this stronger assumption were true, future and past negotiations between a given seller and different buyers would contain information about S in the current negotiation. We do not adopt this assumption and instead limit our data to a single buyer per seller and vice versa. Our motivation is twofold: many of the bounds we derive do not immediately apply to cases where an agent negotiates with multiple opponents over time, and attempts to extend these assumptions

action taken. The 25th and 75th percentile of window length are 20 seconds and 2.55 hours; the 10th and 90th percentiles are zero seconds (which can happen if the sequence ends immediately through the auto accept/decline mechanism) and 25.34 hours. We say a sequence has *overlapping buyers* if the window overlaps that of any other sequence for that same seller/item and some other buyer (or if some other buyer bought the item through the Buy-It-Now option on a day that overlaps the window). A sequence has *overlapping sellers* if the window for a given buyer with a given seller overlaps the window for that same buyer and some other seller of a similar product, defined as any product within the same broad category from the four categories listed in Table 1. 42.5% of negotiations in the original data have overlapping buyers, 23.5% have overlapping sellers, and some have both, leading to 54.3% overall that have either.

[14]While we cannot rule out such behavior, note that, in eBay bargaining, when a buyer is awaiting the seller's response to an offer, the buyer is liable to pay that price if the seller accepts. This liability may increase somewhat the buyer's incentives to avoid transactions outside of the negotiation to avoid being accountable for buying the item twice. In our data, slightly more than half of negotiation time consists of the buyer waiting for the seller. To calculate this, we compute, for each negotiation, the percentage of the time window in which the buyer is awaiting a seller's response. Averaged across negotiations, this number is 59%.

quickly become unwieldy (see Section 3.4), making a set of unique pairs appealing. This set of unique pairs could be formed using the last (or a random) seller among those with whom a given buyer negotiates, rather than the *first*. Our motivation for selecting the *first* is to avoid mechanically selecting all negotiations ending in agreement, which may be more likely to involve high-value buyers facing lower-value sellers, where assumptions such as independence or positive correlation (described below) may be less likely to hold.[15]

3.3. Sharpness In our presentation of bounds below, we will say a CDF, F, is in the identified set of buyer values under a given set of assumptions if there exists a data generating process (DGP) satisfying those assumptions such that F is the CDF of buyer values. If no such DGP exists, F is not in the identified set under those assumptions. Our sharpness notion is one sided: under the assumptions and data used to derive a given lower bound, that bound and any CDF lying above it are in the identified set. Similarly, under the assumptions used to derive a given upper bound, that bound and any CDF lying below it are in the identified set.[16] Hence, these bounds hold uniformly over the input variable of the CDF and sharpness also implies that the bounds are weakly increasing. For some bounds, our proofs do not imply that the lower and upper bounds would constitute the sharp identified set if we were to consider the upper and lower bound assumptions jointly. These issues are related to sharpness discussions in Chesher and Rosen (2017) and Molinari (2020).[17]

3.4. Unconditional Bounds on Value Distributions. We now describe a range of assumptions about equilibrium behavior yielding sharp bounds on marginal value distributions.

[15]To understand this selection, suppose buyers who negotiate with multiple sellers exit the market after agreeing with one seller. If we form a dataset consisting of, for each buyer, the *last* seller with whom the buyer negotiates, we are guaranteed to include in the data the one seller with whom the buyer agrees (if indeed the buyer eventually agrees with some seller). By instead selecting the *first* seller with whom the buyer negotiates, this *might* be a negotiation ending in agreement (this will be true if the buyer only negotiates with one seller in total), but the sample selection protocol is not *guaranteed* to include each agreement. While avoiding this selection is our motive for focusing on the *first* seller a buyer engages with and vice versa, in practice this restriction has a negligible effect on our results, which are similar when using the last (or a random) seller a buyer negotiates with (and vice versa); see Appendix Tables A2 and A3.

[16]Throughout the paper we will use the term "lower bound" on a CDF to refer to a bound lying graphically below that CDF (and vice-versa for "upper bound"), although a graphical lower bound is in fact an upper bound on the random variable in the stochastic dominance sense.

[17]For example, even for our unconditional bounds, it is not necessarily true that a CDF lying between the lower and upper bound is in the identified set, which is a well known issue with interval data (Molinari 2020).

Our first and weakest assumption is the following:

Assumption A1 (Revealed Preferences)**.** *The seller (i) never accepts (or counters) at a price $P < S$ and (ii) never quits at a price $P > S$. The buyer (iii) never accepts (or counters) at a price $P > B$ and (iv) never quits at a price $P < B$.*

These assumptions are similar to those in Haile and Tamer (2003) for English auctions.[18] Importantly, A1 imposes only a weak rationality condition, not requiring that agents behave according to any equilibrium concept, although the conditions are weak enough to be satisfied by standard concepts, such as Bayes Nash or perfect Bayes.

Important implications of this assumption are that $X_Q^S \leq S \leq X_{AC}^S$ and $X_{AC}^B \leq B \leq X_Q^B$. These inequalities imply what we call our *unconditional* bounds on F_S and F_B:

$$P(X_{AC}^S \leq x) \leq F_S(x) \leq P(X_Q^S \leq x) \tag{3}$$

$$P(X_Q^B \leq x) \leq F_B(x) \leq P(X_{AC}^B \leq x) \tag{4}$$

Theorem 1. *(3) gives a sharp lower bound for F_S under A1.i and a sharp upper bound for F_S under A1.ii. (4) gives a sharp lower bound for F_B under A1.iv and a sharp upper bound for F_B under A1.iii.*

All proofs are in the online appendix, Freyberger and Larsen (2024). The proof follows immediately from $X_Q^S \leq S \leq X_{AC}^S$ and $X_{AC}^B \leq B \leq X_Q^B$. For sharpness of Theorem 1, given that we place no restrictions on behavior other than A1, nothing in the data or assumptions rules out the possibility that the play of the game is such that $X_{AC}^S = S$, and similarly for the other bounds. These bounds can be relatively tight in some cases and loose in others. Section F of the online appendix offers Monte Carlo simulations illustrating this point.

The bounds are nonparametric. They are weakly increasing and lie in $[0,1]$, and thus correspond to some CDF. The bounds will be valid even if the game has multiple equilibria or if the data is not all generated by the same equilibrium or by any standard notion of

[18] There, the authors assume that a bidder (i) never bids above her value and (ii) never lets another agent win at a price she is willing to beat. An important distinction is that upper and lower bounds exist *for each* observation in Haile and Tamer (2003). In contrast, in the two-sided bargaining game we study, a given negotiation may end with one side of the bounds unobserved (e.g., because no agent quits). We handle this complication by relying on *probabilities* of events, rather than on empirical CDFs of prices/bids alone.

equilibrium play. Furthermore, if the true DGP does not in fact entail sellers all drawing from the same distribution F_S — that is, if sellers (or, analogously, buyers) are asymmetric — the bounds will remain valid for the mixture distribution of values in the data.

These bounds place no restrictions on the dependence between B and S. For example, the bounds allow for the possibility that B and S are correlated through game-level heterogeneity that is either unobservable or observable to the econometrician.[19] One form of heterogeneity is $S = W + \tilde{S}$ and $B = W + \tilde{B}$, where $\tilde{S}$, $\tilde{B}$, and W are independent, and where W is known to both agents but not the econometrician. In this scenario, S and B are independent conditional on W, but, from the analyst's perspective, are correlated *across instances of the game* through W. A related possibility is multiplicative separability, where $S = W\tilde{S}$ and $B = W\tilde{B}$.[20] We do not impose either structure but highlight these below as special cases allowed for by our moderate assumptions and ruled out by our strongest assumptions (and by some existing theoretical models).

The inequalities in (3) and (4) also hold conditional on buyer and seller prices, but these stronger implication do not carry any further identifying power without imposing additional restrictions on the joint distribution of S, B, and prices. In the next subsection, we will explore whether theoretical bargaining models of the environment we study suggest such restrictions to tighten the unconditional bounds.

3.5. What Assumptions Are Suggested by Bargaining Theory? There are very few extensive-form equilibria studied in the literature from bargaining games close in generality to the one we study — a bargaining game with two-sided incomplete information and a continuous value distribution where both parties can make offers.[21] The two models closest to our setting are those of Perry (1986) and Cramton (1992). Appendix G describes

[19]In the case of used cell phones that we examine in Section 5, unobserved heterogeneity can include aspects of the seller's reputation or of the cell phone that both the buyer and seller observe but not the econometrician (e.g., a cracked screen in a listing photo or a protective covering included with the phone).

[20]Multiplicative or additive separability are two structures for unobserved heterogeneity commonly assumed in empirical auction work (e.g., Krasnokutskaya 2011; Freyberger and Larsen 2022).

[21]See Table A9 of Larsen (2021) for a breakdown of the theoretical literature modeling extensive-form, incomplete-information bargaining games. This literature largely focuses on models where only one side has a private value, only one side is allowed to make offers, or agents have only two possible values. eBay's limit of three offers per party is another feature of our setting not completely captured in existing models.

these in detail. Our bounds allow for a much wider range of possible outcomes than the equilibria in these two papers; indeed, these equilibria are among infinitely many that our bounds accommodate. We focus on these examples only because they are the two existing examples general enough to relate to our framework. Table 2 summarizes the assumptions and theoretical models we consider.

Table 2: Assumptions' Relationships to Models of Two-Sided Incomplete Information

Model/Environment	Uncond. (A1)	Mon. (A2)	Indep. (A3)	Sto. Mon. (A4)	Pos. Corr. (A5)
Cramton (1992)	✓	✓	✓	✓	✓
Cramton (1992) + Unobs. Het.	✓			✓	✓
Perry (1986)	✓	✓	✓	✓	✓
Perry (1986) + Unobs. Het.	✓			✓	✓
Assumption Type	Behavioral	Behavioral	Environment	Behavioral	Environment

Notes: Table shows which assumptions are satisfied in Cramton (1992) and Perry (1986) models, as well as in modifications of those models we derive with unobserved heterogeneity (Appendix G). Final row distinguishes between assumptions about behavior vs. those about informational environment.

3.5.1. Monotonicity. In both Perry (1986) and Cramton (1992), offers satisfy a property we refer to as *monotonicity*, describing how an agent's first offer relates to her own value:

Assumption A2 (Monotonicity)**.** *(i)* $\overline{supp}(S \mid P_1^S = y) \leq \underline{supp}(S \mid P_1^S = y')$ *for all* $y \leq y'$, *and (ii)* $\overline{supp}(B \mid P_1^S = y, P_2^B = z) \leq \underline{supp}(B \mid P_1^S = y, P_2^B = z')$ *for all* y *and all* $z \leq z'$.

A2.i describes own-offer monotonicity for sellers: for $y \leq y'$, a seller with $P_1^S = y$ has a weakly lower value than a second seller with $P_1^S = y'$, and hence the lowest price at which the second seller counters or accepts is an upper bound on the value of the first seller.[22] This is precisely what is represented by $X_{AC}^{S*}(y)$, defined in Section 3.1. Similarly, A2.i implies that the highest quit price among sellers with first offers less than y provides a lower bound on the value of the seller who has $P_1^S = y$. Part (ii) of A2, monotonicity for the buyer, only has to hold *conditional* on the seller's first offer. A sufficient condition for A2.i is that P_1^S is a deterministic and strictly increasing function of S, and similarly, a sufficient condition for A2.ii is that P_2^B is a deterministic and strictly increasing function of B conditional on P_1^S.[23] Using these arguments and the notation from Section 3.1, we obtain

[22] Notice that $A2.i$ implies that sellers posting identical first prices have the same valuations.

