

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

Joachim Freyberger and Bradley J. Larsen[†]

August 24, 2021

Abstract

This study provides a structural analysis of detailed, alternating-offer bargaining data from eBay, deriving bounds on buyers and sellers private value distributions using a range of assumptions on behavior. These assumptions range from very 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 in theory and practice. For the median product, bargaining ends in impasses in 43% of negotiations even when the buyer values the good more than the seller.

*We thank Sharon Shiao, Evan Storms, and Caio Waisman for outstanding research assistance, and Brigham Frandsen, Matt Gentzkow, and Lars Lefgren for helpful comments, as well as seminar and conference participants at Berkeley Haas, Universitat de les Illes Balears, University of Melbourne, and BEET 2018. 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.

[†]Freyberger: University of Bonn; freyberger@uni-bonn.de. Larsen: Stanford University and NBER; bjarsen@stanford.edu.

1 Introduction

Bilateral bargaining is one of the oldest and most common forms of trade. A large theoretical literature, and a growing structural empirical literature, examines various aspects of bargaining situations, but the modeling choices of theorists and empiricists diverges 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 bargaining sequences from consumers negotiating online. We propose an approach to bound buyers’ and sellers’ values and the degree of inefficient *impasse* in the market. Unlike Nash bargaining, our approach is robust to the presence of incomplete information.

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

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

The challenge we face is, first, S and B are not observed in the data, and second, there is no theoretical characterization of equilibria in the game we study (bilateral negotiations in which both parties potentially have incomplete information and both parties can make offers), and hence no obvious way to back out estimates of F_S and F_B from observed bargaining actions.¹ Indeed,

¹Previous theoretical analysis of incomplete-information bargaining has highlighted a variety of challenges. For example, 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

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 price negotiations in consumer markets, where agents meet and negotiated infrequently and where it is arguably especially unrealistic to model agents as perfectly informed about the game structure or opponents' values (as Nash bargaining presumes).

To study this setting empirically, we propose a bounds approach based on an *incomplete* model. We first derive bounds on the marginal distributions of buyer and seller values, F_B and F_S . Our approach starts with weak rationality assumptions on agents' behavior. We then build on these assumptions by proposing stronger conditions on behavior and on the information environment. Our first additional assumption is that a seller's first offer (the Buy-It-Now price she chooses on eBay) is stochastically increasing in her value. Then we consider the stronger assumption that it is weakly increasing in her value. Next we propose an assumption that the buyer's value is stochastically increasing in the seller's first offer, capturing a notion of dependence between buyer and seller values. We then consider the stronger assumption that the two are independent. We propose similar assumptions that reverse the role between buyer and seller. Each of these assumptions yields a set of bounds on the marginal distribution of buyer and seller values. Building on these assumptions, we derive bounds on the first-best probability of trade that can inform us about the degree of inefficient impasse in the eBay data and the degree to which agents indeed face uncertainty about the gains from trade in this marketplace.

The bounds we derive under any given set of assumptions are sharp. The bounds are also nonparametrically identified. We propose estimation of these bounds using kernel estimators. We estimate the bounds separately product-by-product, limiting to products for which we have at least 100 bargaining sequences. The validity of the assumptions underlying our bounds can be analyzed by looking for cases where the bounds cross. We find evidence that our strongest assumptions (monotonicity of the seller's first offer or independence of the buyer's value and seller's first offer) are violated. Bounds based on our weaker assumptions, such as stochastic monotonicity or monotonicity of the buyer's first offer conditional on the seller's, do not cross.

We also exploit an additional piece of information unique to the eBay setting that allows us to test the validity and strength of our bounds. On eBay, sellers are permitted (but not required)

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 statements remain true several decades later.

to report an *auto-accept* and *auto-decline* price to the platform. Any prices outside this range are automatically accepted or declined by eBay, without the seller needing to manually respond to the offer. These secret boundaries are themselves bounds on the marginal distribution of seller values. We demonstrate that, in the sample where sellers report auto accept/decline prices, our preferred bounds correspond surprisingly well to these secret bounds (which are not explicitly used anywhere in computing our bounds).

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

A lower bound on the first-best probability of trade, compared to the probability of the trade in the data, contains information about the degree of inefficient impasse. For example, for a popular computer product in our sample, agents agree in the real-world negotiations 22.3% of the time. Under our preferred assumptions, we find that the counterfactual first-best probability of trade is 0.374, suggesting that 40% of the time ($1 - 0.223/0.374$), agents fail to reach an agreement even when the buyer truly values the good more than the seller. For the median product, this percentage of inefficient impasse is 43%, and ranges from 8.3% to 68.9% across all products.

An upper bound on the first-best probability of trade provides information on the uncertainty agents face. If agents have no uncertainty about whether gains from trade exist, they would agree 100% of the time in a first-best world. We find that, for the median product, an upper bound on the first-best probability of trade is below 1 for most products, suggesting that consumers in this market do indeed face uncertainty about whether gains from trade exist.

Our study contributes to both the theoretical and empirical bargaining literatures. Theory work on incomplete-information bargaining studies this topic either by explicitly modeling the extensive form of the game or by applying mechanism design tools. When agents' values are independent, even our strongest behavioral assumptions (monotonicity of agents' first offers in their values) are satisfied in the equilibria of extensive-form games focused on in the literature (e.g. Perry 1986, Grossman and Perry 1986, and Cramton 1992).² We demonstrate, however, that in the presence of unobserved game-level heterogeneity (i.e., features of the negotiation that shift or scale the values

²See Ausubel et al. (2002) for a survey of the incomplete-information bargaining literature.

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 game-level heterogeneity.

In the mechanism design literature, our study is most closely related to Myerson and Satterthwaite (1983), who demonstrated that when agents face uncertainty about whether gains from trade exist, no incentive-compatible, individually rational mechanism will realize the first-best gains from trade without running a deficit. In his extensive-form game, Cramton (1992) showed that the first-best *probability* of trade is attainable if agents burn surplus to signal their values. Our study examines two dimensions of these points. First, we quantify how close eBay participants get to the first-best probability of trade. Second, we study how much uncertainty agents face on eBay about whether gains from trade exist.

Our study relates to a small but recently growing literature estimating structural models of incomplete-information bargaining games. The most closely related studies are those of 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 negotiation in 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 these previous studies make assumptions on the optimality of negotiators' behavior or their knowledge of the 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, as in Bajari et al. (2007). Larsen (2021) relied on auction outcomes in addition to bargaining data, and used an identification and estimation approach which is a special case of one of the many bounds we propose herein (our *independence* bounds for one party — the seller only).³ In contrast, the

³Two structural empirical studies that also examined the used-car setting are Larsen and Zhang (2018) and Larsen

methodology we develop does not rely on any auction data, only sequential-offer bargaining data, and extends well beyond the independence case.

Several structural empirical studies have focused on models where bargaining consists of a take-it-or-leave-it offer (e.g. Silveira 2017, studying bargaining in judicial settings) or sequential bargaining with all offers made 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 (2015) studied incomplete-information bargaining that takes the form of a k double auction, where each party simultaneously make a single offer. Our study focuses instead on a setting where multiple offers from both parties can and frequently do occur in the data, and hence the frameworks of these previous papers are not appropriate for our setting.

Our work also relates to a literature that exploits eBay as a laboratory for studying fundamental questions of price discovery and efficiency in large, decentralized markets. The structural literature examining efficiency of eBay trading mechanisms has largely focused on auctions (e.g. Hendricks and Sorensen 2018; Bodoh-Creed et al. 2021). Backus et al. (2020) and Keniston et al. (2021) offered descriptive analyses of eBay bargaining data, using a superset of the data we analyze here, and documented a number of patterns consistent with the existence of incomplete information and cognitive limitations in this marketplace, underscoring the benefit of our flexible approach to bounding agents’ values without assuming a complete model of fully rational equilibrium behavior.

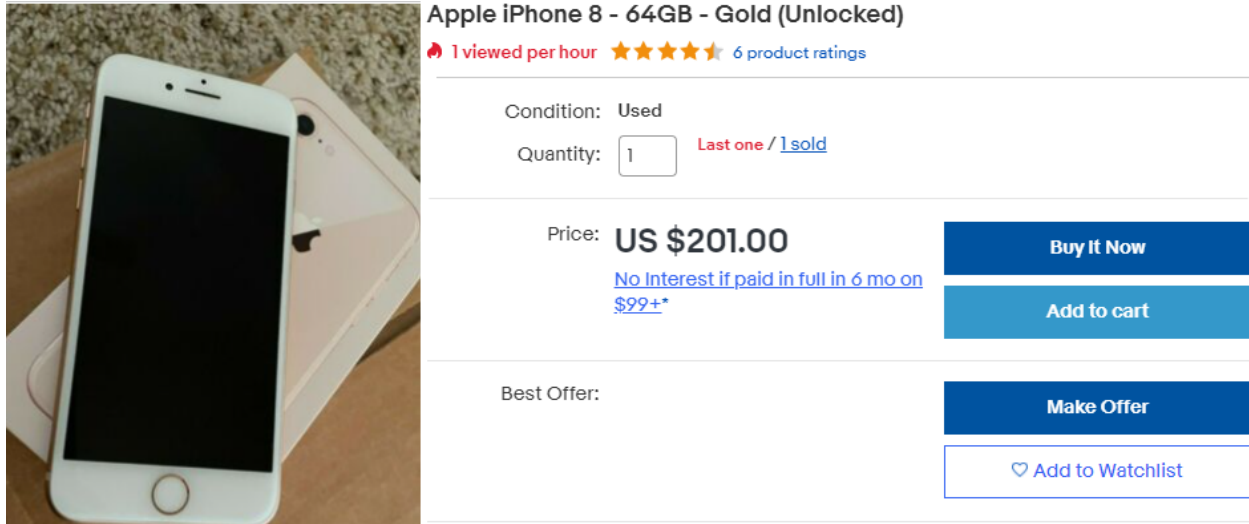
2 eBay’s Best Offer Platform

While eBay is most well known for its auction and Buy-It-Now (posted-price) transactions, a fast-growing sales mechanism on the platform is the *Best Offer* option, which was introduced around 2006. This option allows a seller and buyer to haggle over an item’s price. When creating a posted-price listing, the seller is given the option to “allow offers”, and this option has, for many years, been enabled by default on every listing. If this option is selected, a prospective buyer viewing the listing will see the Buy-It-Now price (which we will refer to here as the list price) as well as a *Make Offer* button, illustrated for an iPhone 8 listing in Figure 1.

Clicking this Make Offer button allows the buyer to propose an offer to the seller. The seller can then respond by declining, accepting, or making a counteroffer. If the seller accepts, the parties

et al. (2021). 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, while plausible in the used-car market setting — which only includes experienced professional negotiators — are unlikely hold in the consumer negotiations we study herein.

Figure 1: Illustration of Best Offer Listing



trade at that accepted price. If the seller counters, it is then the buyer's turn to accept, decline, or counter. If the seller declines, the buyer may still choose to make another counteroffer. The buyer and seller are each limited to three offers, and the buyer is permitted to purchase at the list price at any time.⁴ If any agent fails to respond to an offer within 48 hours, the offer expires (essentially equivalent to the offer having been declined).

When creating the listing, the seller is permitted to specify an *auto-accept* and *auto-decline* price. Any buyer offer above the auto-accept price is automatically accepted by the platform, without the seller needing to manually respond. And any buyer offer below the auto-decline price is automatically declined by the platform. In our analysis below, we take advantage of these secretly reported thresholds to examine the validity of our bounds.

Buyers in this marketplace are typically retail consumers. Sellers in this marketplace may be a business or an individual. Thus, our data comes from a mix of business-to-consumer (B2C) and consumer-to-consumer (C2C) price negotiations. Consumers thus play an important role in this market, in contrast to the professional negotiators studied in 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 to analyzing the game that does not require a complete model of equilibrium behavior.

Our sample is a subset of the data created for the descriptive analysis in Backus et al. (2020). The full dataset contains the 25 million bargaining sequences that occurred on the U.S. eBay site

⁴This three-offer limit was in place at the time period from our data comes. In more recent years, eBay moved to a five-offer limit.

from June 2012 through May 2013. For this paper, we focus on a subset of bargaining sequences corresponding to products that have well-labeled product identifiers. These include products such as “Apple iPhone 8 64 GB” or “Xbox 360”. We observe each item’s condition type (used vs. new), and we consider a *product* to be a combination of this condition type and its product identifier. For each product, we construct a *reference price* using all non-Best-Offer posted-price sales of that same product during our sample period.⁵ As we perform estimation separately for each product, the reference price plays no role other than as a normalization, putting each product on a similar scale by dividing prices/offers by its reference price.

We limit our sample to listings to which a buyer makes an offer. This means we drop cases where a buyer arrives at a Best Offer listing but leaves the page without making an offer, or selects to purchase at the Buy-It-Now price. 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 also impose several other sample restrictions. In our analysis, we wish to analyze bargaining within a given bilateral pair. Our data includes buyer and seller identifiers, and thus we can see if a given buyer negotiated with one seller of a product and then, after failing to reach an agreement, negotiated with another seller of the same product. In these cases, we have no way to knowing whether the buyer wanted multiple copies of the same product or only one. We also observe sellers negotiating with multiple buyers on the same listing. We focus on only the first seller of a given product with which a given buyer negotiates, and, among these negotiating pairs, the first buyer which which a given seller negotiates. This gives us a unique set of buyers and sellers in each bargaining sequence. We also limit to products for which we observe at least 100 negotiation sequences. After imposing these sample restrictions, we are left with 70,807 bargaining sequences corresponding to 363 products.

The data does not specify the exact title of each product (only an anonymous product identifier) but does specify the category of the product on the eBay platform. Table 1 displays descriptive statistics for the top-selling product in each of the eight categories appearing in our final sample. These products are varied in their reference price (the second column). As described above, this price is computed using non-Best-Offer fixed price sales on eBay. These prices range from \$15 for the music product to \$432 for the home and garden product. The final price as a fraction of the

⁵We restrict the sample to products that have at least ten non-Best-Offer posted-price sales for the construction of the reference price.