[23] While these are sufficient conditions, they are not necessary. In particular, our assumption allows sellers

$$\int \mathbf{1}(X_{AC}^{S*}(y) \leq x) dF_{P_1^S}(y) \leq F_S(x) \leq \int \mathbf{1}(X_Q^{S*}(y) \leq x) dF_{P_1^S}(y) \tag{5}$$

$$\int \mathbf{1}(X_Q^{B*}(y,z) \leq x) dF_{P_1^S, P_2^B}(y,z) \leq F_B(x) \leq \int \mathbf{1}(X_{AC}^{B*}(y,z) \leq x) dF_{P_1^S, P_2^B}(y,z) \tag{6}$$

Theorem 2. *(5) gives a sharp lower bound for F_S under A1.i and A2.i and a sharp upper bound for F_S under A1.ii and A2.i. (6) gives a sharp lower bound for F_B under A1.iv and A2.ii and a sharp upper bound for F_B under A1.iii and A2.ii*

These *monotonicity* bounds are derived as follows: Under A2 and conditional on $P_1^S = y$, we have $X_Q^{S*}(y) \leq S \leq X_{AC}^{S*}(y)$, and the objects $X_{AC}^{S*}(y)$ and $X_Q^{S*}(y)$ are non-random. We plug these objects into (1) to obtain (5), where $\mathbf{1}(\cdot)$ is the indicator function. The buyer bounds follow similarly. The bounds improve upon the unconditional ones by comparing the accept/counter or quit actions of agents across instances of the game. Appendix F shows simulations where the bounds do vs. do not improve upon the unconditional ones. Appendix G shows that Perry (1986) and Cramton (1992) satisfy monotonicity.

3.5.2. Independence. This assumption relates an agent's *value* to the opponent's *offer*:

Assumption A3 (Independence)**.** *(i) S is independent of P_2^B conditional on P_1^S, and (ii) B is independent of P_1^S.*

As the seller makes the first move, a natural assumption is that the seller's first offer depends on S. A3.ii takes this one step further and assumes that the seller's first offer does not depend on B; A3.i describes a similar condition for the seller's value, but conditional on the seller's first offer. Through this relationship between an agent's value and an opponent's offer, A3 captures a notion of independence between values. This assumption is satisfied in the Perry (1986) and Cramton (1992) environments.

A3 implies the following *independence* bounds:

$$\int \max_z m_{AC}^S(x,y,z) dF_{P_1^S}(y) \leq F_S(x) \leq \int \min_z m_Q^S(x,y,z) dF_{P_1^S}(y) \tag{7}$$

$$\max_{y'} P(X_Q^B \leq x \mid P_1^S = y') \leq F_B(x) \leq \min_{y'} P(X_{AC}^B \leq x \mid P_1^S = y') \tag{8}$$

(or buyers) with identical valuations to make different offers, but only in very specific settings. For example, if $S \in \{0,2\}$ and $P_1^S = S + U$ where $U \sim U[0,1]$, then A2.i still holds.

where $m_{AC}^S(x,y,z) = P(X_{AC}^S \leq x \mid P_1^S = y, P_2^B = z)$, $m_Q^S(x,y,z) = P(X_Q^S \leq x \mid P_1^S = y, P_2^B = z)$.

Theorem 3. *(7) gives a sharp lower bound for F_S under A1.i and A3.i and a sharp upper bound for F_S under A1.ii and A3.i. (8) gives a sharp lower bound for F_B under A1.iv and A3.ii and a sharp upper bound for F_B under A1.iii and A3.ii*

These bounds are obtained by applying (1) and (2) and then using A1 and A3. The bounds can be narrow or wide in practice; Appendix F illustrate both cases via simulations and discusses data features affecting the bounds' width.[24]

While the models of Cramton (1992) and Perry (1986) satisfy both monotonicity and independence, it is not hard to modify their environments to violate these assumptions. Appendix G shows that additively separable unobserved heterogeneity would violate both assumptions in both models, as listed in Table 2.[25] Table 2 also offers a categorization of assumptions. Our assumptions restricting how an agent's value relates to her own offers (potentially conditional on previous offers) can roughly be considered assumptions about *behavior*. Those that instead restrict how an agent's value relates to offers of her opponent roughly capture a notion of the *informational environment*. For example, independence of B and P_1^S is similar to an assumption of independence between B and S (but independence of B and P_1^S is more useful, as it directly relates an unobserved value to an observable action). This is only a rough taxonomy, however, as all assumptions involve agents' behavior.

Moreover, either type of assumption may be violated in a given dataset because it fails to describe behavior or an environment well, or because of a data weakness (e.g. unobserved

[24] The independence case shows how our bounds would become unwieldy, as alluded to in Section 3.2, if we were to attempt to exploit data from a single agent negotiating with multiple opponents over time, even if these negotiations are non-overlapping. Consider a seller negotiating with multiple buyers over time, and assume S and P_1^S are constant across negotiations. The generalization of the independence assumption would be that, conditional on P_1^S, S is independent of *each* P_2^B the seller faces, and our bounds would then condition on the (potentially large) set of such offers — a set that may differ in size across instances of the game. As highlighted in Section 3.2, this complication is one reason why we limit our data to only one buyer per seller and vice versa. The independence case is only one such example; other bounds become unwieldy as well without this restriction.

[25] Assumption A3.ii could (but will not necessarily) be violated by our choice, described in Section 3.1, to restrict the sample to cases where $D_2^B = c$. Consider the following period-2 buyer protocol: if $B \geq P_1^S$, accept with probability p_{BA} and counter otherwise; if $B < P_1^S$, quit with probability p_{BQ} and counter otherwise. When $p_{BA} = p_{BQ}$, if Assumption A3.ii holds without conditioning on $D_2^B = c$ it also holds with this conditioning. However, if $p_{BQ} \neq p_{BA}$, conditioning on $D_2^B = c$ could violate Assumption A3.ii in this example.

heterogeneity). This raises an important point for empirical work: theoretical equilibrium models of bargaining may be unhelpful for empirics if their results do not hold in the presence of unobserved heterogeneity across instances of the game. It is precisely empirical challenges such as this that motivate our incomplete modeling approach, which can help bridge the gap between restrictive, extensive-form (and complete) models and analysis of bargaining in actual negotiation data.

3.6. Weakening Monotonicity and Independence. Our next assumption generalizes monotonicity:

Assumption A4 (Stochastic Monotonicity)**.** *(i) $P(S \leq x \mid P_1^S = y)$ weakly decreases in y $\forall$ x, and (ii) $P(B \leq x \mid P_1^S = y, P_2^B = z)$ weakly decreases in z $\forall$ y, x.*

This assumption is implied by monotonicity. Like monotonicity, stochastic monotonicity describes how an agent's offer is related to her value and allows us to exploit comparisons across instances of the game. A4 means that an agent's value is more likely high when her first offer is high.[26] Combined with (1) and (2), we obtain the following *stochastic monotonicity* bounds:

$$\int \max_{y' \geq y} P(X_{AC}^S \leq x \mid P_1^S = y') dF_{P_1^S}(y) \leq F_S(x) \leq \int \min_{y' \leq y} P(X_Q^S \leq x \mid P_1^S = y') dF_{P_1^S}(y) \tag{9}$$

$$\int \max_{z' \geq z} m_Q^B(x, y, z') dF_{P_1^S, P_2^B}(y, z) \leq F_B(x) \leq \int \min_{z' \leq z} m_{AC}^B(x, y, z') dF_{P_1^S, P_2^B}(y, z) \tag{10}$$

where $m_Q^B(x,y,z) = P(X_Q^B \leq x \mid P_1^S = y, P_2^B = z)$, $m_{AC}^B(x,y,z) = P(X_{AC}^B \leq x \mid P_1^S = y, P_2^B = z)$.

Theorem 4. *(9) gives a sharp lower bound for F_S under A1.i and A4.i and a sharp upper bound for F_S under A1.ii and A4.i. (10) gives a sharp lower bound for F_B under A1.iv and A4.ii and a sharp upper bound for F_B under A1.iii and A4.ii.*

Our next assumption weakens independence:

Assumption A5 (Positive correlation)**.** *(i) $P(S \leq x \mid P_1^S = y, P_2^B = z)$ is weakly decreasing in z for all y and x, and (ii) $P(B \leq x \mid P_1^S = y)$ is weakly decreasing in y for all x.*

[26] In another recent application of a similar assumption, Frandsen and Lefgren (2021) exploit stochastic monotonicity to bound treatment effects of charter school attendance.

A5 states that one agent's value is stochastically increasing in the other agent's first offer, capturing a notion of correlation between buyer and seller values.[27] A5 is implied by A3. Under A5 we obtain

$$\int \max_{z' \geq z} m_{AC}^{S}(x,y,z')dF_{P_1^S,P_2^B}(y,z) \leq F_S(x) \leq \int \min_{z' \leq z} m_{Q}^{S}(x,y,z')dF_{P_1^S,P_2^B}(y,z) \tag{11}$$

$$\int \max_{y' \geq y} P(X_Q^B \leq x \mid P_1^S = y')dF_{P_1^S}(y) \leq F_B(x) \leq \int \min_{y' \leq y} P(X_{AC}^B \leq x \mid P_1^S = y')dF_{P_1^S}(y) \tag{12}$$

Theorem 5. *(11) gives a sharp lower bound for F_S under A1.i and A5.i and a sharp upper bound for F_S under A1.ii and A5.i. (12) gives a sharp lower bound for F_B under A1.iv and A5.ii and a sharp upper bound for F_B under A1.iii and A5.ii*

Appendix G demonstrates that, when modifying Cramton (1992) and Perry (1986) to include unobserved heterogeneity, independence and monotonicity can be violated even while these weaker conditions hold. Table 2 summarizes these results.

3.7. Combining Assumptions on Marginal Distributions. A2–A5 can be combined to obtain tighter bounds. For example, we can combine A2 and A3 — monotonicity and independence — our two strongest assumptions. Or we can combine A4 and A5 — stochastic monotonicity and positive correlation — two weaker assumptions. Appendix C derives these combined bounds and proves sharpness.

The choice of which assumptions to adopt should be based on which seem reasonable in a given setting. For example, seller monotonicity and buyer independence may be inappropriate in settings with suspected unobserved heterogeneity, where stochastic monotonicity and positive correlation would be more appropriate. Other assumptions that may be less sensitive to unobserved heterogeneity include buyer monotonicity, which, through conditioning on the seller's first offer, may purge some unobserved heterogeneity. In Section 5, we estimate each of the bounds and test for crossings.

[27] While these bounds exhibit a form of opponent-offer stochastic monotonicity, we refer to them as the *positive correlation* bounds to distinguish them from A4, (own-offer) stochastic monotonicity.

4 Estimation

We now describe estimators for the bounds. For some of the bounds, the natural plug-in estimators are unbiased or have an outward bias. In other cases, however, the plug-in estimator are inward biased and might be artificially tight (and unbiased estimators do not exist; Hirano and Porter 2012). We therefore modify the plug-in estimator to be either outward biased or half-median-unbiased in the spirit of Chernozhukov et al. 2013. We describe the basics of our estimators here and relegate additional details, along with discussions of inference and testing, to Appendix E. An observation $i = \{1,\ldots,n\}$ is the bargaining sequence for a given buyer-seller pair negotiating over a given product. We do all estimation separately by product and thus omit any notation denoting products.

4.1. Preliminary Ingredients for Estimation. For each i, the variables $X^S_{AC,i}$, $X^S_{Q,i}$, $X^B_{AC,i}$, and $X^B_{Q,i}$ are observed. We estimate the conditional probability $P(X^B_Q \leq x \mid P^S_1 = y') = E\left[\mathbf{1}(X^B_Q \leq x) \mid P^S_1 = y'\right]$ for each value of x using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel and bandwidth $n^{-1/4}$. This bandwidth choice implies undersmoothing, which facilitates inference. Let $\widehat{P}(X^B_Q \leq x \mid P^S_1 = y')$ denote the estimator. We proceed analogously for $P(X^B_{AC} \leq x \mid P^S_1 = y')$, $P(X^S_{AC} \leq x \mid P^S_1 = y')$, and $P(X^S_Q \leq x \mid P^S_1 = y')$.

Similarly, we estimate the function $m^B_Q(x,y,z)$ using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel and bandwidth $n^{-1/5}$. Due to the higher dimension of $m^B_Q(\cdot)$, the bandwidth converges at a slower rate. Denote the estimator by $\hat{m}^B_Q(x,y,z)$. We estimate $\hat{m}^B_{AC}(x,y,z)$, $\hat{m}^S_{AC}(x,y,z)$, and $\hat{m}^S_Q(x,y,z)$ analogously.

$X^{S*}_{AC}(y)$ and $X^{S*}_Q(y)$ can be estimated with sample analogs, $\hat{X}^{S*}_{AC}(y) = \min_{i:P^S_{1,i} \geq y} X^S_{AC,i}$ and $\hat{X}^{S*}_Q(y) = \max_{i:P^S_{1,i} \leq y} X^S_{Q,i}$. Notice that $\hat{X}^{S*}_{AC}(y) \geq X^{S*}_{AC}(y)$ and $\hat{X}^{S*}_Q(y) \leq X^{S*}_Q(y)$, which implies that the estimated seller monotonicity bounds are outward biased (i.e., they are on average too wide) and thus are conservative when it comes to bounding the true CDF.

Estimating $X^{B*}_{AC}(y,z)$ and $X^{B*}_Q(y,z)$ is more complicated, as these condition on a specific value of the continuous variable P^S_1. Let $N(y) = \{z \in \mathbb{R} : |z-y| \leq h_n(y)\}$ be a neighborhood of y of a size $h_n(y)$ dependent on n and decreasing to 0 as $n \to \infty$. Define

$\hat{X}^{B*}_{AC}(y,z) = \max_{i:P^B_{2,i}\leq z, P^S_{1,i}\in N(y)} X^B_{AC,i}$ and $\hat{X}^{B*}_{Q}(y,z) = \min_{i:P^B_{2,i}\geq z, P^S_{1,i}\in N(y)} X^B_{Q,i}$. As opposed to $\hat{X}^{S*}_{AC}(y)$, the bias of $\hat{X}^{B*}_{AC}(y,z)$ cannot be signed. To obtain a conservative estimator, we assume that $X^{B*}_{AC}(y,z)$ is Lipschitz continuous. That is, there exists $C \geq 0$ such that $|X^{B*}_{AC}(y,z) - X^{B*}_{AC}(y',z)| \leq C|y-y'|$ for all y,y',z. Then $\hat{X}^{B*}_{AC}(y,z) \leq \overline{\text{supp}}(X^B_{AC} : P^B_2 \leq z, P^S_1 \in N(y)) \leq X^{B*}_{AC}(y,z) + Ch_n(y)$. To choose $h_n(y)$, we use a matching approach. Let K_n be the number of neighbors and let $h_n(y)$ be such that $\sum_{i=1}^n \mathbf{1}(|P^S_{1,i} - y| \leq h_n(y)) = K_n$. We choose $K_n = n^{1/4}$. If the density of $P^S_1(y)$ is bounded and bounded away from 0 in a neighborhood of y, then $h_n(y)$ is proportional to $n^{-3/4}$ and therefore goes to 0 as $n \to \infty$. Finally, we let $\tilde{X}^{B*}_{AC}(y,z) = \hat{X}^{B*}_{AC}(y,z) - \eta_n$ where $\eta_n = n^{-1/2}$, which ensures that $\tilde{X}^{B*}_{AC}(y,z) \leq X^{B*}_{AC}(y,z)$ with probability approaching 1. Similarly, we use $\tilde{X}^{B*}_{Q}(y,z) = \hat{X}^{B*}_{Q}(y,z) + \eta_n$.

4.2. Estimation of Bounds. With the ingredients from above, we estimate the bounds. For brevity, we describe here the estimation of each lower bound; estimators for upper bounds are analogous. The unconditional lower bound estimators are simply the empirical analogs of (3) and (4): $\frac{1}{n}\sum_{i=1}^n \mathbf{1}(X^S_{AC,i} \leq x)$ and $\frac{1}{n}\sum_{i=1}^n \mathbf{1}(X^B_{Q,i} \leq x)$, both of which are unbiased.

For the monotonicity bounds, note that $\int \mathbf{1}(X^{S*}_{AC}(y) \leq x) dF_{P^S_1}(y) = E_{P^S_1}[\mathbf{1}(X^{S*}_{AC}(P^S_1) \leq x)]$, where $E_{P^S_1}[\cdot]$ is an expectation over P^S_1. We estimate the lower bounds from (5) and (6) by $\frac{1}{n}\sum_{i=1}^n \mathbf{1}(\hat{X}^{S*}_{AC}(P^S_{1,i}) \leq x)$ and $\frac{1}{n}\sum_{i=1}^n \mathbf{1}(\tilde{X}^{B*}_{Q}(P^S_{1,i}, P^B_{2,i}) \leq x)$, which are both conservative.

We estimate the stochastic monotonicity lower bounds in (9) and (10) by
$\frac{1}{n}\sum_{i=1}^n \max_{y'\in\omega_1(P^S_{1,i})} \widehat{P}(X^S_{AC} \leq x \mid P^S_1 = y')$ and $\frac{1}{n}\sum_{i=1}^n \max_{z'\in\omega_2(P^B_{2,i})} \hat{m}^B_Q(x, P^S_{1,i}, z')$, where we define $\omega_1(P^S_{1,i}) = \{y : y \geq P^S_{1,i}, y \in Q_{0.05}(P^S_{1,i}) \cup \{P^S_{1,i}\}\}$, and for any random variable Y, $Q_\alpha(Y)$ is the interval between the α and $1-\alpha$ quantiles of Y; $\omega_2(P^B_{2,i})$ is defined analogously. Notice that the sample analog estimator of the seller's stochastic monotonicity lower bound is $\frac{1}{n}\sum_{i=1}^n \max_{y'\geq P^S_{1,i}} \widehat{P}(X^S_{AC} \leq x \mid P^S_1 = y')$, but we use the additional constraint that $y \in Q_{0.05}(P^S_{1,i})$ since $P(X^S_{AC} \leq x \mid P^S_1 = y')$ can be poorly estimated at the support boundary. Moreover, to ensure that we maximize over a nonempty set, we always include $P^S_{1,i}$. Applying this tail truncation yields conservative estimates; the same is true for all estimators in the paper using this truncation. These estimators might still be inward biased due to the maxima. To address this, we modify our estimators to be half-median-unbiased (Ap-

pendix E). Simulations in Appendix F compare bias-corrected and uncorrected estimators.

We estimate the independence lower bounds in (7) and (8) by $\frac{1}{n}\sum_{i=1}^{n}\max_{z\in Q_{0.05}(P^B_{2,i})}\hat{m}^S_{AC}(x,P^S_{1,i},z)$ and $\max_{y'\in Q_{0.05}(P^S_{1,i})}\widehat{P}(X^B_Q\leq x\mid P^S_1=y')$, and the positive correlation lower bounds in (11) and (12) by $\frac{1}{n}\sum_{i=1}^{n}\max_{z'\in\omega_2(P^B_{2,i})}\hat{m}^S_{AC}(x,P^S_{1,i},z')$ and $\frac{1}{n}\sum_{i=1}^{n}\max_{y'\in\omega_1(P^S_{1,i})}\widehat{P}(X^B_Q\leq x\mid P^S_1=y')$. Appendix E discusses corresponding half-median-unbiased estimators. Appendices C and E address estimators combining assumptions.[28]

5 Bounding Values in eBay Bargaining

5.1. Bounds on Buyer and Seller Values for Cell Phones. We apply these estimators to bound buyer and seller value distributions separately for all 36 products. We normalize prices by the product's reference price to aid interpretation. To illustrate, we first focus on one product: the most popular product from the cell phone category from Table 1.

5.1.1. Bounds on Seller Values. Figure 2 shows bounds on F_S for this product under different assumptions. Every panel also shows the unconditional bounds for comparison. Dashed lines are upper and solid lines are lower bounds. Next to point estimates, we also show 95% one-sided pointwise confidence bands.

The unconditional bounds — which rely on our weakest assumptions — can be wide, the upper bound in particular. This is because it is constructed using prices at which a seller quits (walking away from bargaining), which are unobserved if a sequence either ends in agreement or if the buyer quits (rather than the seller).[29] Depending on the product, some assumptions do little to improve the unconditional bounds. For example, for this product, the stochastic monotonicity bounds (top left panel), are nearly as wide as the unconditional

[28] While we do not address this in this paper, one could condition bounds arguments on covariates and apply this conditioning in the estimated conditional probability functions. In this case, a parametric approximation, such as a probit, may be preferred to the nonparametric estimators we propose here. The matching approach described above could then be used to estimate support bounds for the monotonicity assumptions.

[29] In such cases, we cannot rule out the seller having a value of zero, and hence the upper bound CDF is (graphically) high. The lower bound, on the other hand, relies on prices at which the seller accepts or counters, and at least one of these prices is *always* available in the data (P^S_1). The lower bound will thus, by construction, be surjective (i.e., mapping to each value in [0,1]).