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

Category	Reference Price (\$)	n	$P(\text{sale})$	Final Price Over List (if trade)	Buyer Price Over List (if no trade)	Seller Price Over List (if no trade)
Electronics	51.44	577	0.31	0.77	0.6	0.96
Cameras	60.23	159	0.38	0.76	0.54	0.95
Sports	180.98	190	0.31	0.82	0.68	0.97
Video Games	80.78	487	0.29	0.83	0.65	0.96
Musical	15.4	123	0.72	0.74	0.54	0.85
Home/Garden	432.54	150	0.2	0.92	0.73	0.98
Cell Phones	224.82	2,501	0.13	0.88	0.73	0.98
Computers	131.68	497	0.22	0.87	0.7	0.98
All Products	213.75	70,807	0.21	0.84	0.68	0.97

Notes: First eight rows show descriptive statistics for the top-selling product in each category. Final row shows same statistics for the full sample of 363 products (70,807 observations). The object n represents the number of observations.

list price, when trade occurs, ranges from 0.74 to 0.92. When trade fails, the highest price offered by the buyer as a fraction of the list price (the second-to-last column) ranges from 0.54 to 0.73, whereas the lowest price offered by the seller in these disagreement cases ranges from 0.85 to 0.98. One object of interest in this study is the fourth column, the probability that negotiation ends in agreement. The probability varies widely — from 0.13 for the cell phone product (which was the subject of 2,501 negotiating sequences) to 0.72 for the music product (which was the subject of 123 negotiating sequences). Our empirical approach allows us to compute counterfactual bounds on what the probability of trade would be in a first-best world.

3 Bounds on Values in a Bargaining Game

In this section we present bounds on buyers’ and sellers’ values in an alternating-offers bargaining game, the game used on eBay’s Best Offer platform.⁶ We begin with bounds under minimal assumptions, using revealed preference arguments only. We then introduce assumptions on the strategic behavior of the agents and the dependence of the buyer and seller values to tighten these bounds.

3.1. Bargaining Game Setup and Notation. A seller with value $S \sim F_S$ and buyer with value $B \sim F_B$ engage in an alternating-offer bargaining game.⁷ If the buyer and the seller agree to trade at a price P , the buyer’s payoff is $B - P$ and the seller’s payoff is P (less any bargaining

⁶The bounds we derive can be modified to allow for protocol other than alternating offers.

⁷Throughout, we use uppercase letters to denote random variables and lowercase to denote realizations.

costs). If they break up the bargaining (i.e., some agent chooses to quit), the seller gets S and the buyer gets 0 (again, less any costs of bargaining). We remain agnostic about the form of any bargaining costs, such as discounting or per-period additive costs, which are the two most common forms assumed in the extensive-form bargaining literature. Maintaining the assumptions that we do in this paper (which we make explicit below), the bounds we derive are robust to ignoring such costs. Throughout the paper, we maintain the assumption that, in a given instance of the game, S and B are realized before any actions take place and are held fixed throughout the game.⁸

To match the structure of eBay’s Best Offer platform, we treat the first bargaining offer as coming from the seller. This first offer is the list price posted by the seller. The buyer makes the second offer. The seller can then choose to accept, counter, or quit. We refer to each turn as a *period*. We denote the beginning period of the game as $t = 1$. The buyer then responds at $t = 2$, and so on, with t odd being the seller’s turn and t even being the buyer’s.

Denote the offers of the seller in period t by P_t^S (if t is a period in which seller gets to make an offer) and the offer by the buyer as P_t^B (if t is a period in which the buyer gets to make an offer).⁹ Denote the decision of the buyer in period 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 when t is the seller’s turn. Either player choosing to accept or quit ends the game. For a given instance of the game, the data available to the econometrician consists of the sequence of offers made and any decision made by the buyer or seller to accept or quit.

We define the following random variables at the level of the *bargaining sequence* rather than at the level of a *period* (t) *within* a bargaining sequence: Let $D^S = A$ (without a t subscript) if the seller ever accepts or counters in a given bargaining sequence, and $D^S = Q$ if the seller ever quits. Similarly, let $D^B = A$ if the buyer ever accepts or counters and $D^B = Q$ if the buyer ever quits. Let $X_{AC}^S = \min\{\{P_t^B : D_{t+1}^S = A\}, \min_t\{P_t^S : D_t^S = C\}\}$. Thus, X_{AC}^S is the smallest offer the seller ever makes or accepts in a given bargaining sequence (where the notation “AC” stands for “accepting or countering”). Note that X_{AC}^S is always defined because the seller always makes the first offer (so $D_1^S = C$). Also, let $X_Q^S = \max\{\{P_t^B : D_{t+1}^S = Q\}, 0\}$ be the offer a seller quits at if indeed the seller quits, and 0 otherwise. Notice that it is only possible for there to be at most one offer in a

⁸We 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 (2020) show in laboratory experiments is certainly a possibility. We also do not consider the possibility of a buyer’s outside option changing during the bargaining game.

⁹The superscripts S and B are not strictly necessary given that odd t correspond to the seller and even t to the buyer, but we maintain the superscripts for additional clarity.

bargaining sequence at which a seller quits. The definition ensures that X_Q^S is well defined even if $D^S \neq Q$. Let $X_{AC}^B = \max\{\{P_t^S : D_{t+1}^B = A\}, \max_t\{P_t^B : D_t^B = C\}, 0\}$ be the largest price a buyer accepted or offered in a given bargaining sequence. Finally, let $X_Q^B = \min\{\{P_t^S : D_{t+1}^B = Q\}, \infty\}$ be the offer a buyer quits at, if indeed the buyer quits, and ∞ if $D^B \neq Q$. The infinite support points are conservative and can easily be replaced with less conservative assumptions. In the same spirit, we could use $-\infty$ rather than 0 as the lower bound of the support of seller values in defining X_Q^S .

For all of our results in the body of the paper, we focus on the case where the buyer always counters in the second period ($D_2^B = C$). In the proofs, found in Appendix A, we derive bounds for the case that is slightly more general (but more cumbersome in terms of notation) in which the buyer may accept or quit in the second period rather than counter, immediately ending the game with no buyer offers occurring.

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

$$P(S \leq x) = \int P(S \leq x \mid P_1^S = y) dF_{P_1^S}(y) \quad (1)$$

$$P(B \leq x) = \int P(B \leq x \mid P_1^S = y, P_2^B = z) dF_{P_1^S, P_2^B}(y, z) \quad (2)$$

where $F_{P_1^S}$ is the CDF of sellers' first offers, P_1^S , and $F_{P_1^S, P_2^B}$ is the joint distribution of sellers' and buyers' first offer, P_1^S and P_2^B . When written as a function, $P(\cdot)$ represents the probability of a given event (e.g., $P(S \leq x)$).

As a final piece of notation, for given random variables X and Y , let $\overline{\text{supp}}(Y \mid X \geq a)$ and $\underline{\text{supp}}(Y \mid X \geq a)$ be the maximum and minimum of the support of Y given $X \geq a$, respectively. Define $X_{AC}^{S*}(y) \equiv \underline{\text{supp}}(X_{AC}^S : P_1^S \geq y)$ and $X_Q^{S*}(y) \equiv \overline{\text{supp}}(X_Q^S : P_1^S \leq y)$. Thus, $X_{AC}^{S*}(y)$ is the smallest accept/counter price of sellers conditional on events where the first offer of sellers is at least y . $X_Q^{S*}(y)$ has a similar interpretation. For all (y, z) on the support of (P_1^S, P_2^B) , define $X_{AC}^{B*}(y, z) \equiv \overline{\text{supp}}(X_{AC}^B : P_2^B \leq z, P_1^S = y)$ and $X_Q^{B*}(y, z) \equiv \underline{\text{supp}}(X_Q^B : P_2^B \geq z, P_1^S = y)$. Conditional on P_1^S , $X_{AC}^{B*}(y, z)$ is the largest accept/counter price of buyers conditional on events where the second offer of buyers is at most z . $X_Q^{B*}(y, z)$ has a similar interpretation.

3.2. Unconditional Bounds on Value Distributions. We now describe the range of assumptions we make about equilibrium behavior that yield bounds on buyer and seller marginal value distributions. Our first and weakest assumption is the following:

Assumption 1 (Revealed Preferences). (i) The seller never accepts (or counters) at a price $P < S$ or quits at a price $P > S$, and (ii) the buyer never accepts (or counters) at a price $P > B$ or quits at a price $P < B$.

These revealed preference assumptions are similar to those employed in Haile and Tamer (2003) for English auctions. There, the authors assume that (i) a bidder never bids above her value, analogous in our setting to a buyer never accepting or countering at a price above B ; and (ii) a bidder never lets another agent win the auction at a price she would have been willing to beat, analogous in our setting to the assumption that a buyer never quits when doing so yields a payoff (i.e., 0) that is lower than the payoff from accepting ($B - P$). Our seller assumptions have a similar interpretation.¹⁰ Importantly, Assumption 1 imposes only weak rationality conditions, and does not impose that agents behave according to any particular equilibrium concept, although the conditions are weak enough to be satisfied by standard equilibrium concepts, such as Bayes Nash or Perfect Bayes Equilibrium.

We maintain Assumption 1 everywhere in the paper, without restating it each time it is used. Its important implications, used throughout the paper, are

$$X_Q^S \leq S \leq X_{AC}^S \tag{3}$$

$$X_{AC}^B \leq B \leq X_Q^B. \tag{4}$$

These inequalities immediately imply bounds on F_S and F_B , which we refer to as our *unconditional* bounds:

$$P(X_{AC}^S \leq x) \leq F_S(x) \leq P(X_Q^S \leq x) \tag{5}$$

$$P(X_Q^B \leq x) \leq F_B(x) \leq P(X_{AC}^B \leq x) \tag{6}$$

We formally state these bounds as the following theorem:

Theorem 1. Under Assumption 1.i, (5) bounds F_S , and under Assumption 1.ii, (6) bounds F_B .

The proof of this result (Appendix A) follows immediately from $X_Q^S \leq S \leq X_{AC}^S$ and $X_{AC}^B \leq$

¹⁰An important distinction, however, is that, in the auction setting of Haile and Tamer (2003), upper and lower bounds exist *for each* observation in the data (a lower bound is given by a buyer's bid and an upper bound is given by exploiting the minimum bid increment). In contrast, in the eBay two-sided bargaining game, a given realization of the game may end with many potential bargaining actions never being observed (for example, if the negotiation ends in agreement, no upper bound on the buyer's value is observed). We handle this extra complication of our setting by relying on *probabilities* of certain events occurring, rather than on empirical CDFs of prices/bids alone.

$B \leq X_Q^B$. These bounds are sharp. Specifically, given that we place no restrictions on behavior other than Assumption 1, nothing in the data or in the assumptions rules out the possibility that the play of the game is such that $X_{AC}^S = S$. For the upper bound, nothing in the data or in the assumptions rules out the possibility that, in any sequence in which the seller quits ($D^S = Q$), $X_Q^S = S$, and in any sequence in which the seller does not quit ($D^S \neq Q$), the seller has a value of $S = 0$.¹¹ This same line of reasoning applies to the buyer bounds. These bounds can be relatively tight in some cases and quite loose in others. Appendix C offers Monte Carlo simulations illustrating this point.

The bounds are also nonparametric. Each of the above bounds are weakly increasing and lie in $[0, 1]$, and thus can correspond themselves to a cumulative distribution function. The bounds will be valid even if the game has multiple equilibria, and, in particular, even if the data is not all generated by the same equilibrium or by any standard notion of equilibrium play. Furthermore, if the true data generating process does not in fact entail sellers all drawing for the same distribution F_S — that is, if sellers (or, analogously, buyers) are asymmetric — then the bounds will still remain valid for the mixture distribution of values in the data.

These bounds do not place any 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. One form of such 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 perspective of the analyst, are correlated *across instances of the game* through the presence of W . A related possibility is multiplicative separability, where $S = W\tilde{S}$ and $B = W\tilde{B}$.¹² These two structures are not imposed anywhere in our paper but we highlight them below as special cases that are allowed for by our moderate assumptions and ruled out by our strongest assumptions (and by some existing theoretical models).

3.3. Bounds Based on Agent’s Value and Agent’s Offer. Our next assumptions describe how an agent’s first offer relates to her own value:

Assumption 2 (Monotonicity). (i) P_1^S is weakly increasing in S , and (ii) P_2^B is weakly increasing

¹¹Throughout 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.

¹²Multiplicative or additive separability are two structures commonly assumed in empirical auction work (e.g., Krasnokutskaya 2011; Freyberger and Larsen 2020).

in B conditional on P_1^S .

Assumption 2.i describes own-offer weak monotonicity for the seller. Under this assumption, for $y < y'$, a seller with $P_1^S = y$ must have a weakly lower value than a second seller who has chooses $P_1^S = y'$, and therefore the lowest price at which the second seller counters or accepts is an upper bound on the value of the first seller. This is precisely what is represented by $X_{AC}^{S*}(y)$, defined in Section 3.1. Similarly, Assumption 2.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 Assumption 2, monotonicity for the buyer, is weaker, as it is *conditional* on the seller's first offer. These arguments yield the following bounds:

$$\int \mathbf{1}(X_{AC}^{S*}(y) \leq x) dF_{P_1^S}(y) \leq F_S(x) \leq \int \mathbf{1}(X_Q^{S*}(y) \leq x) dF_{P_1^S}(y) \quad (7)$$

$$\int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_1^S, P_2^B}(y, z) \leq F_B(x) \leq \int \mathbf{1}(X_{AC}^{B*}(y, z) \leq x) dF_{P_1^S, P_2^B}(y, z) \quad (8)$$

where $\mathbf{1}(\cdot)$ represents the indicator function. These bounds are derived as follows: Under Assumption 2, we have $X_Q^{S*}(y) \leq S \leq X_{AC}^{S*}(y)$, and, conditional on P_1^S , the objects $X_{AC}^{S*}(y)$ and $X_Q^{S*}(y)$ are non-random. We plug these objects into the iterated expectations representation from (1) to obtain (7). The buyer bounds follow similarly.

We refer to these bounds as *monotonicity* bounds. These bounds improve upon the unconditional bounds by comparing the accept/counter or quit actions of agents across instances of the game. Appendix C includes Monte Carlo simulations that illustrate cases where the bounds can be narrow and cases where they will not improve upon the unconditional bounds.

Assumption 2 may be too strong for our online bargaining setting; a weaker assumption that also exploits comparisons across instances of the game is own-offer stochastic monotonicity:

Assumption 3 (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 $z \forall y, x$.

Assumption 3 means that an agent's value is more likely high when her first offer is high, and this assumption is thus implied by Assumption 2.¹³ Combining Assumption 3 with the iterated expectation representations from (1) and (2), we obtain the following, which we refer to as the

¹³In another recent application of a similar assumption to obtain partial identification, Frandsen and Lefgren (2021) exploit stochastic monotonicity to bound treatment effects of attending a charter school.

stochastic monotonicity bounds:

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

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

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

We state the monotonicity and stochastic monotonicity bounds as the following theorem. The proof of this theorem derives the bounds step by step.

Theorem 2. *The following bounds hold: (i) under Assumption 2.i, (7) bounds F_S ; (ii) under Assumption 2.ii, (8) bounds F_B ; (iii) under Assumption 3.i, (9) bounds F_S ; and (iv) under Assumption 3.ii, (10) bounds F_B .*

These bounds are again sharp: under their corresponding assumptions, it is impossible to rule out that the bounds hold with equality. To provide intuition and motivation for these assumptions, in Appendix D we walk through equilibria considered in Perry (1986) and Cramton (1992). Our bounds allow for a much wider range of possible outcomes than these equilibria; indeed, these equilibria are among the infinitely many outcomes that our bounds allow for. We focus on these examples only because, to our knowledge, they are some of the few extensive-form equilibria studied anywhere in the literature from a bargaining game that comes close in generality to the game we study — a bargaining game with two-sided incomplete information and a continuous value distribution where both parties are allowed to make offers in equilibrium.¹⁴ Monotonicity (and hence, stochastic monotonicity as well) is satisfied in the environments of Perry (1986) and Cramton (1992), which assume independent private values with no unobserved game-level heterogeneity. But monotonicity can be violated in a more general environments, such as additive (e.g., $S = W + \tilde{S}$ and $B = W + \tilde{B}$) or multiplicative (e.g., $S = W\tilde{S}$ and $B = W\tilde{B}$) unobserved heterogeneity, as described in Section 3.2. Importantly, stochastic monotonicity — our weaker assumption — will still be satisfied in the presence of unobserved heterogeneity.

¹⁴See 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. In Appendix D, in addition to the models of Perry (1986) and Cramton (1992), we also highlight the equilibrium of Grossman and Perry (1986), which is less general than the other two in that only one party has incomplete information. Recent work by Keniston et al. (2021) derives a Perfect Bayes Equilibrium of a two-sided incomplete-information, alternating-offer game in which offers along the equilibrium path do not depend on an agent's value, but rather split the difference between the two most recent offers.

This raises an important point for empirical work: theoretical equilibrium models of bargaining, even if they describe behavior well, may be unhelpful for empirical work 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 and our more moderate assumptions, which can help bridge the gap between restrictive, extensive-form (and complete) models and analysis of bargaining in actual negotiation data.

3.4. Bounds Based on Agent’s Value and Opponent’s Offer. We now consider how one player’s *value* relates to another player’s *offer*:

Assumption 4 (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 in the game, a natural assumption is that the seller’s first offer depends on S . Assumption 4.ii takes this one step further and assumes that the seller’s first offer does not depend on the *buyer’s* value; Assumption 4.i describes a similar condition for the seller’s value, but this condition is weaker, as it is conditional on the seller’s first offer. Through this relationship between an agent’s value and an opponents’ offer, Assumption 4 captures a notion of independence between buyer and seller values.

We obtain the following bounds under Assumption 4, which we refer to as *independence* bounds:

$$\int \max_z m_{AC}^S(x, y, z) dF_{P_1^S}(y) \leq F_S(x) \leq \int \min_z m_Q^S(x, y, z) dF_{P_1^S}(y) \quad (11)$$

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

where $m_{AC}^S(x, y, z) = P(X_{AC}^S \leq x \mid P_1^S = y, P_2^B = z)$ and $m_Q^S(x, y, z) = P(X_Q^S \leq x \mid P_1^S = y, P_2^B = z)$. These bounds are obtained by combining Assumption 4 with (3) and (4), and applying the iterated expectation representation of (1) and (2). These bounds can be narrow or wide in practice; Monte Carlo simulations in Appendix C illustrate both cases and discusses data features affecting the bounds’ width.

Like Assumption 2, Assumption 4 may be too strong for the eBay bargaining platform. A weaker alternative is the following:

Assumption 5 (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 .*

Assumption 5 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. Assumption 5 is implied by Assumption 4. Under Assumption 5 we obtain

$$\int \max_{z' \geq z} m_{AC}^S(x, y, z') dF_{P_1^S, P_2^B}(y, z) \leq F_S(x) \leq \int \min_{z' \leq z} m_Q^S(x, y, z') dF_{P_1^S, P_2^B}(y, z) \quad (13)$$

$$\int \max_{y' \geq y} P(X_Q^B \leq x \mid P_1^S = y') dF_{P_1^S}(y) \leq F_B(x) \leq \int \min_{y' \leq y} P(X_{AC}^B \leq x \mid P_1^S = y') dF_{P_1^S}(y) \quad (14)$$

We refer to these bounds as the *positive correlation* bounds to distinguish them from Assumption 3, stochastic monotonicity.

We formally state the independence and positive correlation bounds as follows:

Theorem 3. *The following bounds hold: (i) under Assumption 4.i, (11) bounds F_S ; (ii) under Assumption 4.ii, (12) bounds F_B ; (iii) under Assumption 5.i, (13) bounds F_S ; and (iv) under Assumption 5.ii, (14) bounds F_B .*

The bounds described in Theorem 3 are sharp. Returning to our examples of theoretical models, Perry (1986) and Cramton (1992), we note that independence (and hence also positive correlation) is satisfied in those models. However, in a modified version of their settings, with unobserved game-level heterogeneity, Assumption 4 can be violated, even while the weaker condition, Assumption 5, still holds. We demonstrate these results in Appendix D. As discussed above, we highlight these equilibria only as examples; our bounds do not rely on these models in any form.

3.5. Combining Assumptions on Marginal Distributions. Assumptions 2–5 can be combined with one another to obtain tighter bounds. For example, we can combine Assumptions 2 and 4 — monotonicity and independence — our two strongest assumptions. Or we can combine Assumptions 3 and 5 — stochastic monotonicity and positive correlation — two weaker assumptions. Appendix B derives bounds based on such combinations. As with all of our bounds, these bounds are sharp.

4 Estimation

In this section we describe estimators for the bounds from Section 3. In our data, an observation $i = \{1, \dots, n\}$ consists of a buyer-seller pair and their corresponding bargaining sequence.

4.1. Preliminary Ingredients for Estimation. For each i , the variables $X_{AC,i}^S$, $X_{Q,i}^S$, $X_{AC,i}^B$, and $X_{Q,i}^B$ are directly observed. Using these variables, we estimate the conditional probability

$P(X_Q^B \leq x \mid P_1^S = y') = E[\mathbf{1}(X_Q^B \leq x) \mid P_1^S = y']$ for each value of x using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel and bandwidth $n^{-1/5}$. This bandwidth choice yields, under certain assumptions, the optimal rate of convergence (Stone 1982). Let $\widehat{P}(X_Q^B \leq x \mid P_1^S = y')$ denote the estimator. We proceed analogously for $P(X_{AC}^B \geq x \mid P_1^S = y')$, $P(X_{AC}^S \leq x \mid P_1^S = y')$, and $P(X_Q^S \geq x \mid P_1^S = y')$. We restrict all estimators to be in the interval $[0, 1]$ and rearrange them such that the estimated functions are monotone in x .

Similarly, we estimate the function $m_Q^B(x, y, z)$ using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel and bandwidth $n^{-1/6}$. Here, due to the higher dimension, the optimal bandwidth converges at a slower rate. Again, we restrict the estimators to be in the interval $[0, 1]$ and rearrange them such that they are monotone in x . We denote the estimator by $\widehat{m}_Q^B(x, y, z)$. We estimate $\widehat{m}_{AC}^B(x, y, z)$, $\widehat{m}_{AC}^S(x, y, z)$, and $\widehat{m}_Q^S(x, y, z)$ analogously.

Finally, we need estimates of $X_{AC}^{S*}(y)$, $X_Q^{S*}(y)$, $X_{AC}^{B*}(y, z)$, and $X_Q^{B*}(y, z)$. The variables $X_{AC}^{S*}(y)$ and $X_Q^{S*}(y)$ can be estimated with sample analogs. That is,

$$\widehat{X}_{AC}^{S*}(y) = \max_{i: P_{1,i}^S \geq y} X_{AC,i}^S \quad \text{and} \quad \widehat{X}_Q^{S*}(y) = \max_{i: P_{1,i}^S \leq y} X_{Q,i}^S$$

$X_{AC}^{B*}(y, z)$ and $X_Q^{B*}(y, z)$ are more complicated to estimate because we condition on a specific value of the continuous variable P_1^S . To do so, let $N(y) = \{z \in \mathbb{R} : |z - y| \leq h_n(y)\}$ be a neighborhood of y where the neighborhood size $h_n(y)$ is sample-size dependent and decreases to 0 as $n \rightarrow \infty$. Now define

$$\widehat{X}_{AC}^{B*}(y, z) = \max_{i: P_{2,i}^B \leq z, P_{1,i}^S \in N(y)} X_{AC,i}^B \quad \text{and} \quad \widehat{X}_Q^{B*}(y, z) = \max_{i: P_{2,i}^B \geq z, P_{1,i}^S \in N(y)} X_{Q,i}^B$$

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 $\sum_{i=1}^n \mathbf{1}(|P_{1,i}^S - y| \leq h_n(y)) = K_n$. We choose $K_n = n^{1/4}$. If the density of $P_1^S(y)$ 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 $n \rightarrow \infty$.

While we use kernel estimators for each of the above functions, we find similar results using series estimators. Note also that we do not explicitly discuss conditioning on covariates in these estimators, as we perform estimation product-by-product and, under the assumptions we maintain, our bounds are robust to omitting such controls. However, our bounds arguments could be applied conditional on covariates by including them in the estimated conditional probability functions. In this case, a parametric approximation, such as a probit model, may be preferred to the nonpara-

metric estimators we propose here. The matching approach described above could then be used to estimate the support bounds for the monotonicity assumptions.

4.2. Estimation of Bounds. With the ingredients from above, we can estimate the bounds from Section 3. For brevity, we describe here the estimation of each lower bound; the estimators for the upper bounds are analogous. To estimate the unconditional lower bounds, we simply plug in the empirical analogs of (5) and (6):

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_{AC,1}^S \leq x) \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \mathbf{1}(X_{Q,i}^B \leq x).$$

To estimate the monotonicity bounds, note that

$$\int \mathbf{1}(X_{AC}^{S*}(y) \leq x) dF_{P_1^S}(y) = E_{P_1^S}[\mathbf{1}(X_{AC}^{S*}(P_1^S) \leq x)].$$

We therefore estimate the lower bounds from (7) and (8) by

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}(\hat{X}_{AC}^{S*}(P_{1,i}^S) \leq x) \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\hat{X}_Q^{B*}(P_{1,i}^S, P_{2,i}^B) \leq x).$$

We estimate the stochastic monotonicity lower bounds in (9) and (10) by

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \hat{P}(X_{AC}^S \leq x \mid P_1^S = y')$$

and

$$\frac{1}{n} \sum_{i=1}^n \max_{z' \in \{z: z \geq P_{2,i}^B, Q_{0.05}(P_{2,i}^B) \leq z \leq Q_{0.95}(P_{2,i}^B)\} \cup \{P_{2,i}^B\}} \hat{m}_Q^B(x, P_{1,i}^S, z')$$

where $Q_\alpha(P_{1,i}^S)$ and $Q_\alpha(P_{2,i}^B)$ denote the α quantiles of $P_{1,i}^S$ and $P_{2,i}^B$, respectively. Notice that the sample analog estimator of the seller's stochastic monotonicity lower bounds is

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \geq P_{1,i}^S} \hat{P}(X_{AC}^S \leq x \mid P_1^S = y')$$

Here, we use the additional constraint that $Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)$ because the function $P(X_{AC}^S \leq x \mid P_1^S = y')$ can be poorly estimated at the boundary of the support. To ensure the set is never empty, we always include $P_{1,i}^S$. Note that applying this tail truncation yields conservative estimates of the bounds; the same is true for all estimators in the paper that use this truncation.

We estimate the independence lower bounds in (11) and (12) by

$$\frac{1}{n} \sum_{i=1}^n \max_{Q_{0.05}(P_{2,i}^B) \leq z \leq Q_{0.95}(P_{2,i}^B)} \hat{m}_{AC}^S(x, P_{1,i}^S, z) \quad \text{and} \quad \max_{Q_{0.05}(P_{1,i}^S) \leq y' \leq Q_{0.95}(P_{1,i}^S)} \hat{P}(X_Q^B \leq x \mid P_1^S = y').$$

Finally, we estimate the positive correlation lower bounds in (13) and (14) by

$$\frac{1}{n} \sum_{i=1}^n \max_{z' \in \{z: z \geq P_{2,i}^B, Q_{0.05}(P_{2,i}^B) \leq z \leq Q_{0.95}(P_{2,i}^B)\} \cup \{P_{2,i}^B\}} \hat{m}_{AC}^S(x, P_{1,i}^S, z')$$

and

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \hat{P}(X_Q^B \leq x \mid P_1^S = y').$$

The estimators of our bounds obtained by combining assumptions together are similar, and are described in Appendix B.