Figure 2: Bounds on Seller Distribution for Cell Phone

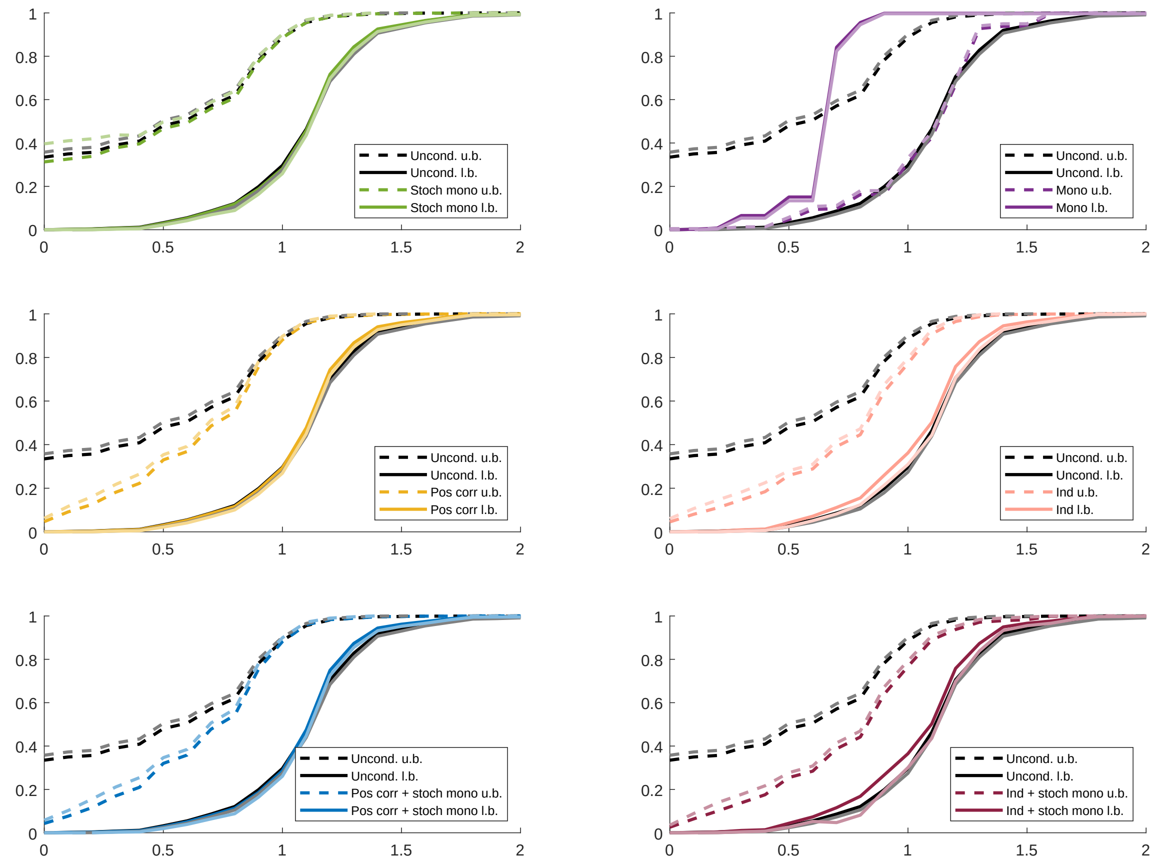

Notes: Bounds on F_S for the most popular cell phone product. Top two panels show stochastic monotonicity bounds (left) and monotonicity bounds (right). Middle panels show positive correlation bounds (left) and independence bounds (right). Bottom panels show combined positive correlation + stochastic monotonicity bounds (left) and combined independence + stochastic monotonicity bounds (right). Every panel shows unconditional bounds for comparison. Upper bounds are dashed lines and lower bounds are solid lines. Faded lines represent 95% confidence bands (constructed via subsampling; see Appendix E) for the bounds represented in the corresponding color. All prices are scaled by product's reference price, and thus units on horizontal axis are fraction of the reference price.

bounds. Stochastic monotonicity implies that a seller with a high first offer is likely to have a high value. This assumption will lead to a tightening of the lower bound if, for example, in some instances of the game in which sellers have relatively high first offers the seller eventually ends up accepting a relatively *low* offer.[30] If instead sellers with high first offers

[30]This can happen due to the randomness in the buyer to which the seller is matched or other DGP features for later offers, which we make no assumptions about. A similar argument applies to the positive correlation bounds (left middle panel). These bounds rely on assuming the seller's value is stochastically increasing in the buyer's first offer conditional on the seller's first offer, and will only tighten the lower bound, say, if,

always end up accepting higher prices, the stochastic monotonicity assumption (even if true) will not tighten the bounds.

The monotonicity bounds (top right) illustrate the potential for crossings; here the lower bound lies entirely above the upper bound. As demonstrated in Section 3.5 and Appendix G, monotonicity can be violated by unobserved game-level heterogeneity, which can generate non-monotonicities in the relationship between S and P_1^S. This finding highlights the importance of our weaker assumption (stochastic monotonicity), which is not rejected by the data. In the bottom right panel of Figure 2, we display the tightest bounds for this product that do not cross (independence with stochastic monotonicity). Monte Carlo simulations in Appendix F demonstrate that any of these bounds — even the unconditional bounds — can be quite narrow or wide, depending on the DGP.

Figure 3: Bounds on Seller Distribution and Auto Accept/Decline Prices

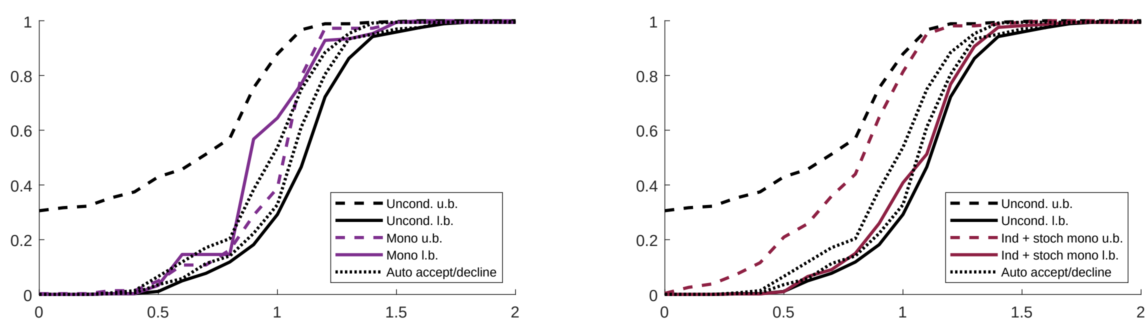

Notes: Bounds on F_S using only negotiations in which seller reported a non-zero auto-accept and auto-decline price. Left panel shows monotonicity bounds and right panel shows combined independence and stochastic monotonicity bounds, with upper bounds as dashed lines and lower bounds as solid lines. Every panel also shows the unconditional bounds for comparison. Empirical CDFs of auto-accept and auto-decline prices are shown in gray dotted lines. All prices are scaled by the reference price for the product, and thus units on the horizontal axis are fraction of the reference price.

For seller values, secret auto-accept and auto-decline prices serve as a novel validation check — private information of the seller that offers an immediate upper and lower bound on the true F_S.[31] Here we re-estimate F_S bounds limiting to the 363 negotiations of this

conditional on cases where $P_1^S = y$, sellers who face high P_2^B sometimes accept relatively low offers. If this is not the case, positive correlation, while not rejected by the data, will do little to improve the bounds.

[31] A given seller's auto-accept price is a weak upper bound on her value and the auto-decline price a weak lower bound. If, rather than alternating offers, the protocol involved only a take-it-or-leave-it offer by the buyer, the seller's optimal choice would be to set both the auto-accept and auto-decline prices equal to her value. The auto-accept/decline prices are not used anywhere in identifying or estimating our bounds; we

product in which the seller reported these secret thresholds. Figure 3 shows the results, along with empirical CDFs of auto-accept and auto-decline prices (gray dotted lines). The left panel shows the monotonicity bounds, which still cross. The right panel (independence with stochastic monotonicity) is our primary interest, as these are the tightest bounds we obtain without the (overly strong) assumption of seller monotonicity. Here we observe that our assumptions are consistent with the tight bounds implied by the secret prices: the lower bound implied by the auto-accept price CDF lies below our upper bound and the auto-decline price CDF lies above our lower bound.[32]

5.1.2. Bounds on Buyer Values. Figure 4 displays bounds on F_B for this same product. We find wider bounds for F_B than for F_S. The lower bound is not onto on $[0,1]$ because it depends on quit prices, similar to the F_S upper bound. The combined positive correlation and monotonicity bounds (bottom right) especially help to tighten the buyer bounds for this product. Unlike for F_S, the monotonicity bounds for F_B do not cross. The monotonicity assumption for B appears weaker, requiring only that, *conditional* on the seller's first offer, a buyer's value be higher at higher buyer offers (whereas the seller monotonicity assumption is particularly strong). We reiterate that all bounds are sharp, implying they are the best possible under their corresponding assumptions. Any tightening of the bounds necessarily requires stronger assumptions.

We can examine how the bounds in Figure 4 change — and whether they cross — when we impose a stronger assumption that raises the lower bound: the *list-price-recall condition* described in Section 3.2.[33] For this condition, recall that we incorporate into X_Q^B the concept that a buyer can always purchase at the list price, and therefore, under this condition, $B \leq P_1^S$. Under this condition, we should observe the buyer independence bounds crossing, because, as long as the marginal supports overlap, $B \leq P_1^S$ implies B and P_1^S are not indepen-

intentionally withhold that information to use it in this validation exercise.

[32]Cases where the auto-accept/decline prices are inconsistent with our bounds would be when the F_S lower bound lies above the auto-decline price CDF or the F_S upper bound lies below the auto-accept price CDF. We only find such crossings for bounds involving seller monotonicity. We repeat this exercise for the other product in our sample with at least 200 negotiations with auto-accept/decline prices and find a similar level of consistency with the auto accept/decline price distribution bounds.

[33]The list-price-recall condition only affects the F_B lower bound, the only bound that uses X_Q^B.

Figure 4: Bounds on Buyer Distribution for Cell Phone

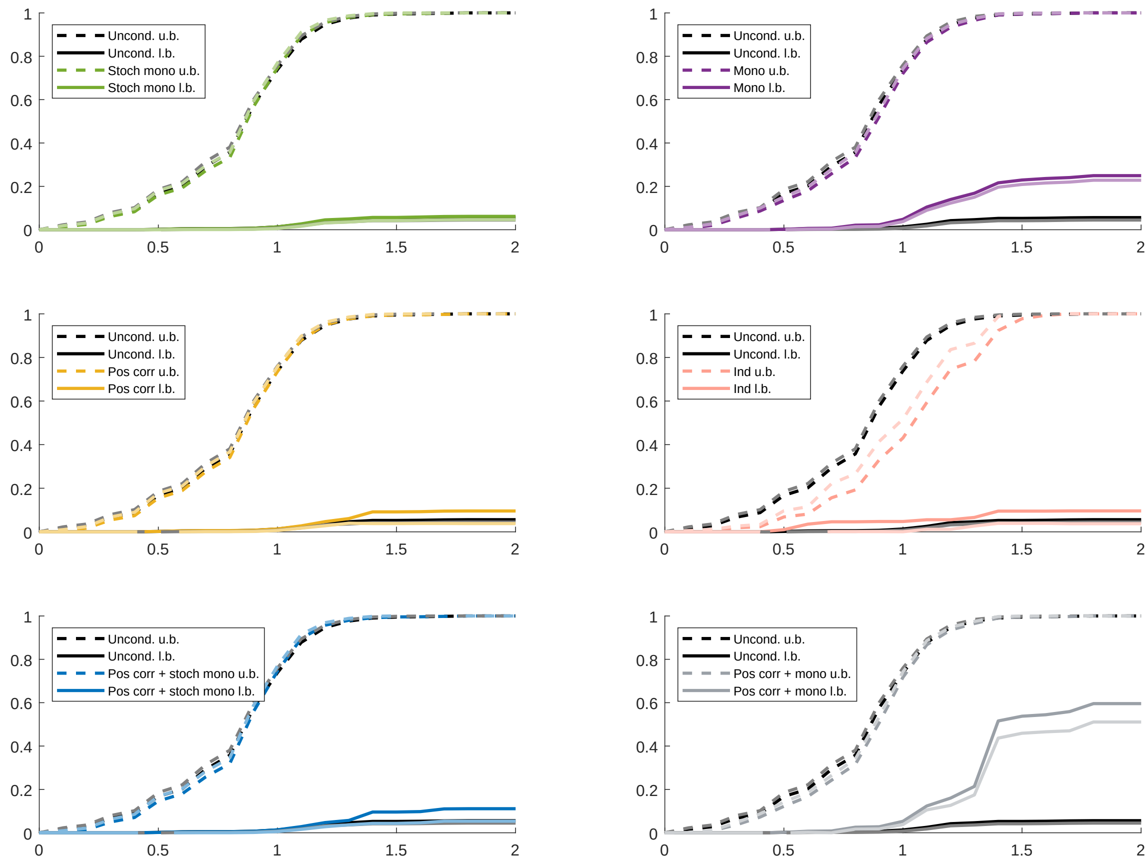

Notes: Bounds on F_B for the most popular cell phone product. Top two panels show stochastic monotonicity bounds (left) and monotonicity bounds (right). Middle panels show positive correlation bounds (left) and independence bounds (right). Bottom panels show combined positive correlation + stochastic monotonicity bounds (left) and combined positive correlation + monotonicity bounds (right). Every panel shows unconditional bounds for comparison. Upper bounds are dashed lines and lower bounds are solid lines. Faded lines represent 95% confidence bands (constructed via subsampling; see Appendix E) for the bounds represented in the corresponding color. All prices are scaled by product's reference price, and thus units on horizontal axis are fraction of the reference price.

dent. Figure 5 shows bounds incorporating this condition. We observe clear evidence of the independence bounds being violated: the lower bound lies nearly entirely above the upper. The unconditional bounds (black lines) are also much tighter under list-price-recall, and do not cross. Exploiting the list-price-recall condition also offers a check on the validity of the tightest bounds we obtained in Figure 4, those relying on monotonicity and positive correlation. Even applying the list-price-recall condition, these bounds do not cross. The fact that the positive correlation bounds do not cross, while the independence bounds do, highlights the benefit of the weaker assumption underlying these bounds.

Figure 5: Bounds on Buyer Distribution with List Price Recall

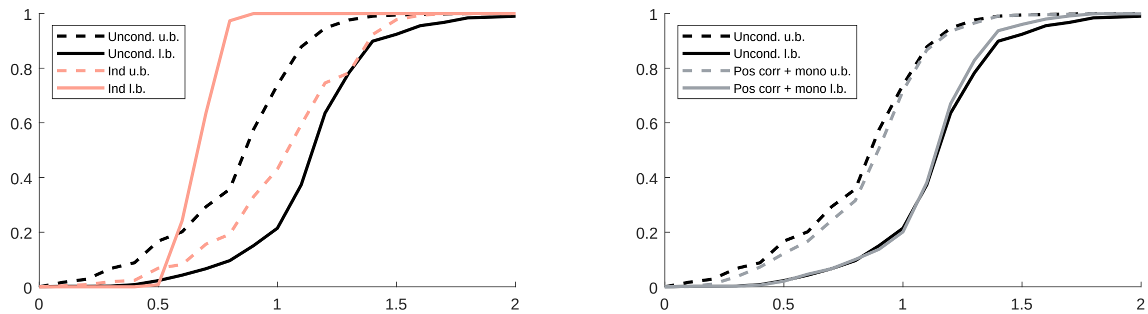

Notes: Bounds on F_B for the most popular cell phone product using the list-price-recall condition. Upper bounds are equivalent to those in Figure 4; only lower bounds change. Left panel shows independence bounds and right panel shows combined positive correlation and monotonicity bounds. Every panel shows unconditional bounds for comparison. Upper bounds are dashed lines and lower bounds are solid lines. All prices are scaled by product's reference price, and thus units on horizontal axis are fraction of the reference price.

5.2. Exploring All Products. We now summarize bounds for all 36 products. In Table 3, we show, under each assumption, the fraction of products for which bounds cross.[34] Crossings indicate a violation of the underlying assumption(s). Some crossings may be due to finite sample estimation error, and we account for this by testing whether any crossing is statistically significant (Frac. Reject); Appendix E describes details. We also compute the *integrated violation error* (IVE), which ranges from 0 to 1 and measures the average difference between the upper and lower bound in cases where they cross.[35]

Table 3 shows that the results from Figures 2 and 4 are representative of products in our data. In particular, we find that seller monotonicity bounds cross for 100% of products, with the upper monotonicity bound violating the lower monotonicity bound by an average of 23% (the IVE).[36] For buyer values, where the monotonicity assumption conditions on

[34]In practice we choose a grid on which to evaluate the bounds and check for crossings (upper bound lying below lower bound) at each grid point. We use 0 to 2.5 in increments of 0.1. To allow for machine rounding errors, we only consider violations exceeding 2e-10.

[35]For a generic upper and lower bound F^U and F^L, the IVE is given by $\int \max\{F^L(x) - F^U(x), 0\} dG(x)$, where the distribution function G is equal to the unconditional lower bound in the case of sellers and the unconditional upper bound in the case of buyers. We choose these distributions as they are surjective on $[0,1]$, as described in Section 5.1.

[36]Appendix C proves that the F_S monotonicity bounds are unchanged when combined with A3 or A5.

Table 3: Bound Crossings Under Different Assumptions

	Seller Bounds			Buyer Bounds		
	Frac. Cross	Frac. Reject	IVE	Frac. Cross	Frac. Reject	IVE
Unconditional (A1)	0	0	0	0	0	0
Monotonicity (A2)	1.00	1.00	0.23	0	0	0
Independence (A3)	0.08	0	0.00	0.42	0.11	0.006
Stochastic Monotonicity (A4)	0	0	0	0	0	0
Positive Correlation (A5)	0	0	0.00	0	0	0
Mon. + Indep. (A2 + A3)	1.00	1.00	0.23	0.78	0.47	0.07
Mon. + Pos. Corr (A2 +A5)	1.00	1.00	0.23	0.03	0	0.00
Stoch. Mon. + Indep. (A4 + A3)	0.19	0.03	0.00	0.42	0.08	0.006
Stoch. Mon. + Pos. Corr (A4 + A5)	0	0	0.00	0	0	0

Notes: Across all products in the estimation sample, table shows the fraction of products for which seller lower bound crosses the upper bound, fraction of products for which crossings are statistically significant, and the IVE (averaged across products). Table shows similar quantities for the buyer bounds.

P_1^S, the bounds never cross. The opposite is true for the independence bounds, which only cross for 8% of products for sellers (none of which are statistically significant) but for 42% of products when bounding F_B, with 11% having statistically significant crossings. In general, we observe that, even for bounds with some statistically significant crossings, the size of the crossings are small, as indicated by the low IVE.[37]

In Table 4, for a given set of bounds, we compute the average width of the bounds for a given product and take the mean, minimum, and maximum of these average widths across products.[38] The unconditional seller bounds can be relatively tight for some products, with an average probability gap of 0.340 for the minimum-width product, and quite wide for others, with an average gap of 0.524 for the maximum product. The average width for some bounds — such as the stochastic monotonicity bounds — is quite similar to that of the unconditional bounds. Given that all of the bounds are sharp under their corresponding assumptions, these widths could only be reduced under stronger assumptions. We find

[37] Imposing the list-price-recall condition leads to additional products having statistically significant crossings only for F_B bounds involving independence (consistent with our findings in Section 5.1.2).

[38] For a given type of bounds for a given product, we integrate the upper bound minus the lower bound against the density of the unconditional lower bound in the case of seller values and the unconditional upper bound in the case of buyer values, as with the IVE. This average width metric is similar to the IVE, ranging from 0 to 1, with a lower number meaning the bounds are tighter.

Table 4: Statistics Across Products on Width of Bounds

	Seller Bounds			Buyer Bounds		
	Min	Mean	Max	Min	Mean	Max
Unconditional (A1)	0.340	0.416	0.524	0.408	0.428	0.459
Monotonicity (A2)	–	–	–	0.292	0.367	0.429
Independence (A3)	0.173	0.281	0.433	0.097	0.230	0.341
Stochastic Monotonicity (A4)	0.334	0.413	0.516	0.399	0.418	0.451
Positive Correlation (A5)	0.277	0.369	0.511	0.367	0.413	0.437
Mon. + Indep. (A2 + A3)	–	–	–	0.089	0.217	0.324
Mon. + Pos. Corr (A2 +A5)	–	–	–	0.240	0.343	0.398
Stoch. Mon. + Indep. (A4 + A3)	0.171	0.269	0.380	0.109	0.238	0.346
Stoch. Mon. + Pos. Corr (A4 + A5)	0.271	0.369	0.499	0.364	0.407	0.432

Notes: Table shows the minimum, mean, and max (across products) of the average width of the bounds, where the average width for a given product is computed by the upper minus lower bound integrated against the density of the unconditional lower bound for sellers or the unconditional upper bound for buyers. These statistics are computed for a given set of bounds only for products for which the bounds do not cross (thus, the width of bounds and any bounds relying on seller monotonicity is omitted, as these always cross).

that stochastic monotonicity and positive correlation assumptions for the seller improve the bounds for some products, decreasing the minimum average width to 0.271. For buyer bounds, monotonicity and positive correlation improve tightness relative to the unconditional bounds, with a minimum average width of 0.240.

6 Quantifying Inefficient Impasse

6.1. Motivation. We now consider bounds on the counterfactual first-best trade probability. In a first-best world, a buyer with value B and seller with value S will trade whenever $B \geq S$. The first-best is unattainable when buyers and sellers have independent private values with overlapping supports (Myerson and Satterthwaite 1983) and may be attained under certain conditions with correlated values (McAfee and Reny 1992). The first-best quantity of trade will thus be weakly higher than the realized volume of trade in the data; the empirical question is, how much higher? A lower bound of $P(B \geq S)$ can be compared to the sale probability in the data to bound the degree of inefficient impasse occurring in reality. As important as the object $P(B \geq S)$ is for quantifying inefficiency, existing empirical tools and theoretical models are insufficient for identifying it, and data from bargaining settings where agents have private information will not typically contain data on those private values themselves (S and B) except in lab experiments.

As explained in Section 3.2, we allow for the possibility that buyers and sellers may have a continuation value after failed trades (and may, for example, re-enter the marketplace). Buyer/seller values are then interpretable as willingness to pay/sell. This interpretation is not problematic for studying inefficient impasse because we are quantifying, for a fixed set of buyer-seller pairs, how well real-world bilateral bargaining performs relative to the first-best. We do not model or study the process by which those pairs came to be matched, which would be necessary to study the efficiency of the searching and matching process. Both types of efficiencies — the efficiency of the search and matching process and the efficiency of the incomplete-information bilateral bargaining conditional on matching — are important components of the efficiency of the market as a whole. Given the complexities of incomplete-information bargaining (a continuum of equilibria with no full characterization; Ausubel et al. 2002), this paper focuses only on the latter: the inefficiency of the bargaining *conditional* on the buyer-seller pairs matched to one another in the data.

It also bears emphasis that the goal of our exercise is not to construct a policy counterfactual — indeed, the counterfactual we examine, the first-best, may not even be feasible (Myerson and Satterthwaite 1983) — but rather to quantify how well real-world bargaining performs relative to this benchmark. We hope this quantification can inform several audiences. For the platform, it may offer insights into how much value is left on the table in the current sales mechanism and whether it might be worth investigating alternatives, such as encouraging the use of the auto accept/decline feature or encouraging more communication between agents (which has been shown in lab experiments to reduce inefficient impasse in incomplete-information bargaining; Valley et al. 2002). For empirical bargaining researchers, who typically model negotiated prices as arising from a protocol with no inefficient impasse (some form of Nash bargaining; e.g., Crawford and Yurukoglu 2012), this quantification exercise may offer some insights into how reasonable such an abstraction may be, albeit only within the eBay platform. For bargaining theory, we hope the exercise can motivate future models capturing elements of inefficient impasse and unobserved heterogeneity.

6.2. Bounds on Surplus. In this section we construct bounds on $P(B-S \geq x)$, the distribution of the surplus $B-S$, and use these to bound $P(B \geq S)$, which is $P(B-S \geq x)$ evaluated at $x=0$. Two of these bounds follow immediately from steps in Section 3. The other two are new, building on ideas from the marginal distribution bounds in Section 3.