5 Bounding Values in eBay Bargaining

5.1. Bounds on Buyer and Seller Values for Cell Phones. We now apply the estimators derived above to obtain bounds on the distributions of buyer and seller values. As described in Section 2, we do this estimation separately for each product in our data, limiting to products for which we observe at least 100 negotiation sequences. We normalize prices by the product’s reference price for ease of interpretation. For illustrative purposes, we begin by focusing on one particular product — the most popular cell phone from Table 1. Figure 2 displays the bounds on the buyer value CDF for this cell phone product under different assumptions. Every panel also shows the unconditional bounds for comparison. Upper bounds are shown with dashed lines and lower bounds with solid lines.

These results demonstrate that the unconditional bounds — which rely on very weak assumptions — can be very wide, the lower bound in particular. This is because it is constructed using prices at which a buyer quits (walking away from bargaining), which are unobserved if a sequence either ends in agreement or if the seller quits (rather than the buyer). In such cases, we cannot rule out the buyer having a very large value, and hence the lower bound is very low. The upper bound, on the other hand, relies on prices at which the buyer accepts or counters, and at least one of these prices is *always* available in the data we analyze, as we focus only on listings that include

a buyer proposing an offer (P_2^B).¹⁵

Depending on the product, some additional 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 bounds. This does not imply the assumption is violated, rather that is just too weak to tighten the bounds for this product. Stochastic monotonicity implies that, conditional on the seller’s first offer, a buyer with a higher first offer is more likely to have a higher value. This assumption will lead to a tightening of the lower bound if, for example, in some instances of the game in which buyers have relatively high first offers the buyer ends up accepting a relatively *low* seller offer. This can happen due to the randomness in which seller the buyer is matched to, or to other features of the game that generate later offers, which we make no assumptions about. If this is not true — that is, if buyers with higher values always end up accepting higher prices, the stochastic monotonicity assumption (though satisfied) will not tighten the bounds.

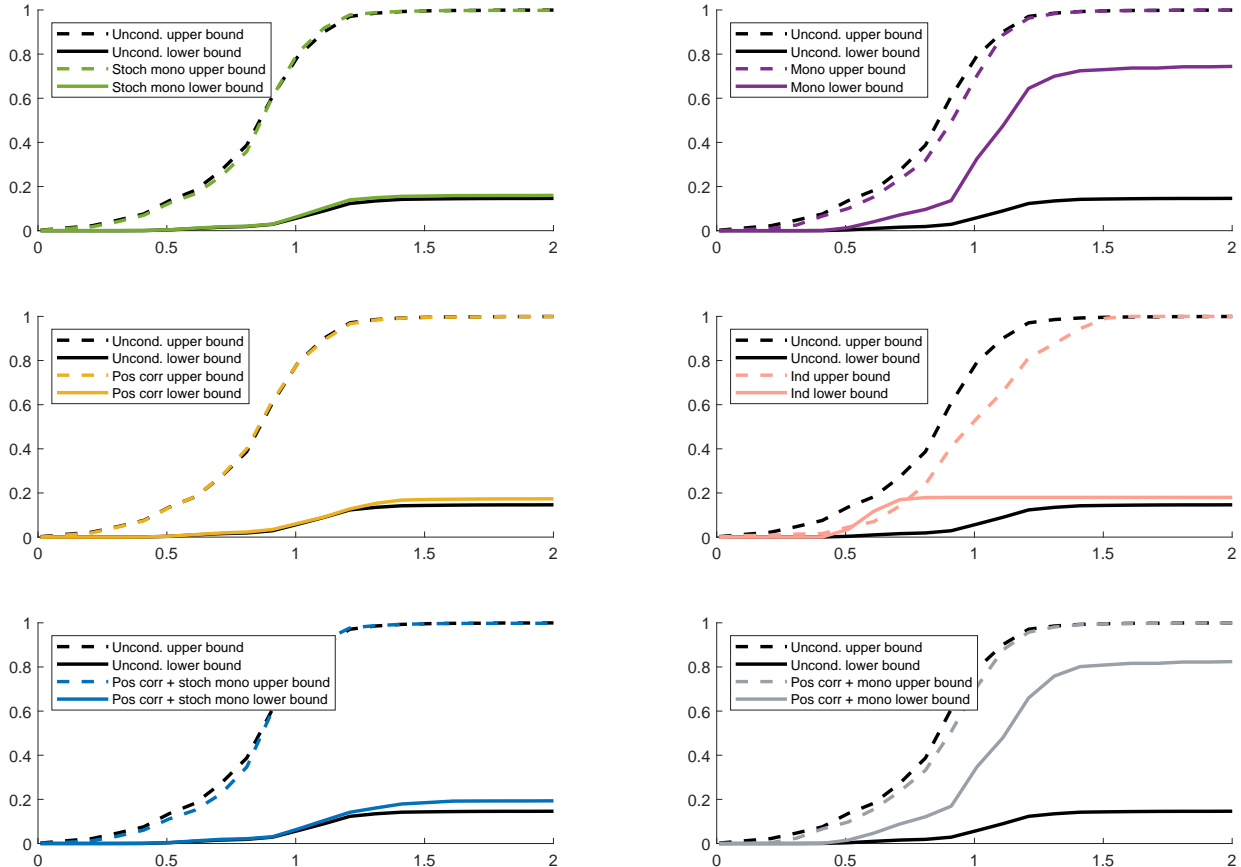
A similar argument applies to the positive correlation bounds (left middle panel). These bounds rely on the assumption that the buyer’s value is stochastically increasing in the seller’s first offer, and will only lead to a tighter upper bound, say, if some games in which the seller makes a relatively high first price end with a buyer quitting at a relatively low offer. If this is not the case, the positive correlation bounds, while not rejected by the data, will do little to improve the bounds.

We reiterate that all of our bounds are sharp, which implies that they are the best possible bounds under their corresponding assumptions. Any tightening of the bounds necessarily requires stronger assumptions. The monotonicity bounds (shown in the top right panel), illustrate this point, as they drastically improve upon the unconditional bounds, especially the lower bound. This suggests that there are buyers in the data who propose relatively high P_2^B and yet end up accepting relatively low offers later in the game.

Figure 2 also illustrates the potential for bounds to cross, especially under somewhat stronger assumption such as independence (middle right panel). Here we observe the lower bound crossing the upper at low buyer values. Therefore, this assumption is not consistent with our data. This assumption does not allow for any unobserved game-level heterogeneity, such as features of the cell phone that both the buyer and seller observe (e.g., a cracked screen in a listing photo or a protective covering included with the phone). Such heterogeneity can generate positive correlation (from the econometrician’s perspective) in the buyer’s value and seller’s first offer, and indeed our positive correlation bounds, which allow for this type of heterogeneity, do not cross. This highlights the

¹⁵The upper bound will thus, by construction, be surjective (i.e., mapping to each value in $[0,1]$).

Figure 2: Bounds on Buyer Distribution for Cell Phone



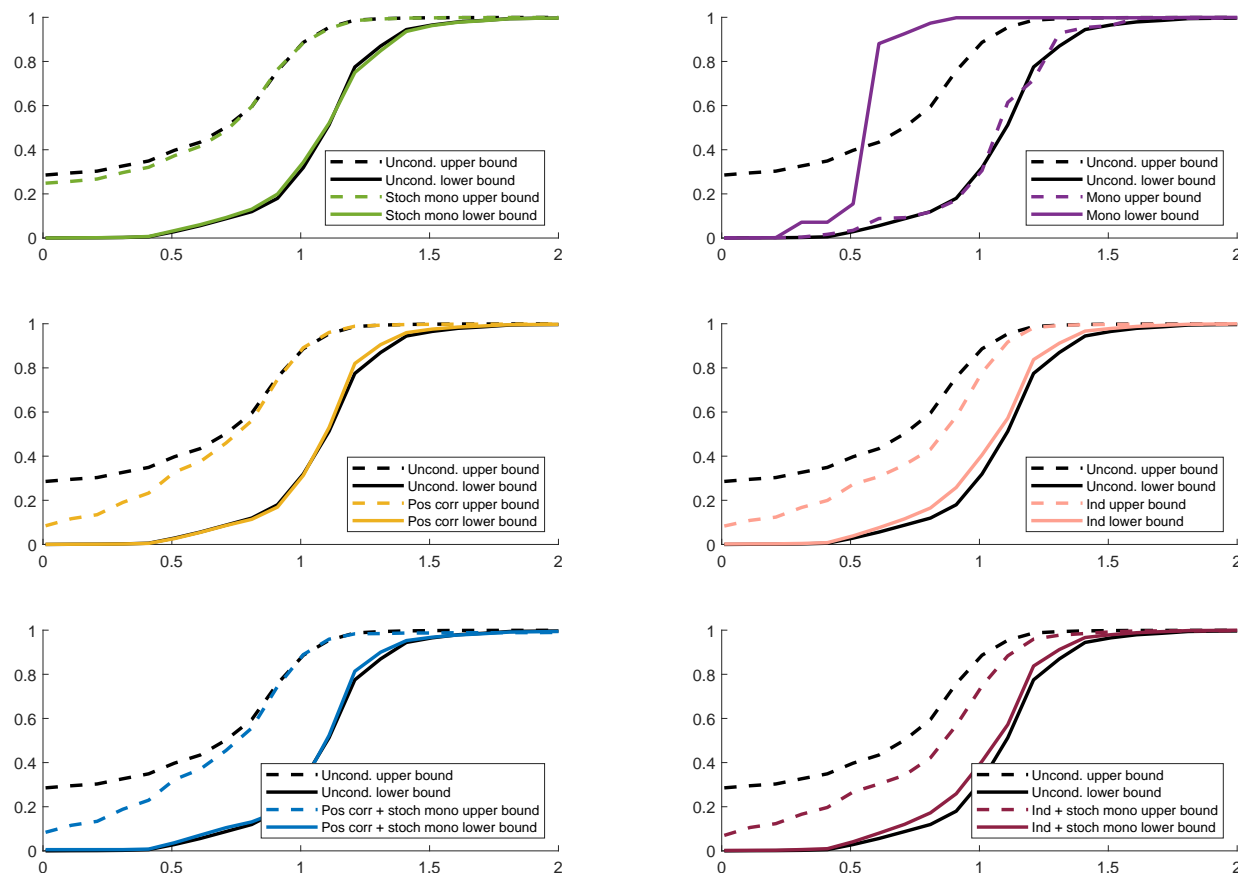
Notes: Each panel displays the estimated bounds on the buyer distribution for the most popular cell phone product in the data. The top two panels show the stochastic monotonicity bounds (left) and monotonicity bounds (right). The middle panels show the positive correlation bounds (left) and independence bounds (right). The bottom panels show the combined positive correlation + stochastic monotonicity bounds (left) and combined positive correlation + monotonicity bounds (right). Every panel also shows the unconditional bounds for comparison. In each panel, upper bounds are shown with dashed lines and lower bounds with solid 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.

value of our moderately weak bounds, which the data from this product do not reject. The bottom right panel of Figure 2 displays the combined positive correlation and monotonicity bounds — the tightest bounds we can impose for this product that do not cross. These bounds improve slightly upon the monotonicity bounds alone.

In Figure 3 we display analogous estimates for the seller distribution using the same set of cell phone listings. Here we observe that the bounds are much narrower overall, and the upper bound is less tightly estimated than the lower bound (whereas the opposite occurs for the buyer distribution). This is because it is the upper bound on the seller value distribution that can be identified by accept and counter behavior of the seller (and the seller always has at least one such price — the list price,

P_1^S) and it is the upper bound that is identified by quitting behavior of the seller. Quits for the seller will not be observed if the negotiation ends in agreement or if the buyer ends the negotiation by quitting. The stochastic monotonicity and positive correlation assumptions, as well as their combination, shown in the three left panels, help to tighten this upper bound slightly.

Figure 3: Bounds on Seller Distribution for Cell Phone



Notes: Each panel displays the estimated bounds on the seller distribution for the most popular cell phone product in the data. The top two panels show the stochastic monotonicity bounds (left) and monotonicity bounds (right). The middle panels show the positive correlation bounds (left) and independence bounds (right). The bottom panels show the combined positive correlation + stochastic monotonicity bounds (left) and combined independence + stochastic monotonicity bounds (right). Every panel also shows the unconditional bounds for comparison. In each panel, upper bounds are shown with dashed lines and lower bounds with solid 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 the seller distribution for this product, we observe that the monotonicity assumption is grossly violated. This can be seen by the solid line (the lower bound) lying above the dashed line (the upper bound), and above the unconditional upper bound. This finding is the opposite of what we found for the buyer distribution. This is because, for the seller distribution, the monotonicity assumption is particularly strong. It requires that a seller with a higher first offer must have a

higher value, and hence behavior across sellers with different first offers can be used to tighten the bound on any seller’s value. The monotonicity assumption in the buyer case is weaker, requiring only that, *conditional* on the seller’s first offer, a buyer’s value be higher at higher buyer offers. This finding — that seller own-offer monotonicity is violated — highlights the importance of our weaker assumption (stochastic monotonicity), which is not rejected by the data.

The independence assumption for the seller distribution is not violated, unlike the buyer distribution case. This is because the independence assumption for the seller is weaker than for the buyer. It requires only that the seller value be independent of the buyer’s first offer (P_2^B) *conditional* on the seller’s first offer (P_1^S). This weaker assumption allows the bounds to capture some degree of unobserved game-level heterogeneity, for example, that would violate buyer independence.

In the bottom right panel, we display the tightest bounds for this product that do not cross, which are the combined independence and stochastic monotonicity assumptions. The Monte Carlo simulations in Appendix C demonstrate that any of these bounds — even the unconditional bounds — can be quite narrow or wide, depending on features of the data.