Our weakest bounds for $P(B-S \geq x)$ rely only on revealed preferences (A1), which implies $P(B-S \geq x) \geq P(X^B_{AC} - X^S_{AC} \geq x)$. Buyers will generally not accept or counter at a price that is strictly higher than the seller (as this would be strange behavior indeed). Therefore, $P(X^B_{AC} - X^S_{AC} \geq x)$ typically corresponds to $P(X^B_{AC} - X^S_{AC} = x)$, which, at $x=0$, is equal to $P(sale)$ (the sale probability in the data), as this represents cases where one agent accepts a price the other proposes. The upper bound is similar: A1 implies $P(B-S \geq x) \leq P(X^B_Q - X^S_Q \geq x)$. At $x=0$, this latter probability is always equal to 1, because only one party (the buyer or seller) can quit in a given negotiation. When the seller quits, $X^B_Q = \infty$, and when the buyer quits, $X^S_Q = 0$. Thus, revealed preferences alone yield uninformative bounds; the most we learn is that $P(B \geq S) \in [P(sale), 1]$.

We next consider bounds that rely on buyer monotonicity (A2.ii) as well as the following weak assumption on $B-S$:

Assumption A6. *(Surplus stochastic monotonicity).* $P(B-S \geq x \mid P^S_1 = y, P^B_2 = z)$ *is increasing in* z *for all* y.

To interpret this assumption, consider two negotiations, both of which involve $P^S_1 =$ \$300. Suppose that in the first negotiation $P^B_2 =$ \$200 and the parties trade, while in the second negotiation $P^B_2 =$ \$290 and trade fails. A6 implies that gains from trade likely exist in the second negotiation, even though trade fails.[39] A6 is akin to our stochastic

[39] Sufficient (but not necessary) conditions for A6 are a strict version of A2.ii (buyer monotonicity) and A3.i (seller independence), two assumptions for which we do not find large crossings in Section 5. To see that these assumptions are sufficient, suppose we can write $P^B_2 = f(B, P^S_1)$ where $f(\cdot, P^S_1)$ is increasing for all P^S_1 with inverse function $g(\cdot, P^S_1)$. Then the conditional probability statement in A6 can be written $P(B-S \geq x \mid P^S_1 = y, B = g(z,y)) = P((z,y) - S \geq x \mid P^S_1 = y, B = g(z,y))$. This latter statement is equivalent to $P(g(z,y) - S \geq x \mid P^S_1 = y)$, which is increasing in z for all y.

monotonicity assumption on *values* applied instead to the *difference in values*.[40] Combined with buyer monotonicity, it yields the following:

$$P(B-S\geq x)\geq \int \max_{z'\leq z} P(X^{B*}_{AC}(y,z)-X^{S}_{AC}\geq x\mid P^S_1=y,P^B_2=z')dF_{P^S_1,P^B_2}(y,z) \tag{13}$$

$$P(B-S\geq x)\leq \int \min_{z'\geq z} P(X^{B*}_{Q}(y,z)-X^{S}_{Q}\geq x\mid P^S_1=y,P^B_2=z')dF_{P^S_1,P^B_2}(y,z) \tag{14}$$

Theorem 6. *(13) gives a sharp lower bound for $P(B-S\geq x)$ under A1.i, A1.iii, A2.ii, and A6. (14) gives a sharp upper bound for $P(B-S\geq x)$ under A1.ii, A1.iv, A2.ii, and A6.*

A stronger variant of the last assumption is the following:

Assumption A7. *(Surplus weak monotonicity). $\overline{supp}(B-S\mid P^S_1=y,P^B_2=z)\leq \underline{supp}(B-S\mid P^S_1=y,P^B_2=z')$ for all y and for all $z\leq z'$.*

Revisiting the example discussed after A6, A7 implies that gains from trade definitely (rather than only likely) existed in the second negotiation. While stronger than A6, A7 is weaker than assuming monotonicity for *both* the buyer and seller.[41] As A6 is akin to a stochastic monotonicity assumption applied to $B-S$, A7 is akin to weak monotonicity. Combined with buyer monotonicity, A7 yields the following:

$$P(B-S\geq x)\geq \int \mathbf{1}(X^{B*-S}_{AC}(y,z)\geq x)dF_{P^S_1,P^B_2}(y,z) \tag{15}$$

$$P(B-S\geq x)\leq \int \mathbf{1}(X^{B*-S}_{Q}(y,z)\geq x)dF_{P^S_1,P^B_2}(y,z) \tag{16}$$

where $X^{B*-S}_{AC}(y,z)=\overline{\text{supp}}(X^{B*}_{AC}(y,z)-X^S_{AC}:P^B_2\geq z,P^S_1=y)$ and $X^{B*-S}_{Q}(y,z)=\underline{\text{supp}}(X^{B*}_{Q}(y,z)-X^S_{Q}:P^B_2\leq z,P^S_1=y)$.

Theorem 7. *(15) gives a sharp lower bound for $P(B-S\geq x)$ under A1.i, A1.iii, A2.ii, and A7. (16) gives a sharp upper bound for $P(B-S\geq x)$ under A1.ii, A1.iv, A2.ii, and A7.*

Finally, the strongest assumptions we consider for bounding $P(B-S\geq x)$ rely on monotonicity (A2) for both the buyer *and* seller, implying that, conditional on $P^S_1=y$ and $P^B_2=z$,

[40] Note that we do not rely on assumptions about the correlation structure between S and B, as these are unhelpful when examining the *difference* $B-S$.

[41] Strict versions of A2.i and A2.ii are sufficient for A7 to hold, but not necessary. Specifically, suppose we can write $P^S_1=f_1(S)$ and $P^B_2=f_2(B,P^S_1)$ where $f_1(\cdot)$ and $f_2(\cdot,P^S_1)$ are increasing with inverse functions $g_1(\cdot)$ and $g_2(\cdot,P^S_1)$, respectively. Then $B-S=g_2(P^B_2,P^S_1)-g_1(P^S_1)$ and, conditional on P^S_1, this latter difference is an increasing function of P^B_2.

$X_{AC}^{B*}(y,z) - X_{AC}^{S*}(y) \leq B - S \leq X_{Q}^{B*}(y,z) - X_{Q}^{S*}(y)$. Section 5 shows that seller (but not buyer) monotonicity is rejected by the data. We nonetheless estimate these bounds to illustrate that they are indeed too strong and can cross. Estimators for all of our bounds on $P(B - S \geq x)$ are similar to those for F_S and F_B and are discussed in Appendix E.

6.3. Inefficient Impasse Results. We evaluate bounds on $P(B - S \geq x)$ at $x = 0$, or, equivalently, $P(B \geq S)$, the first-best trade probability. We are primarily interested in the lower bound on $P(B \geq S)$ but we evaluate the upper bound as well to look for crossings.

We display estimates of lower bounds on $P(B \geq S)$, along with confidence intervals on these lower bounds, under these assumptions in Table 5, using the most popular product in each major category as in Table 1. For the electronics product, $P(sale)$ is 0.441 and the lower bound on the counterfactual $P(B \geq S)$ under surplus stochastic monotonicity and buyer monotonicity (column 2) is 0.467, suggesting that the real-world bargaining misses some efficient trades. However, the confidence interval on this 0.467 contains $P(sale)$, and thus the evidence of inefficient impasse under this set of assumptions is relatively weak. Column 3 shows bounds relying on surplus weak monotonicity and buyer monotonicity. Here we find confidence intervals that lie above $P(sale)$ for each product, suggesting that, maintaining these assumptions, bargaining is indeed inefficient. The column 4 bounds are higher still, but these invoke seller monotonicity and hence are rejected by the data.

We can construct a measure of *inefficient impasse*, which we define as the fraction of cases where gains from trade exist and yet trade fails, or $P(\textit{no sale} \mid B \geq S) = 1 - P(sale)/P(B \geq S)$. A lower bound on this quantity is obtained by plugging in a lower bound on $P(B \geq S)$. The bounds in Table 5 suggest that the real-world bargaining for these products exhibits inefficient impasse, but not as much as implied by the (overly strong) assumption of seller monotonicity. For example, for the popular cell phone product, the sale probability in the data is 0.276, but in column 3 we find that it would be as high as 0.508 in a first-best world, suggesting that, when the buyer values the phone more than the seller, the pair still fails to reach an agreement 45.6% of the time (i.e., $1 - 0.276/0.508 = 0.456$).

Figure 6 extends this analysis to all 36 products. In each panel, we order products on

Table 5: Lower Bounds on First-Best Trade Probability for Most Popular Products

	n	(1) $P(sale)$	(2) Surplus Stoch Mon, Buyer Mon	(3) Surplus Weak Mon, Buyer Mon	(4) Seller Mon, Buyer Mon
Consumer Electronics	401	0.441	0.467	0.701	0.853
		[0.393,0.490]	[0.406,0.532]	[0.606,0.756]	[0.771,0.895]
Video Games/Consoles	330	0.427	0.436	0.591	0.724
		[0.374,0.481]	[0.476,0.604]	[0.483,0.650]	[0.640,0.779]
Cell Phones/Accessories	1,147	0.276	0.270	0.508	0.889
		[0.250,0.302]	[0.241,0.302]	[0.450,0.540]	[0.870,0.912]
Computers/Tablets/Networking	318	0.368	0.363	0.604	0.742
		[0.315,0.421]	[0.310,0.419]	[0.494,0.661]	[0.678,0.828]

Notes: For the most popular products with each category, table displays lower bounds on $P(B \geq S)$ under different assumptions. 95% confidence intervals (obtained via subsampling) are shown in square braces below each estimate.

the horizontal axis by their sale probability in the data. On vertical axes, an "×" represents the estimated upper bound and a hollow circle the lower bound.[42] Panel A displays bounds relying on surplus stochastic monotonicity and buyer monotonicity. As in Table 5, these assumptions are too weak to yield informative bounds, as they are close to $P(sale)$ and 1. In panel C, we find that assuming monotonicity for both agents yields bounds that are tighter, but violated for three products (highlighted with purple circles).[43] These crossings are expected, as seller monotonicity bounds cross for all products (Table 3). The Goldilocks-like assumptions are in panel B (surplus weak monotonicity and buyer monotonicity), where we observe informative bounds that do not cross. The lower bounds suggest that the real-world bargaining indeed exhibits inefficient impasse, ranging from 18.0% to 54.2% across products, with a median of 37.3%.

As discussed in Section 3.7, the choice of which bounds to favor should be guided by the details of a particular setting, where possible. We favor the bounds in Figure 6.B because both buyer monotonicity and surplus weak monotonicity allow us to condition on the

[42] To facilitate uncovering possible violations, we tighten the upper bound by invoking the list-price-recall condition. This condition only affects the upper bound on $P(B \geq S)$ because only the upper bound is related to X_Q^B. The gain from invoking this condition is small here, lowering the upper bound only by 3.1 percentage points for the average product in panel B and by less than 0.4 percentage points in panels A and C.

[43] Figure 6 only shows bounds on $P(B - S \geq x)$ at $x = 0$; the bounds may cross at other x as well.

Figure 6: Bounds on $P(B \geq S)$, All Products

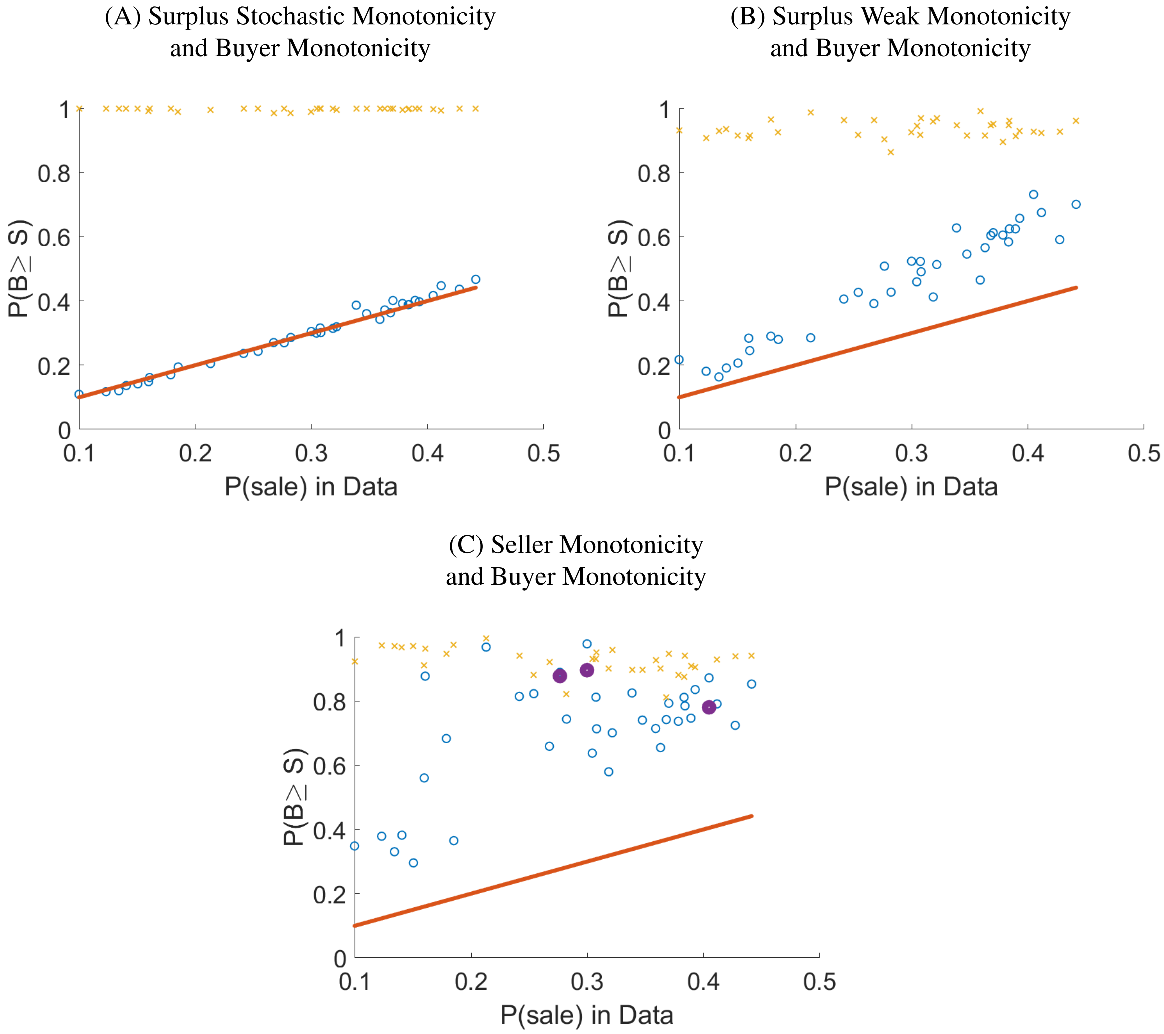

Notes: Figure display upper bounds (marked with "×" and lower bounds (marked with hollow circles) for the counterfactual first-best trade probability ($P(B \geq S)$) under different assumptions, for each product in the estimation sample. Horizontal axes rank products by sale probability in the data. The solid line represents the 45-degree line. Solid purple dots in panel C highlight the lower bound for products where bounds cross.

seller's first offer, allowing us to potentially condition on some unobserved heterogeneity, which seems wise in the eBay setting. We have additional confidence that these assumptions may be capturing accurate general properties of eBay bargaining in that, like the buyer monotonicity bounds alone (Table 3), they consistently do not cross for any product.

6.4. When is Bargaining More Efficient? We now evaluate inefficient impasse, $P(\textit{no sale} \mid$

$B \geq S$), separately in various cuts of the data in which agents, items, or other aspects of the negotiation satisfy particular conditions (such as whether or not the agents exchange any messages). We rely on our lower bounds based on surplus weak monotonicity and buyer monotonicity (Theorem 7). For each condition we consider, we limit the analysis to products for which we observe at least 100 negotiations where the condition is satisfied and 100 where it is not. We then evaluate the difference in the inefficient impasse lower bound between those two subsets. This exercise is exploratory and speculative in nature. We have no model for *why* bargaining would be more efficient under certain conditions. Our goal here is to provide a number of findings that may motivate future investigation. A final important caveat for this analysis is that apparent differences in efficiency may be driven by unobservable selection. For example, sellers who select to use the auto-accept feature (which we analyze below) may be those who are already more prone to agree in bargaining, independent of the auto-accept feature.

The results are shown in Table 6. We introduce this exercise first by comparing observations in which agents exchanged some text communication during the negotiation.[44] The first row of Table 6 shows that the inefficient impasse lower bound among observations with communication (the "Yes" column) is 0.349. This bound is higher, 0.394, among observations without communication (the "No" column). The difference is -0.045, which is not statistically significant.

Because we only have lower bounds on inefficient impasse, we can only make limited claims about relative efficiency. To illustrate this, consider an inefficient impasse level of 0.37, which lies halfway between the inefficient impasse lower bound of the Yes and No groups. This level is achievable in observations with communication, as 0.37 lies above the inefficient impasse lower bound for that group. This is not true for observations without communication, where the true inefficient impasse is bounded below by 0.394. Thus, we

[44] Our data indicates whether some agent sends a message, not who sends it. Only two products contained at least 100 observations with and at least 100 without a message. Even for these products, messages were relatively rare (occurring in only 12.6% of observations). For the samples used in rows 2–4 of Table 6.A, the average product has 23.1% of observations with an eBay store, 54.0% corresponding to an U.S. buyer, and 32.0% using auto accept/decline prices

Table 6: Heterogeneity in Inefficient Impasse Lower Bound

A. Inefficient Impasse Lower Bound Where Condition Met vs. Not							
Condition	Yes	S.E.	No	S.E.	Diff.	S.E.	# Prod
Has $\geq$ One Message	0.349	(0.0381)	0.394	(0.0165)	-0.045	(0.0415)	2
Seller is an eBay Store	0.346	(0.0222)	0.387	(0.0125)	-0.042	(0.0255)	6
Buyer is from U.S.	0.298	(0.0099)	0.341	(0.0304)	-0.043	(0.0319)	19
Has Auto Accept/Decline	0.302	(0.0203)	0.360	(0.0147)	-0.058	(0.0250)	35
High Num. of Photos	0.320	(0.0112)	0.329	(0.0118)	-0.010	(0.0163)	32
High Seller Average Rating	0.321	(0.0107)	0.323	(0.0113)	-0.002	(0.0156)	36
High Seller Num. of Reviews	0.337	(0.0110)	0.296	(0.0109)	0.041	(0.0154)	35
High Seller Experience	0.329	(0.0115)	0.326	(0.0104)	0.003	(0.0155)	8
High Buyer Experience	0.355	(0.0110)	0.314	(0.0105)	0.040	(0.0152)	24
B. Inefficient Impasse Lower Bound Where Condition Met vs. Obs. with Low Seller & Low Buyer Experience							
High Seller/High Buyer Exp	0.322	(0.0233)	0.326	(0.0219)	-0.004	(0.0320)	5
High Seller/Low Buyer Exp	0.334	(0.0234)	0.326	(0.0219)	0.008	(0.0321)	5
Low Seller/High Buyer Exp	0.353	(0.0240)	0.326	(0.0219)	0.027	(0.0325)	5
C. Inefficient Impasse Lower Bound For Products Where Condition Met vs. Not							
New Product (vs. Used)	0.266	(0.0077)	0.385	(0.0605)	-0.119	(0.0610)	2
High Reference Price Prod	0.359	(0.0085)	0.362	(0.0085)	-0.003	(0.0120)	36

Notes: Analysis using various subsamples. For each row in panels A and B, we limit to products with at least 100 negotiations where the condition is satisfied and 100 where it is not, resulting in fewer than 36 products for many conditions. This is true even for conditions relying on within-product medians ("High" rows), because some products have multiple observations at the median, resulting in fewer than 100 observations weakly greater than (or fewer than 100 strictly less than) the median. We compute inefficient impasse ($P(no\ sale \mid B \geq S) = 1 - P(sale)/P(B \geq S)$) within the subsample where the condition is satisfied ("Yes" columns) vs. not ("No" columns) as well as the difference between these values in the two subsamples for a given product. The table reports averages of these quantities across products, with standard errors on these differences shown in the "S.E." columns, computed via the delta method. The final column shows the number of products available for a given condition. Panel B measures differences relative to negotiations with low seller and low buyer experience. Panel C instead computes differences between products that are used vs. new (using the 2 product identifiers that appear in our sample both as used and new products) in the first row. Panel C second row compares differences between products with a reference price above the the median (across products) vs. below the median.

can thus infer that there exists a range of inefficient impasse (levels between 0.349 and 0.394) that *may* be reached in negotiations of the Yes group and that *cannot* be reached in negotiations of the No group. These comparisons are only suggestive of possibilities; we cannot reach strong conclusions as to which group has a higher true degree of inefficient impasse because the truth for both groups may exceed their respective lower bounds substantially; the ranking of the true inefficient impasse between the two groups is not pinned down. Nevertheless, for simplicity, we will still say, for example, that a group exhibits "more efficiency" or "has less inefficient impasse" when its lower bound is higher.

This particular finding about communication, though insignificant, is consistent with laboratory experiments in Valley et al. (2002) and with descriptive evidence from a natural experiment on eBay's Germany site (Backus et al. 2023) in which the sale probability increased when the site began allowing communication. Importantly, here we find only a correlation, not a causal statement, as agents choose whether to convey such messages.

Table 6 shows an insignificant difference in the inefficient impasse bound in negotiations with the seller being an eBay store vs. not.[45] We find an insignificant difference of a similar size when comparing buyers who are in the U.S. vs. not.[46] We do find a significant difference (an decrease in the bound of 5.8 percentage points) when the seller reports auto-accept/decline prices. This feature eliminates the need for the seller to consider very low/high offers, saving time. Our results suggest that this commitment of the seller *a priori* to an acceptable range of offers may reduce inefficient impasse.[47]

The remaining conditions in panel A of Table 6 refer to "high" vs "low" characteristics, e.g., a high number of photos in the listing. We define a listing to have a high number of photos if it has weakly more photos than the median listing for that product. Other "high" vs. "low" conditions are defined analogously.[48] Listings with more photos have

[45] A *store* is a status larger sellers may pay for, giving them access to special marketing tools.

[46] Previous work has documented mixed results comparing negotiations in the U.S. and elsewhere. Roth et al. 1991, studying lab ultimatum games, found higher offers and lower acceptance rates in the U.S. than in Japan or Israel. Keniston et al. 2024 showed that agents have preferences for splitting the difference between previously proposed bargaining offers in the U.S., Spain, and India, and in cross-country tariff negotiations.