5.2. Exploring All Products. We now apply our bounds to all 363 products in our sample. In Table 2, we show, under each assumption or set of assumptions, the fraction of products for which a violation of the bounds occurs. Here we define a violation as any product for which the estimated upper bound crosses the estimated lower bound at any point other than in the extreme tails.¹⁶ Such a crossing indicates that the assumption(s) underlying the bounds is violated. We also compute the *integrated violation error* (IVE), which takes on values from 0 to 1 and measures the average difference between the upper and lower bound in cases where they cross.¹⁷

Table 2 shows that the results from Figures 2–3 are quite representative of products in our data. In particular, we find that the monotonicity bounds for seller values (Assumption 2.i) cross for 96% of products, with the upper monotonicity bound violating the lower monotonicity bound by an average of 16.8% (the IVE). For buyer values, where the monotonicity assumption is weaker, the bounds never cross. The opposite is true for the independence bounds (Assumption 4), which

¹⁶In order to focus on violations that do not occur primarily at the tails of the distribution, Table 2 only considers violations that occur at points in the support where the upper bound is above 0.05 and the lower bound is below 0.95. Thus, denoting a generic upper and lower bound by F^U and F^L , we consider the bounds as crossing at least once if there exists some x such that $F^U(x) \geq 0.05$, $F^L(x) \leq 0.95$, and $F^U(x) < F^L(x)$.

¹⁷For a generic upper and lower bound by 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 the unconditional lower bound in the case of sellers and the unconditional buyer bound in the case of buyers. We choose these distributions as they are surjective on $[0, 1]$ (mapping to every point in $[0, 1]$), as the seller lower bound depends on P_1^S and the buyer upper bound depends on P_2^B , which are both always observed in our sample.

Table 2: Measuring Bound Crossings Under Different Assumptions

	Seller Bounds		Buyer Bounds	
	Fraction	Integrated	Fraction	Integrated
	Bounds	Violation	Bounds	Violation
	Crossing	Error	Crossing	Error
Unconditional (A1)	0	0	0	0
Monotonicity (A2)	0.9642	0.1683	0	0
Stochastic Monotonicity (A3)	0	0	0	0
Independence (A4)	0.0248	0.0001	0.4876	0.0246
Positive Correlation (A5)	0.0055	0	0	0
Pos. Corr. & Stoch. Mon. (A5 + A3)	0.0055	0	0	0
Indep. & Stoch. Mon. (A4 + A3)	0.0689	0.0004	0.4298	0.0199
Pos. Corr. & Mon. (A5 + A2)	0.9642	0.1683	0.0055	0
Indep. & Mon. (A4 + A2)	0.9642	0.1683	0.8595	0.1546

Notes: Across all products in the full sample, table shows the fraction of products for which seller lower bound crosses the upper bound, as well corresponding average integrated violation error of these crossings. Table shows similar quantities for the buyer bounds.

cross for 48.8% of products when bounding buyer values but only 2.5% of products for seller values, where the assumption is weaker. Even with 48.8% of products exhibiting at least one crossing under the buyer independence assumption, the IVE is only 2.5% on average. However, when combined with monotonicity (the final row of Table 2), the IVE is 15.5% on average for buyer values. Bounds based on combinations of assumptions naturally cross at least as often as any of the the underlying bounds. Our moderate assumptions, such as stochastic monotonicity (Assumption 3), never cross for buyer or seller values for any product. Positive correlation bounds (Assumption 5) never cross for the buyer CDF, and cross for fewer than 1% of products for the seller CDF, with an IVE that rounds to zero.

In Table 3, for which a given set of bounds, we display statistics across products on the tightness of the bounds, limiting to products for which the bounds do not cross. We first compute the average width of the bounds for a given product by integrating the upper bound minus the lower bound.¹⁸ This average width metric is similar to the IVE, ranging from 0 to 1, with a lower number meaning the bounds are tighter.

We find that the unconditional seller bounds can be relatively tight for some products, with an average probability gap of 0.328 for the minimum-width product, and quite wide for others, with an average gap that is nearly twice as high (0.631) for the maximum-gap product. The average width for some bounds — such as the stochastic monotonicity bounds — is quite similar to that

¹⁸We integrate this difference 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.

Table 3: Statistics Across Products on Width of Bounds

	Seller Bounds			Buyer Bounds		
	Min	Mean	Max	Min	Mean	Max
Unconditional (A1)	0.328	0.428	0.631	0.350	0.403	0.448
Monotonicity (A2)	0.000	0.008	0.040	0.193	0.303	0.409
Stochastic Monotonicity (A3)	0.320	0.417	0.627	0.330	0.396	0.439
Independence (A4)	0.137	0.281	0.521	0.033	0.176	0.307
Positive Correlation (A5)	0.144	0.377	0.580	0.321	0.388	0.426
Pos. Corr. & Stoch. Mon. (A5 + A3)	0.146	0.373	0.582	0.316	0.387	0.432
Indep. & Stoch. Mon. (A4 + A3)	0.121	0.266	0.510	0.061	0.190	0.329
Pos. Corr. & Mon. (A5 + A2)	0.000	0.008	0.040	0.179	0.282	0.393
Indep. & Mon. (A4 + A2)	0.000	0.008	0.040	0.012	0.137	0.269

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.

of the unconditional bounds. Given that these bounds (and all of the bounds we study) are sharp under their corresponding assumptions, this implies that a bargaining model may need to satisfy stronger assumptions in order to lead to tighter bounds. We find that independence and positive correlation assumptions for the seller do much to improve the bounds for some products, decreasing the minimum average width to 0.14. For the buyer bounds, monotonicity drastically improves the tightness of the bounds over the unconditional bounds, with a minimum average width of 0.193.¹⁹

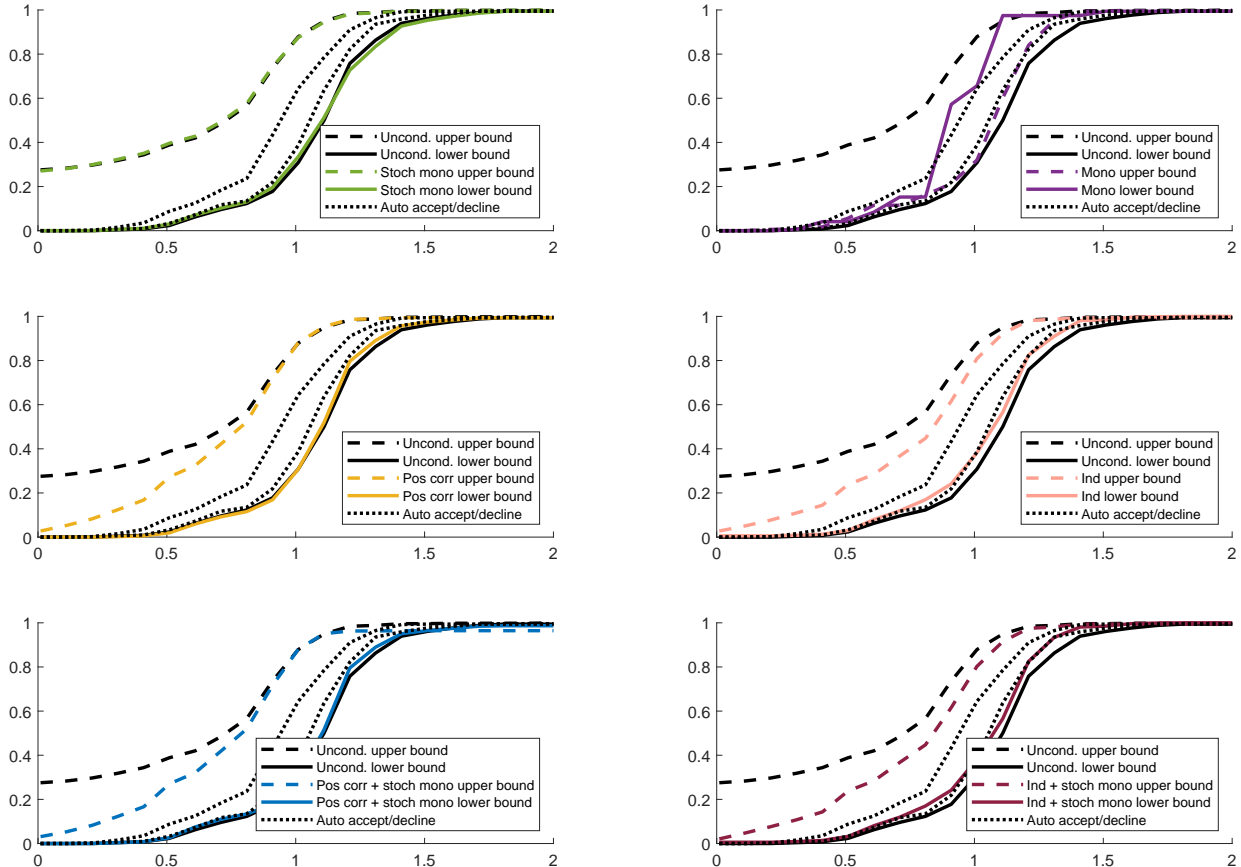
5.3. Auto-Accept/Decline Prices: A Novel Check on Validity. We now exploit the additional information available in sellers’ auto accept/decline prices. This information serves as a valuable check on the validity of our bounds — a piece of private information known to the seller (and reported secretly to the platform) that offers an immediate upper and lower bound on the true CDF of seller values. Indeed, to a degree, this allows us to partially identify the true CDF of seller values and compare this to our bounds.

In Figure 4, we display the same bounds as in Figure 3, for the same cell phone product, but estimated using only the subset of observations for this product for which the seller reported an auto-accept and auto-decline price. The empirical CDF of these auto-accept and auto-decline prices are shown in gray dotted lines in every panel. We observe that each of our assumptions shown in Figure 4 are consistent almost everywhere with the tight bounds implied by the secret prices.

In Table 4, we extend this auto accept/decline analysis to the full set of products that each

¹⁹The seller monotonicity and buyer independence bounds also drastically improve upon the unconditional bounds, but as highlighted above, these bounds cross for many products and the Table 3 results only reflect those products for which bounds do not cross.

Figure 4: Bounds on Seller Distribution Compared to Auto Accept/Decline Prices



Notes: This figure is equivalent to Figure 3 but estimated using only the subset of observations for this product for which the seller reported a non-zero auto-accept and auto-decline price. The empirical CDF of these auto-accept and auto-decline prices are shown in gray dotted lines in every panel. All prices are scaled by the reference price for the product, and thus units on the horizontal axis are fraction of the reference price.

have at least 100 sequences in which auto accept/decline prices are recorded, which includes 16 products. Here we record a crossing of the seller lower bound as a case where it lies above the CDF of auto-decline prices (which is an upper bound on the true seller value distribution). Similarly, we define a crossing of the seller upper bound as a case where it lies below the CDF of auto-accept prices (which is a lower bound on the true seller CDF). We count these violations and compute the IVE, as in Table 2.

In Table 4, we observe a large fraction of products in which the seller monotonicity bounds violate the auto-decline or auto-accept CDFs, which is not surprising given that the monotonicity bounds themselves frequently cross. We observe that our moderate assumptions, such as stochastic monotonicity, positive correlation, or the combination of the two, even if they cross the auto-accept/decline prices, do so by very little, having low IVE (well below 1%). The seller independence

Table 4: Measuring Bound Crossings Relative to Auto Accept/Decline Prices

	Lower Bound v. Auto-Decline		Upper Bound v. Auto-Accept	
	Fraction Bounds Crossing	Integrated Violation Error	Fraction Bounds Crossing	Integrated Violation Error
	Unconditional (A1)	0	0	0
Monotonicity (A2)	0.9375	0.0493	0.875	0.0085
Stochastic Monotonicity (A3)	0	0	0	0.0001
Independence (A4)	0.4375	0.0047	0.0625	0
Positive Correlation (A5)	0.125	0.0001	0	0
Pos. Corr. & Stoch. Mon. (A5 + A3)	0.125	0.0002	0.0625	0.0001
Indep. & Stoch. Mon. (A4 + A3)	0.5	0.0059	0.0625	0.0002
Pos. Corr. & Mon. (A5 + A2)	0.9375	0.0493	0.875	0.0085
Indep. & Mon. (A4 + A2)	0.9375	0.0493	0.875	0.0085

Notes: Among products with at least 100 sequences for which the seller reported an auto-accept and auto-decline price, table shows the fraction of products for which seller lower bound crosses the auto-decline price CDF and the corresponding average integrated violation error of these crossings. Table shows similar quantities for the seller upper bound compared to the auto-accept price CDF.

assumptions cross the auto-decline price CDF for 44% of products, but not by much: the IVE is again less than 1%. Table 4 lends credence to our bounds and the assumptions underlying them — in particular for the moderate assumptions.

6 Quantifying Inefficient Impasse and Uncertainty

We now consider bounds on the counterfactual first-best probability of trade. In a first-best world, a buyer with value B and seller with value S will trade whenever $S \leq B$. The Myerson and Satterthwaite (1983) Theorem demonstrated that achieving the surplus offered by such a mechanism is infeasible when buyers and sellers have incomplete information and overlapping support of values. Cramton (1992) demonstrated that, while the first-best *surplus* is infeasible, it is possible for negotiators to achieve the first-best *quantity* of trade after costly delay. The first-best quantity of trade will be weakly higher than the realized volume of trade in the data; the question is, how much higher? A lower bound of $P(S \leq B)$ can be compared to the probability of sale in the data to examine how much inefficient impasse occurs in the data that would be avoided in a first-best world.²⁰

²⁰Note that the gains from trade (i.e. surplus) and the probability of trade are intimately related; the former is weakly increasing in the latter. The probability of trade is a much more useful object for our empirical purposes because the real-world probability of trade is observable in the data, whereas the real-world surplus is not. Specifically, whenever trade occurs, trade must be the efficient outcome, but the size of those gains is not necessarily identified. Whenever trade fails even though gains from trade exist, the outcome is necessarily inefficient, but the size of the loss is not necessarily identified.

An upper bound on $P(S \leq B)$ contains different information than a lower bound. An upper bound offers a metric for how much uncertainty agents have about whether gains from trade exist. An upper bound that is well below 1 implies that, even in a first-best world, negotiations would sometimes fail. Why would agents still negotiate, even if the first-best outcome is to disagree? The only justification for such behavior is that agents must be uncertain about whether gains from trade actually exist (i.e., about whether $S \leq B$). If agents were to have complete information, they would not engage in bargaining in cases where gains from trade do not exist.²¹

As important as the object $P(S \leq B)$ is for studying the impasse and uncertainty, existing empirical tools and theoretical models are insufficient for identifying it. Real-world bargaining data from settings where agents have private information will not typically contain data on those private values themselves (S and B). Our bounds herein build on the assumptions and bounds on marginal distributions derived above. We describe each of these assumptions in turn and display estimated bounds on $P(S \leq B)$ under these assumptions in Table 5, with lower bounds in panel A and upper bounds in panel B. Below each estimate, we display 95% bootstrapped confidence intervals based on the method of Fang and Santos (2018).²² Here we focus on the eight most popular products — the same products shown in Table 1. Columns 1 through 4 order the bounds from those relying on our weakest to strongest assumptions.

Our first set of bounds rely only on our weakest assumption, Assumption 1, and as such are generally uninformative, corresponding to $[P(\text{sale}), 1]$, where $P(\text{sale})$ is the fraction of negotiations in the data that end in trade. To see this, note that Assumption 1 implies $P(S \leq B) \geq P(X_{AC}^S \leq X_{AC}^B)$. 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, $P(X_{AC}^S \leq X_{AC}^B)$ typically corresponds to $P(X_{AC}^S = X_{AC}^B)$, which is equal to $P(\text{sale})$, as this represents cases where one agent accepts a price the other proposes. For the upper bound, Assumption 1 implies $P(S \leq B) \leq P(X_Q^S \leq X_Q^B)$. This latter probability is always equal to 1, because only one party (the buyer or seller, not both) can quit in a given negotiation. When the seller quits, $X_Q^B = \infty$, and when the buyer quits, $X_Q^S = 0$. Thus, under our unconditional bounds, the most we learn is that $P(S \leq B) \in [P(\text{sale}), 1]$. In

²¹If an agent has incomplete information about her opponent’s value but knows that the support of possible buyer values lies fully above that of seller values (what is referred to in Fudenberg and Tirole (1991) as the “gap” case), then the first-best probability of trade would also be $P(S \leq B) = 1$.

²²Our estimators involve maxima and minima of estimated conditional mean functions and the asymptotic distributions of the estimated upper and lower bounds are therefore non-normal. Even though the estimated bounds are not fully differentiable functions of the estimated conditional mean functions, they are directionally differentiable, and we can obtain confidence intervals by using an extension of the delta method. We calculate the directional derivatives analytically and use 500 bootstrap samples.

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

A. Lower Bounds		(1)	(2)	(3)	(4)
	n	A1 ($P(sale)$)	A6 & A2.ii	A7 & A2.ii	A2.i & A2.ii
Electronics	577	0.319	0.348	0.622	0.842
		[0.281,0.357]	[0.317,0.400]	[0.583,0.662]	[0.813,0.872]
Cameras	159	0.377	0.384	0.509	0.679
		[0.302,0.453]	[0.296,0.446]	[0.431,0.587]	[0.606,0.752]
Sports	190	0.311	0.321	0.563	0.768
		[0.245,0.376]	[0.238,0.380]	[0.492,0.634]	[0.708,0.829]
Video Games	487	0.29	0.328	0.472	0.739
		[0.249,0.330]	[0.274,0.355]	[0.428,0.517]	[0.700,0.778]
Musical	123	0.724	0.794	0.902	0.894
		[0.645,0.803]	[0.721,0.883]	[0.850,0.955]	[0.840,0.949]
Home/Garden	150	0.207	0.211	0.313	0.593
		[0.142,0.271]	[0.142,0.269]	[0.239,0.388]	[0.514,0.672]
Cell Phones	2501	0.13	0.142	0.409	0.933
		[0.117,0.143]	[0.134,0.161]	[0.390,0.429]	[0.923,0.943]
Computers	497	0.223	0.227	0.374	0.757
		[0.187,0.260]	[0.196,0.263]	[0.332,0.417]	[0.719,0.794]
B. Upper Bounds		(1)	(2)	(3)	(4)
	n	A1	A6 & A2.ii	A7 & A2.ii	A2.i & A2.ii
Electronics	577	1	0.998	0.917	0.875
			[0.996,0.998]	[0.894,0.939]	[0.848,0.902]
Cameras	159	1	1	0.981	0.95
			[1.000,1.000]	[0.960,1.000]	[0.916,0.984]
Sports	190	1	1	0.979	0.958
			[1.000,1.000]	[0.958,0.999]	[0.929,0.987]
Video Games	487	1	0.994	0.867	0.77
			[0.988,0.994]	[0.836,0.897]	[0.733,0.807]
Musical	123	1	1	1	0.927
			[1.000,1.000]	[1.000,1.000]	[0.881,0.973]
Home/Garden	150	1	1	0.927	0.847
			[1.000,1.000]	[0.885,0.969]	[0.789,0.905]
Cell Phones	2501	1	0.992	0.677	0.487
			[0.985,0.992]	[0.659,0.695]	[0.468,0.507]
Computers	497	1	0.995	0.94	0.769
			[0.989,0.999]	[0.919,0.961]	[0.731,0.806]

Notes: For the most popular products with each category, table displays lower bounds (panel A) and upper bounds (panel B) on the counterfactual first-best probability of trade ($P(S \leq B)$) under different assumptions. Bootstrapped 95% confidence intervals are shown in square braces below each estimate.

column 1 of Table 5, we therefore display $P(sale)$ in panel A and 1 in panel B.

The bounds in Table 5 that rely on our strongest assumptions are those shown in column 4, where we maintain Assumption 2, monotonicity, for both the buyer *and* the seller. In light of the results shown in Section 5, monotonicity is a reasonable assumption for buyers but not for sellers. We nonetheless estimate bounds assuming monotonicity for *both* parties to illustrate that these bounds are indeed too strong, and can cross. Under this assumption, bounds on $P(S \leq B)$ are

given by $[P(X_{AC}^{S*} \leq X_{AC}^{B*}), P(X_Q^{S*} \leq X_Q^{B*})]$. We display these bounds in column 4 of Table 5. Taking these estimates at face value, we see evidence of gross inefficiencies in these markets. For example, for the cell phone product, sales occur in the data with probability 0.13, but the bounds imply that trade would occur 93% of the time in a first-best world. The upper bounds under these assumptions, shown in Panel B, suggest that incomplete information is present for each of these products: the confidence intervals are well below 1. However, the cell phone row illustrates that these bounds can cross — the upper bound is 0.487 — indicating a violation of the monotonicity assumption. In a number of other rows the confidence intervals cross for the lower and upper bounds. This is not surprising given that Table 2 shows that seller monotonicity (Assumption 2.i) is violated for most products.

We next examine bounds that drop the assumption of seller monotonicity and instead rely on buyer monotonicity (Assumption 2.ii) and on the following weak assumption on $S - B$:

Assumption 6. $P(S - B \leq x \mid P_1^S = y, P_2^B = z)$ is increasing in z for all y .

Sufficient conditions for Assumption 6 are a strict version of Assumption 2.ii (buyer monotonicity) and Assumption 4.i (seller independence), two assumptions for which we do not find large crossings in Section 5.²³ While implied by buyer monotonicity and seller independence, the assumption is far weaker than these two. It is akin to our stochastic monotonicity assumption on *values* applied instead to the *difference in values*. It does not directly rely on any assumptions about the correlation structure between S and B . Indeed, any restriction on that correlation has little bite when examining the difference $S - B$.

We display estimates of bounds on $P(S \leq B)$ under these assumptions in column 2 of Table 5.²⁴ For each product, the point estimates of the lower bound are higher than $P(\text{sale})$ observed in the data. For example, for the electronics product, the probability of sale in the data is 0.319, and the counterfactual first-best probability of trade in column 2 is 0.348, suggesting that the real-world bargaining misses some efficient trades. However, the confidence intervals for most products in panel A contain $P(\text{sale})$, and thus the evidence of inefficient impasse under this assumption

²³To see that these assumptions are sufficient, suppose we can write $P_2^B = f(B, P_1^S)$ where $f(\cdot, P_1^S)$ is increasing for all P_1^S with inverse function $g(\cdot, P_1^S)$. Then the conditional probability statement in Assumption 6 can be written $P(S - B \leq x \mid P_1^S = y, B = g(z, y)) = P(S - g(z, y) \leq x \mid P_1^S = y, B = g(z, y))$. This latter statement is equivalent to $P(S - g(z, y) \leq x \mid P_1^S = y)$, which is increasing in z for all y .

²⁴Estimation of these bounds is similar to that of the marginal distribution bounds. We first estimate $P(X_{AC}^S \leq X_{AC}^{B*} \mid P_1^S = y, P_2^B = z)$ using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel function and bandwidth $n^{-1/6}$. Denote the estimator by $\hat{m}^{S-B}(y, z)$. The estimated lower bound is then $\frac{1}{n} \sum_{i=1}^n \max_{z' \in \{z: z \geq P_{2,i}^B, Q_{0.05}(P_{2,i}^B) \leq z \leq Q_{0.95}(P_{2,i}^B)\} \cup \{P_{2,i}^B\}} \hat{m}^{S-B}(P_{1,i}^S, z')$. The upper bound estimator is analogous.

is relatively weak. Column 2 of panel B shows that the estimated upper bound on the first-best probability of trade is generally very close to 1 (or even equal to 1 for some products). Together, the lower and upper bound estimates in column 2 offer little information beyond the unconditional bounds in column 1.

The final bounds we consider invoke buyer monotonicity and the following:

Assumption 7. P_2^B is weakly decreasing in $S - B$ conditional on P_1^S .

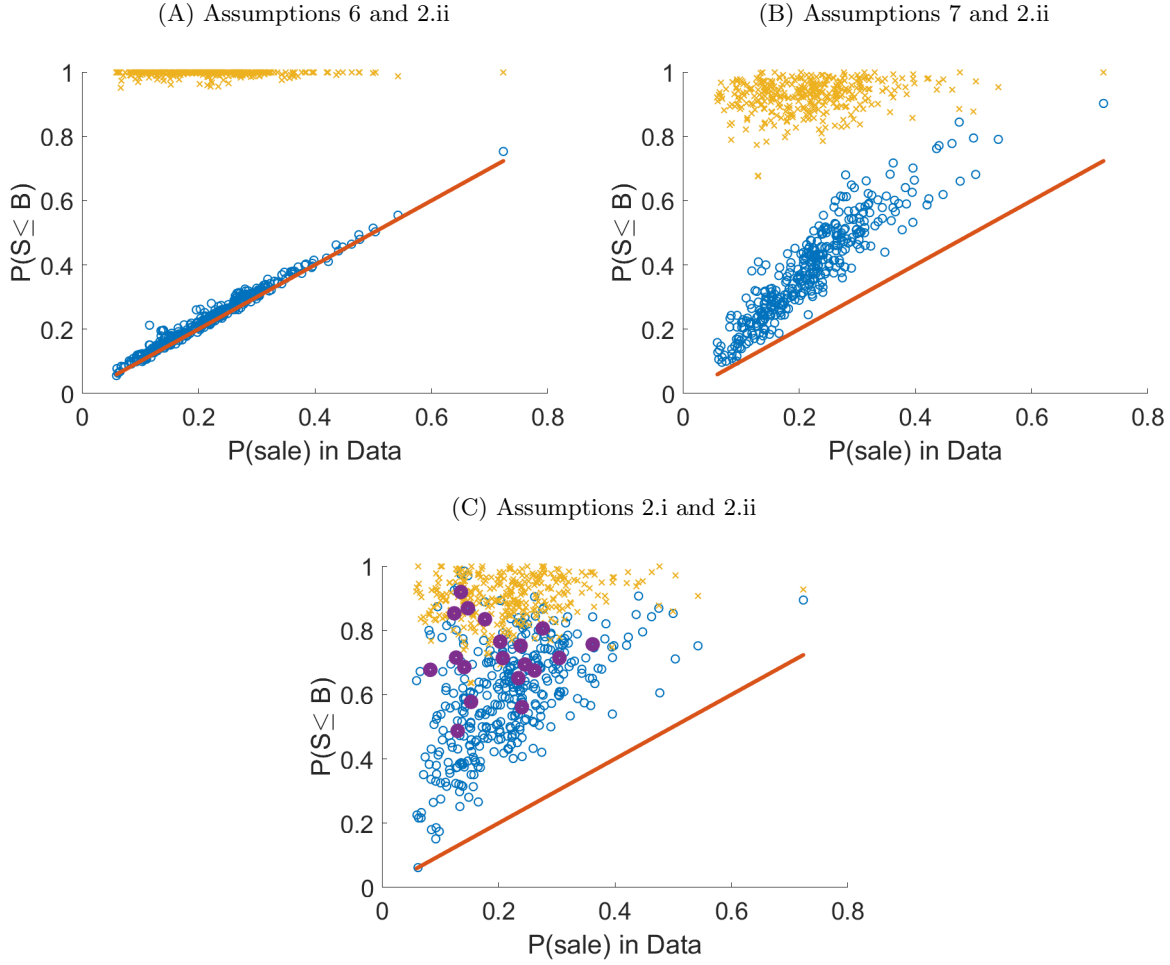
This assumption is a stronger version of Assumption 6, but is still weaker than assuming monotonicity for *both* the buyer and seller: strict versions of Assumption 2.i and 2.ii are sufficient for Assumption 7 to hold, but not necessary.²⁵ As Assumption 6 is akin to a stochastic monotonicity assumption applied to $S - B$, Assumption 7 is akin to weak monotonicity applied to this difference. In column 3 of Table 5, we estimate bounds relying on Assumption 7 and buyer monotonicity.²⁶ Unlike columns 1 and 2, in column 3 we find informative lower bounds for all products, with confidence intervals lying well above $P(\text{sale})$ for all products. We also find informative upper bounds for most products, with confidence intervals lying below 1. We also find that these bounds do not cross, unlike those shown in column 4. These bounds suggest that the real-world bargaining exhibits inefficient impasse, but not as much as implied by the (overly strong) assumption of seller monotonicity. For example, for computers, the probability of trade in the data is 0.223, but in column 3 we find that it would be as high as 0.374 in a first-best world, suggesting that, when the buyer values the computer more than the seller, the pair still fails to reach an agreement 40% of the time (i.e., $1 - 0.223/0.374$). We also find that at least 6% of negotiations would end in impasse even in a first-best world, as the upper bound estimate is 0.94, suggesting that agents do indeed negotiate under uncertainty (and that assuming Nash bargaining, for example, would be incorrect and potentially yield misleading conclusions).