[47] A seller with a realized value and first offer of (s, p_1^S) who reports an auto-accept price of $p^{auto-accept}$ commits to automatically accept buyer offers in $[p^{auto-accept}, p_1^S)$, whereas, if the seller does not report an auto-accept price, we know only that she commits to accept offers at or above p_1^S (as p_1^S is the Buy-It-Now price). Without the auto-accept feature, if this seller receives an offer in $[p^{auto-accept}, p_1^S)$, she could be tempted to make a counteroffer that has a non-zero probability of being rejected by the buyer, potentially ending in breakdown. Thus, the auto-accept feature may aid with efficiency by committing the seller — vis-a-vis her own future self — to an action that results in more trade. This efficiency improvement might be undone, however if the seller's choice of an *auto-decline* price, $p^{auto-decline}$, results in additional scenarios of no trade at offers in $[s, p^{auto-decline}]$ that the seller might have accepted absent the auto-decline price. We leave a more detailed investigation of the auto accept/decline feature to future work.

[48] The number of photos is a choice of the seller; eBay requires at least one. The median number of photos (across observations within a product) is 3.2 photos for the average product.

an inefficient impasses lower bound that is 1.0 percentage point lower, consistent with information about product quality improving efficiency, but this is imprecisely estimated.[49]

The remaining variables in Table 6 refer to the seller's average rating, the seller's number of reviews, and the buyer and seller experience. An agent's experience is the count of all eBay negotiations, regardless of whether trade resulted, that involved that agent and that occurred prior to the current negotiation.[50] We find that a higher quantity of reviews or higher buyer experience are associated with a *higher* inefficient impasse lower bound. These differences are precisely measured (and we find no difference based on the level of the seller's ratings or experience).

A possible explanation for greater apparent inefficiency accompanying more reviews/more experience lies in an interesting feature of incomplete-information bargaining settings. In such settings, there is a clear trade off between efficiency and rent extraction (Myerson and Satterthwaite 1983; Loertscher and Marx 2022).[51] If additional experience (or reviews) gives an agent more power to extract rents from her opponent, and agents exploit this power, we would expect additional experience to harm efficiency, consistent with Table 6. In panel B of Table 6, we peel back these experience results further, reporting the difference between the inefficient impasse lower bound in observations with different combinations of high and low experience. In each case, the difference is relative to negotiations between low-experience sellers and low-experience buyers. We find the largest difference when the buyer is experienced and the seller is not, but this difference is not precisely measured.

Finally, panel C of Table 6 compares inefficiency *across* products (rather than within) based on certain characteristics. We find that the inefficient impasse lower bound is 11.9

[49]Backus et al. (2020) documented that the first buyer offer (on a given listing) arrives more quickly to listings with more photos, arguing that the additional photos may reduce asymmetric information, a problem that reduces efficiency in some models (Deneckere and Liang 2006). Lewis (2011) similarly associates an increase in the number of photos on eBay listings with a reduction in asymmetric information.

[50]The median cutoff conditions defining "High" realizations of these variables are, for the average product, a 99.8 rating (out of 100), 849.8 reviews, 324.2 previous negotiations for the seller, and 13.9 previous negotiations for the buyer.

[51]Under incomplete information, giving one agent more bargaining power increases her payoff but also increases deadweight loss, reducing the size of the total pie. In a complete-information settings (Nash bargaining, say), an increase in the bargaining power of one agent does not imply any change to total surplus.

percentage points lower for negotiations over new products than used products, consistent with bargaining being more efficient for new products, and this difference is nearly significant at the 0.05 level (with a t-stat of 1.95).[52] We find a small and insignificant difference between higher-reference-price products relative to lower-reference-price products.

7 Discussion and Conclusion

This study provides bounds on the private-value distributions of buyers and sellers and on the first-best trade probability from sequential-offer bargaining data on eBay. The bounds are sharp and nonparametric. We rely on revealed preferences and other assumptions on behavior or information without specifying a full model of equilibrium play. Bounds relying on our strongest assumptions (monotonicity of sellers' first offers and independence between buyer's values and seller's first offers) can cross. While these strong assumptions are satisfied in equilibria of two-sided incomplete-information bargaining games analyzed in the theoretical literature, they appear too strong for empirical work. This underscores the importance of our more moderate assumptions, which allow for empirical features such as unobserved game-level heterogeneity.

Our approach circumvents a major theoretical problem arising in bargaining games of incomplete information: signaling. Each action taking by a player signals information to the opposing player, yielding a multiplicity (even a continuum) of equilibria that are qualitatively very different depending on how off-equilibrium beliefs are specified (see Ausubel et al. 2002). The bounds we propose do not rely on any specification of beliefs, equilibrium refinements, or equilibrium selection, allowing us to study how well bargaining performs in this real-world market without strongly constraining the answer a priori.

Given that the bounds rely on assumptions that, in many cases, are quite weak, they can naturally be wide. We show that, in spite of this, the bounds can highlight which

[52]This comparison uses the single bar codes that appear in our sample both as a new and used item; that is, we observe at least 200 negotiations for used versions and at least 200 for new versions of this bar code. Because we define a *product* as a bar-code-condition-type pair (where condition type means used vs. new), as described in Section 2, there are two *products* in the analysis of the first row of Table 6

behavioral properties are consistent with real-world bargaining and allow us to quantify inefficient impasse — a question that is not possible to address under complete-information frameworks such as Nash bargaining. Under our preferred assumptions, we find evidence of inefficient impasse: for the median product, at least 37% of failed trades are cases where positive trade gains exist. Thus, viewing consumer negotiations in this market through the lens of a complete-information model would be incorrect. We also find that it would be misleading to impose too strong of an assumption on behavior. Though satisfied by existing theoretical models, the strongest assumptions we explore would suggest that inefficient breakdown is far more prevalent. We are able to falsify these stronger assumptions.

It is possible that the assumptions we use — even those that are not the strongest — are still too strong. To assess this, there are relatively few empirical analyses of bargaining under incomplete information to which ours can be directly compared — and none, to our knowledge, from real-world negotiations involving consumers. However, several studies offer useful comparisons. First, Valley et al. (2002) studied laboratory participants in a two-sided incomplete-information bargaining game, and found that participants fail to trade in 46% of cases where trade gains exist. They found that this impasse is reduced substantially (to 15%) when negotiators are allowed to communicate. Bochet et al. (2023) and Huang et al. (2023) also studied two-sided incomplete-information experimentally, finding a corresponding level of inefficient impasse of 30% and 17%, respectively. In a structural model, Larsen (2021), studying professionals negotiating over used-car inventory, found that at least 21% of first-best trades fail, and Larsen et al. (2022) found that skilled mediators substantially reduce this inefficiency.[53] Relative to these numbers, our estimates of inefficient impasse for the median product suggest that the performance of bargaining in real-world consumer settings is in the ballpark of (but perhaps more inefficient than) those involving laboratory participants or business-to-business negotiations. We see our findings as benchmarks to which future studies of bargaining in various contexts may be compared.

[53] These numbers are found in (or can be constructed from) Table 1 of Valley et al. (2002), Figure 3 of Bochet et al. (2023), Table 2 of Huang et al. (2023), and Table 3 of Larsen (2021). Note that in Huang et al. (2023), gains from trade always exist, which is not the case for the other experimental studies.

References

Ambrus, A., Chaney, E., and Salitskiy, I. (2018). Pirates of the Mediterranean: An empirical investigation of bargaining with asymmetric information. *Quantitative Economics*, 9(1):217–246.

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

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

Backus, M., Blake, T., Pettus, J., and Tadelis, S. (2023). Communication, learning, and bargaining breakdown: An empirical analysis. *Management Science*, forthcoming.

Binmore, K., Osborne, M., and Rubinstein, A. (1992). Noncooperative models of bargaining. *Handbook of Game Theory*, 1:179–225.

Blundell, R., Gosling, A., Ichimura, H., and Meghir, C. (2007). Changes in the distribution of male and female wages accounting for employment composition using bounds. *Econometrica*, 75(2):323–363.

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

Bodoh-Creed, A., Boehnke, J., and Hickman, B. (2021). How efficient are decentralized auction platforms? *Review of Economic Studies*, 88(1):91–125.

Chernozhukov, V., Lee, S., and Rosen, A. M. (2013). Intersection bounds: Estimation and inference. *Econometrica*, 81(2):667–737.

Chesher, A. and Rosen, A. M. (2017). Generalized instrumental variable models. *Econometrica*, 85(3):959–989.

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

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.

Desai, P. S. and Jindal, P. (2024). Getting a break in bargaining: An upside of time delays. *Marketing Science*, forthcoming.

Frandsen, B. R. and Lefgren, L. J. (2021). Partial identification of the distribution of treatment effects with an application to the Knowledge is Power Program (KIPP). *Quantitative Economics*, 12(1):143–171.

Freyberger, J. and Larsen, B. J. (2022). Identification in ascending auctions, with an application to digital rights management. *Quantitative Economics*, 13:505–543.

Freyberger, J. and Larsen, B. J. (2024). Supplementary appendix for "How well does bargaining work in consumer markets? A robust bounds approach". `https://arxiv.org/abs/`.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. Cambridge: MIT Press.

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

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*, 56(3):601–611.

Haile, P. A. and Tamer, E. (2003). Inference with an incomplete model of english auctions. *Journal of Political Economy*, 111(1):1–51.

Hendricks, K., Sorensen, A., and Wiseman, T. (2021). Dynamics and efficiency in decentralized online auction markets. NBER Working Paper 25002.

Hirano, K. and Porter, J. R. (2012). Impossibility results for nondifferentiable functionals. *Econometrica*, 80(4):1769–1790.

Ho, K. and Rosen, A. (2017). Partial identification in applied research: Benefits and challenges. In Honore, B., Pakes, A., Piazzesi, M., and Samuelson, L., editors, *Advances in Economics and Econometrics: Eleventh World Congress (Volume 11)*, pages 307–359. Cambridge University Press, Cambridge, UK.

Huang, J., Kessler, J. B., and Niederle, M. (2023). Fairness concerns have less impact when agents are less informed. *Experimental Economics*, forthcoming.

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

Keniston, D., Larsen, B., Li, S., Prescott, J., Silveira, B., and Yu, C. (2024). Fairness in incomplete-information bargaining: Theory and widespread evidence from the field. *International Economic Review*, forthcoming.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies*, 78(1):293–327.

Larsen, B., Lu, C., and Zhang, A. L. (2022). Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations. NBER Working Paper 29159.

Larsen, B. and Zhang, A. (2021). Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry. Working paper, Stanford.

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

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of eBay motors. *American Economic Review*, 101(4):1535–1546.

Li, H. and Liu, N. (2022). Nonparametric identification and estimation of double auctions with bargaining. Working paper, Xiamen University.

Loertscher, S. and Marx, L. M. (2022). Incomplete information bargaining with applications to mergers, investment, and vertical integration. *American Economic Review*, 112(2):616–649.

Manski, C. F. (1989). Anatomy of the selection problem. *Journal of Human Resources*, pages 343–360.

Manski, C. F. (1990). Nonparametric bounds on treatment effects. *American Economic Review*, 80(2):319–323.

Manski, C. F. and Pepper, J. V. (2000). Monotone instrumental variables: With an application to the returns to schooling. *Econometrica*, 68(4):997–1010.

McAfee, R. P. and Reny, P. J. (1992). Correlated information and mechanism design. *Econometrica*, pages 395–421.

Molinari, F. (2020). Microeconometrics with partial identification. *Handbook of Econometrics*, 7:355–486.

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

Perry, M. (1986). An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica*, 54(2):313–321.

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *American Economic Review*, pages 1068–1095.

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

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

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