In Figure 5, we extend this analysis to each of our 363 products. In each panel, we order products on the horizontal axis according to their sales probability in the data. On vertical axes, an “x” represents the estimated upper bound and a hollow circle the estimated lower bound. These hollow circles are made solid for products where the estimated bounds cross (i.e., where the lower

²⁵To see that these assumptions are sufficient, suppose we can write $P_1^S = f_1(S)$ and $P_2^B = f_2(B, P_1^S)$ where $f_1(\cdot)$ and $f_2(\cdot, P_1^S)$ are increasing with inverse functions $g_1(\cdot)$ and $g_2(\cdot, P_1^S)$, respectively. Then $S - B = g_1(P_1^S) - g_2(P_2^B, P_1^S)$ and, conditional on P_1^S , this latter difference is a decreasing function of P_2^B .

²⁶Estimation of the bounds in column 3 of Table 5 is similar to those in column 2. First define $\hat{X}_{AC}^{S-B*}(y, z) \equiv \min_{i: P_{2,i}^B \leq z, P_{1,i}^S \in N(y)} (X_{AC,i}^S - X_{AC,i}^B)$, where the neighborhood $N(y)$ is as in Section 4. The estimated lower bound is then $\frac{1}{n} \sum_i \mathbf{1}(\hat{X}_{AC}^{S-B*}(P_{1,i}^S, P_{2,i}^B) \leq 0)$. The estimator of the upper bound is analogous.

Figure 5: Bounds on $P(S \leq B)$, All Products



Notes: Figure displays upper bounds (marked with “x”) and lower bounds (marked with hollow circles) for the counterfactual first-best probability of trade ($P(S \leq B)$) under different assumptions, for each product in the full sample. Each panel ranks products on the horizontal axis by the product’s sale probability in the data. The solid line represents the 45-degree line. In panel A the assumptions are Assumptions 6 and 2.ii; in panel B they are Assumptions 7 and 2.ii; and in panel C they are Assumptions 2.i and 2.ii.

bound lies above the upper bound). Panel A displays bounds relying on Assumptions 6 and 2.ii, corresponding to column 2 of Table 5. Panel B displays bounds relying on Assumptions 7 and 2.ii, corresponding to column 3 of Table 5. Panel C displays bounds relying on Assumptions 2.i and 2.ii, corresponding to column 3 of Table 5. As in Table 5, Panel A shows that the estimates relying on Assumptions 6 and 2.ii are too weak to be informative, as they correspond closely to $P(\text{sale})$ and 1. In Panel C, we find that assuming monotonicity for both agents yields bounds that are tighter, but violated for many products (42 out of 363).

The Goldilocks-like assumptions are those in panel B, where we observe informative bounds

that do not cross. Here the lower bounds suggest that the real-world bargaining indeed exhibits inefficient impasse, and to a large degree for some products. Across all products, the percentage of negotiations in which gains from trade exist and yet agents disagree ranges from 8.3% to 68.9%. For the median product, this inefficient impasse occurs 43% of the time. The upper bounds in panel C offer evidence that agents in these markets face uncertainty about whether efficient trade is possible. The upper bound ranges from 0.673 to 1.000, depending on the product, with a median of 0.936. This suggests that, even in a first-best world, agents would only trade at most 94% of the time, and thus must face some uncertainty about the gains from trade. For most products, this upper bound point estimate is below 1.

7 Discussion and Conclusion

This study provides bounds on the private-value distributions of buyers and sellers and on the first-best probability of trade from real-world, sequential-offer bargaining data from eBay’s Best Offer platform. These bounds are sharp, nonparametric, and robust to the presence of two-sided uncertainty (i.e. both buyer and seller may have private values). Our approach relies on revealed preferences arguments and other assumptions on behavior without specifying a full model of equilibrium play. The assumptions we invoke range from quite weak to strong. We find that 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 bounds 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.²⁷ The bounds we propose do not rely on any specification of beliefs, equilibrium refinement, or equilibrium selection, allowing us to study how well bargaining performs in this real-world market (eBay) without strongly constraining

²⁷In the earliest game-theoretic work on incomplete-information bargaining, a number of equilibrium refinements were introduced, but these were necessarily ad hoc in nature, and in some cases equilibria satisfying these restrictions can fail to exist (see discussion in Ausubel et al. 2002). Binmore et al. (1992) emphasized that bargaining models were one of the main motivations behind much of the refinements literature.

the answer a priori.

Given that the bounds we propose rely only on very weak assumptions, they are naturally wide in some cases. We demonstrate that, in spite of this, the bounds can be useful in allowing us to examine what behavioral properties are satisfied in real-world bargaining. The bounds also allow us to examine questions of impasse and uncertainty — questions that are not possible to address under complete-information frameworks such as Nash bargaining: when trade fails, how often is it the case that the buyer actually values the good more than the seller, and hence the parties should have agreed in a fully efficient world? Under our moderate assumptions, we find evidence of inefficient impasse: for the median product, at least 43% of efficient trades never occur. 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 strong assumptions and focus on our more moderate assumptions.

It is possible that the assumptions we use — even those that are not the strongest — are still too strong. There are relatively few empirical estimates 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 find that participants fail to trade in 46% of cases where gains from trade exist. They found that this impasse is reduced substantially (to 15%) when negotiators are allowed to communicate. Bochet et al. (2020) and Huang et al. (2020) also studied two-sided incomplete-information experimentally, finding a corresponding level of inefficient impasse of 30% and 17%, respectively. Larsen (2021), studying professionals negotiating over used-car inventory, found that at least 21% of first-best trades fail, and Larsen et al. (2021) found that skilled mediators substantially reduce this inefficient breakdown.²⁸ 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 useful benchmarks to which additional studies of bargaining in various contexts may be compared.

²⁸The percentages reported in this paragraph are found in (or can be constructed from) Table 1 of Valley et al. (2002), Figure 3 of Bochet et al. (2020), Table 2 of Huang et al. (2020), and Table 3 of Larsen (2021). Note that in Huang et al. (2020), gains from trade always exist, which is not the case for the other experimental studies.

We view these results as highlighting the importance of allowing for realistic features such as two-sided incomplete information when analyzing bargaining empirically. Analysis allowing for these features has thus far been scarce. Indeed, our analysis and results stand in stark contrast to the standard approach to studying bargaining settings in the empirical economics literature (i.e. complete information/Nash bargaining), which does not allow for the possibility of inefficient impasse.

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.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.
- Binmore, K., Osborne, M., and Rubinstein, A. (1992). Noncooperative models of bargaining. *Handbook of Game Theory*, 1:179–225.
- Bochet, O., Khanna, M., and Siegenthaler, S. (2020). Beyond dividing the pie: An experimental study on bargaining over multiple issues. Available at SSRN 3414288.
- Bodoh-Creed, A., Boehnke, J., and Hickman, B. (2021). How efficient are decentralized auction platforms? *Review of Economic Studies*, 88(1):91–125.
- Cramton, P. (1992). Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59:205–225.
- Desai, P. S. and Jindal, P. (2020). Does bargaining increase product valuation? The upside of bargaining costs. Working paper, available at SSRN: <https://ssrn.com/abstract=3756608>.
- Fang, Z. and Santos, A. (2018). Inference on directionally differentiable functions. *Review of Economic Studies*, 86(1):377–412.
- 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. (2020). Identification in ascending auctions, with an application to digital rights management. *Quantitative Economics*, forthcoming.
- Fudenberg, D. and Tirole, J. (1991). *Game Theory*. Cambridge: MIT Press.

- Grossman, S. and Perry, M. (1986). Sequential bargaining under asymmetric information. *Journal of Economic Theory*, 39(1):120–154.
- 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. and Sorensen, A. (2018). Dynamics and efficiency in decentralized online auction markets. NBER Working Paper 25002.
- Huang, J., Kessler, J. B., and Niederle, M. (2020). Fairness concerns are less relevant when agents are less informed. Working Paper, University of Pennsylvania.
- 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. (2021). Fairness in incomplete-information bargaining: Theory and widespread evidence from the field. NBER Working Paper 29111.
- 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. (2021). Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations. NBER Working Paper 29159.
- Larsen, B. and Zhang, A. (2018). A mechanism design approach to identification and estimation. NBER Working Paper 24837.
- Larsen, B. J. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies*, 88(2):851–882.
- Li, H. and Liu, N. (2015). Nonparametric identification and estimation of double auctions with bargaining. Working paper, Penn State.
- 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.
- 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.
- Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *Annals of Statistics*, 10(4):1040 – 1053.
- 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.

A Proofs

For the proofs in this section we consider a more general setup than in the body of the paper, in which it is possible that the buyer never makes an offer, which is the event $D_2^B \neq C$. This generalization affects seller bounds under Assumption 3 and the buyer bounds under Assumption 2. In case the bounds differ to those in the body of the paper, we state them at the beginning of the corresponding proof.

Proof of Theorem 1.

Proof. Since $X_Q^S \leq S \leq X_{AC}^S$ we have $P(X_{AC}^S \leq x) \leq P(S \leq x) \leq P(X_Q^S \leq x)$. Similarly, since $X_{AC}^B \leq B \leq X_Q^B$, we have $P(X_Q^B \leq x) \leq P(B \leq x) \leq P(X_{AC}^B \leq x)$. \square

Proof of Theorem 2.

Proof. Monotonicity. For the buyer, we will show that

$$\begin{aligned} P(B \leq x) &\geq \int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) + P(X_Q^B \leq x, D_2^B \neq C) \\ P(B \leq x) &\leq \int \mathbf{1}(X_{AC}^{B*}(y, z) \leq x) dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) + P(X_{AC}^B \leq x, D_2^B \neq C) \end{aligned}$$

To do so, first write

$$\begin{aligned} P(B \leq x) &= P(B \leq x, D_2^B = C | D_2^B = C) P(D_2^B = C) + P(B \leq x, D_2^B \neq C) \\ &= \int P(B \leq x | D_2^B = C, P_1^S = y, P_2^B = z) dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) \\ &\quad + P(B \leq x, D_2^B \neq C). \end{aligned}$$

Conditional on $P_1^S = y$, $P_2^B = z$, and $D_2^B = C$, we have $X_{AC}^{B*}(y, z) \leq B \leq X_Q^{B*}(y, z)$ and $X_{AC}^{B*}(y, z)$ and $X_Q^{B*}(y, z)$ are non-random. In addition $X_{AC}^B \leq B \leq X_Q^B$. Therefore,

$$P(B \leq x) \geq \int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) + P(X_Q^B \leq x, D_2^B \neq C).$$

We obtain the upper bound analogously.

For the seller, note that, conditional on y and by Assumption 2, we have $X_Q^{S*}(y) \leq S \leq X_{AC}^{S*}(y)$.

Also note that $X_Q^{S*}(y)$ and $X_{AC}^{S*}(y)$ are non-random. Therefore, applying (1), we have

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

Stochastic Monotonicity. For the buyer, we will show that

$$\begin{aligned} P(B \leq x) &\geq \int \max_{z' \geq z} m_Q^B(x, y, z') dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) + P(X_Q^B \leq x, D_2^B \neq C) \\ P(B \leq x) &\leq \int \min_{z' \leq z} m_{AC}^B(x, y, z') dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) + P(X_{AC}^B \leq x, D_2^B \neq C) \end{aligned}$$

To do so, first write

$$\begin{aligned} P(B \leq x) &= \int P(B \leq x \mid D_2^B = C, P_1^S = y, P_2^B = z) dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) \\ &\quad + P(B \leq x, D_2^B \neq C) \\ &= \int \max_{z' \geq z} P(B \leq x \mid D_2^B = C, P_1^S = y, P_2^B = z') dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) \\ &\quad + P(B \leq x, D_2^B \neq C) \end{aligned}$$

Moreover,

$$\begin{aligned} P(B \leq x) &= \int \min_{z' \leq z} P(B \leq x \mid D_2^B = C, P_1^S = y, P_2^B = z') dF_{P_1^S, P_2^B | D_2^B = C}(y, z) P(D_2^B = C) \\ &\quad + P(B \leq x, D_2^B \neq C) \end{aligned}$$

The bounds now follow from $X_{AC}^B \leq B \leq X_Q^B$.

For the seller, we have, by (1) and Assumption 3, $P(S \leq x) = \int \max_{y' \geq y} P(S \leq x \mid P_1^S = y') dF_{P_1^S}(y)$ and $P(S \leq x) = \int \min_{y' \leq y} P(S \leq x \mid P_1^S = y') dF_{P_1^S}(y)$. The seller bounds then follow from $X_Q^S \leq S \leq X_{AC}^S$. □

Proof of Theorem 3.

Proof. Independence. For the seller, we will show that

$$\begin{aligned} P(S \leq x) &\geq \int \max_z P(X_{AC}^S \leq x \mid P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_1^S | D_2^B = C}(y) P(D_2^B = C) \\ &\quad + P(X_{AC}^S \leq x, D_2^B \neq C) \end{aligned}$$

$$P(S \leq x) \leq \int \min_z P(X_Q^S \leq x \mid P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_1^S \mid D_2^B = C}(y) P(D_2^B = C) \\ + P(X_Q^S \leq x, D_2^B \neq C)$$

To do so, first write

$$P(S \leq x) = P(S \leq x \mid D_2^B = C) P(D_2^B = C) + P(S \leq x, D_2^B \neq C)$$

Moreover

$$P(S \leq x \mid D_2^B = C) = \int P(S \leq x \mid P_1^S = y, D_2^B = C) dF_{P_1^S \mid D_2^B = C}(y) \\ = \int \max_z P(S \leq x \mid P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_1^S \mid D_2^B = C}(y)$$

The lower bound then follows from using $S \leq X_{AC}^S$. Analogous arguments with $S \geq X_Q^S$ yield the upper bound.

For the buyer, we have $P(B \leq x) = \max_y P(B \leq x \mid P_1^S = y)$ and $P(B \leq x) = \min_y P(B \leq x \mid P_1^S = y)$, and the bounds now follow from $X_{AC}^B \leq B \leq X_Q^B$.

Positive Correlation. For the seller, we will show that

$$P(S \leq x) \geq \int \max_{z' \geq z} P(X_{AC}^S \leq x \mid P_1^S = y, P_2^B = z', D_2^B = C) dF_{P_1^S, P_2^B \mid D_2^B = C}(y, z) P(D_2^B = C) \\ + P(X_{AC}^S \leq x, D_2^B \neq C) \\ P(S \leq x) \leq \int \min_{z' \leq z} P(X_Q^S \leq x \mid P_1^S = y, P_2^B = z', D_2^B = C) dF_{P_1^S, P_2^B \mid D_2^B = C}(y, z) P(D_2^B = C) \\ + P(X_Q^S \leq x, D_2^B \neq C)$$

To do so, first write

$$P(S \leq x) = P(S \leq x \mid D_2^B = C) P(D_2^B = C) + P(S \leq x, D_2^B \neq C)$$

Moreover

$$P(S \leq x \mid D_2^B = C) = \int P(S \leq x \mid P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_1^S, P_2^B \mid D_2^B = C}(y, z) \\ = \int \max_{z' \geq z} P(S \leq x \mid P_1^S = y, P_2^B = z', D_2^B = C) dF_{P_1^S, P_2^B \mid D_2^B = C}(y, z)$$

The lower bound then follows from $S \leq X_{AC}^S$. Analogous arguments with $S \geq X_Q^S$ yield the upper bound.

For the buyer, we have $P(B \leq x) = \int \max_{y' \geq y} P(B \leq x \mid P_1^S = y') dF_{P_1^S}(y)$ and $P(B \leq x) = \int \min_{y' \leq y} P(B \leq x \mid P_1^S = y') dF_{P_1^S}(y)$. The bounds then follow from $X_{AC}^B \leq B \leq X_Q^B$. \square

B Bounds Based on Combined Assumptions

B.1. Derivation of Bounds Combining Assumptions. We can combine Assumptions 2–5 in several ways to obtain tighter bounds. As with the proofs of the single-assumption bounds in Appendix A, for each of these combined-assumption bounds we prove the general case where it is possible that the buyer quits or accepts in period 2 of the game (i.e., it is possible that $D_2^B \neq C$).

Independence Plus Monotonicity. Combining independence and monotonicity assumptions — our two strongest assumptions — we obtain the following bounds. First, we note that imposing independence *on top of monotonicity* yields no additional information for the seller distribution bounds, but can tighten the buyer distribution bounds.

Theorem 4. *Suppose Assumptions 2 and 4 hold. Then (7) bounds F_S , and*

$$\begin{aligned} P(B \leq x) &\geq \max_y \left(\int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C \mid P_1^S = y) \right. \\ &\quad \left. + P(X_Q^B \leq x, D_2^B \neq C \mid P_1^S = y) \right) \\ P(B \leq x) &\leq \min_y \left(\int \mathbf{1}(X_{AC}^{B*}(y, z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C \mid P_1^S = y) \right. \\ &\quad \left. + P(X_{AC}^B \leq x, D_2^B \neq C \mid P_1^S = y) \right) \end{aligned}$$

Proof. For the buyer, we have $P(B \leq x) = \max_y P(B \leq x \mid P_1^S = y)$, and

$$\begin{aligned} P(B \leq x \mid P_1^S = y) &= P(B \leq x \mid P_1^S = y, D_2^B = C) P(D_2^B = C \mid P_1^S = y) \\ &\quad + P(B \leq x, D_2^B \neq C \mid P_1^S = y) \\ &= \int P(B \leq x \mid P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) \\ &\quad \times P(D_2^B = C \mid P_1^S = y) + P(B \leq x, D_2^B \neq C \mid P_1^S = y) \\ &\geq \int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C \mid P_1^S = y) \\ &\quad + P(X_Q^B \leq x, D_2^B \neq C \mid P_1^S = y) \end{aligned}$$

Similar arguments yield the upper bound. \square

Positive Correlation Plus Monotonicity. We obtain similar results combining monotonicity and positive correlation. Note that, as with independence, imposing positive correlation *on top of monotonicity* yields no additional information for the seller distribution bounds, but can tighten the buyer distribution bounds.

Theorem 5. *Suppose Assumptions 2 and 5 hold. Then (7) bounds F_S , and*

$$\begin{aligned} F_B(x) &\geq \int \max_{y' \geq y} \left(\int \mathbf{1}(X_Q^{B*}(y', z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C | P_1^S = y') \right. \\ &\quad \left. + P(X_Q^B \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) \\ F_B(x) &\leq \int \min_{y' \leq y} \left(\int \mathbf{1}(X_{AC}^{B*}(y', z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C | P_1^S = y') \right. \\ &\quad \left. + P(X_{AC}^B \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) \end{aligned}$$

Proof. For the buyer, $P(B \leq x) = \int \max_{y' \geq y} P(B \leq x | P_1^S = y') dF_{P_1^S}(y)$ and

$$\begin{aligned} P(B \leq x | P_1^S = y) &= P(B \leq x | P_1^S = y, D_2^B = C) P(D_2^B = C | P_1^S = y) \\ &\quad + P(B \leq x, D_2^B \neq C | P_1^S = y) \\ &= \int P(B \leq x | P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) \\ &\quad \times P(D_2^B = C | P_1^S = y) + P(B \leq x, D_2^B \neq C | P_1^S = y) \\ &\geq \int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C | P_1^S = y) \\ &\quad + P(X_Q^B \leq x, D_2^B \neq C | P_1^S = y) \end{aligned}$$

Analogous arguments yield the upper bound. \square

Independence Plus Stochastic Monotonicity. We now combine the stochastic monotonicity and independence assumptions:

Theorem 6. *Suppose Assumptions 3 and 4 hold. Then*

$$\begin{aligned} F_S(x) &\geq \int \max_{y' \geq y} \left(\max_z m_{AC}^S(x, y', z) P(D_2^B = C | P_1^S = y') \right. \\ &\quad \left. + P(X_{AC}^S \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) \end{aligned}$$

$$F_S(x) \leq \int \min_{y' \leq y} \left(\min_z m_Q^S(x, y', z) P(D_2^B = C | P_1^S = y') \right. \\ \left. + P(X_Q^S \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y)$$

and

$$F_B(x) \geq \max_y \left(\int \max_{z' \geq z} m_Q^B(x, y, z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C | P_1^S = y) \right. \\ \left. + P(X_Q^B \leq x, D_2^B \neq C | P_1^S = y) \right) \\ F_B(x) \leq \min_y \left(\int \max_{z' \geq z} m_{AC}^B(x, y, z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y) P(D_2^B = C | P_1^S = y) \right. \\ \left. + P(X_{AC}^B \leq x, D_2^B \neq C | P_1^S = y) \right)$$

Proof. For the seller, we have $P(S \leq x) = \int \max_{y' \geq y} P(S \leq x | P_1^S = y') dF_{P_1^S}(y)$. Moreover,

$$P(S \leq x | P_1^S = y') = P(S \leq x | P_1^S = y', D_2^B = C) P(D_2^B = C | P_1^S = y') \\ + P(S \leq x, D_2^B \neq C | P_1^S = y') \\ = \max_z P(S \leq x | P_1^S = y', P_2^B = z, D_2^B = C) P(D_2^B = C | P_1^S = y') \\ + P(S \leq x, D_2^B \neq C | P_1^S = y')$$

The lower bound then follows from using $S \leq X_{AC}^S$. Analogous arguments with $S \geq X_Q^S$ yield the upper bound.

For the buyer we have $P(B \leq x) = \max_y P(B \leq x | P_1^S = y)$ and

$$P(B \leq x | P_1^S = y) = P(B \leq x | P_1^S = y, D_2^B = C) P(D_2^B = C | P_1^S = y) \\ + P(B \leq x, D_2^B \neq C | P_1^S = y) \\ = \int P(B \leq x | P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) \\ \times P(D_2^B = C | P_1^S = y) + P(B \leq x, D_2^B \neq C | P_1^S = y) \\ = \int \max_{z' \geq z} P(B \leq x | P_1^S = y, P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y) \\ \times P(D_2^B = C | P_1^S = y) + P(B \leq x, D_2^B \neq C | P_1^S = y)$$

and the lower bound now follows from $B \leq X_Q^B$. Analogous arguments and $X_{AC}^B \leq B$ yield the upper bound. \square

Positive Correlation Plus Stochastic Monotonicity. Here we combine stochastic monotonicity and positive correlation, our weakest combination of assumptions:

Theorem 7. *Suppose Assumptions 3 and 5 hold. Then*

$$\begin{aligned}
F_S(x) &\geq \int \max_{y' \geq y} \left(\int \max_{z' \geq z} m_{AC}^S(x, y', z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y') P(D_2^B = C | P_1^S = y') \right. \\
&\quad \left. + P(X_{AC}^S \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) \\
F_S(x) &\leq \int \min_{y' \leq y} \left(\int \min_{z' \geq z} m_Q^S(x, y', z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y') P(D_2^B = C | P_1^S = y') \right. \\
&\quad \left. + P(X_Q^S \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) dy
\end{aligned}$$

and

$$\begin{aligned}
F_B(x) &\geq \int \max_{y' \geq y} \left(\int \max_{z' \geq z} m_Q^B(x, y', z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y') P(D_2^B = C | P_1^S = y') \right. \\
&\quad \left. + P(X_Q^B \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y) \\
F_B(x) &\leq \int \min_{y' \leq y} \left(\int \min_{z' \geq z} m_{AC}^B(x, y', z') dF_{P_2^B | P_1^S, D_2^B = C}(z|y') P(D_2^B = C | P_1^S = y') \right. \\
&\quad \left. + P(X_{AC}^B \leq x, D_2^B \neq C | P_1^S = y') \right) dF_{P_1^S}(y)
\end{aligned}$$

Proof. We have, for the seller, $P(S \leq x) = \int \max_{y' \geq y} P(S \leq x | P_1^S = y') dF_{P_1^S}(y)$. Moreover,

$$\begin{aligned}
P(S \leq x | P_1^S = y') &= P(S \leq x | P_1^S = y', D_2^B = C) P(D_2^B = C | P_1^S = y') \\
&\quad + P(S \leq x, D_2^B \neq C | P_1^S = y') \\
&= \int P(S \leq x | P_1^S = y', P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y') \\
&\quad \times P(D_2^B = C | P_1^S = y') + P(S \leq x, D_2^B \neq C | P_1^S = y') \\
&= \int \max_{z' \geq z} P(S \leq x | P_1^S = y', P_2^B = z, D_2^B = C) dF_{P_2^B | P_1^S, D_2^B = C}(z|y') \\
&\quad \times P(D_2^B = C | P_1^S = y') + P(S \leq x, D_2^B \neq C | P_1^S = y')
\end{aligned}$$

The lower bound then follows from using $S \leq X_{AC}^S$. Analogous arguments with $S \geq X_Q^S$ yield the upper bound.

Similarly, for the buyer we have $P(B \leq x) = \int \max_{y' \geq y} P(B \leq x | P_1^S = y') dF_{P_1^S}(y)$ and

$$P(B \leq x | P_1^S = y') = P(B \leq x | P_1^S = y', D_2^B = C) P(D_2^B = C | P_1^S = y')$$

$$\begin{aligned}
& + P(B \leq x, D_2^B \neq C \mid P_1^S = y') \\
& = \int P(B \leq x \mid P_1^S = y', P_2^B = z, D_2^B = C) dF_{P_2^B \mid P_1^S, D_2^B = C}(z \mid y') \\
& \quad \times P(D_2^B = C \mid P_1^S = y') + P(B \leq x, D_2^B \neq C \mid P_1^S = y') \\
& = \int \max_{z' \geq z} P(B \leq x \mid P_1^S = y', P_2^B = z, D_2^B = C) dF_{P_2^B \mid P_1^S, D_2^B = C}(z \mid y') \\
& \quad \times P(D_2^B = C \mid P_1^S = y') + P(B \leq x, D_2^B \neq C \mid P_1^S = y')
\end{aligned}$$

The lower bound now follows from $B \leq X_Q^B$. Analogous arguments and $X_{AC}^B \leq B$ yield the upper bound. \square

B.2. Estimation of Bounds Combining Assumptions. Here we focus only on lower bound estimators and the case where $P(D_2^B = C) = 1$. The upper bound estimators are analogous. We first consider the bounds in Theorem 4. For the seller, these bounds do not differ from the bounds that only assume monotonicity. For the buyer lower bound, we first write the lower bound as

$$\begin{aligned}
\max_y \left(\int \mathbf{1}(X_Q^{B*}(y, z) \leq x) dF_{P_2^B \mid P_1^S}(z \mid y) \right) & = \max_y \left(P(X_Q^{B*}(y, P_2^B) \leq x \mid P_1^S = y) \right) \\
& = \max_y P(X_Q^{B*}(P_1^S, P_2^B) \leq x \mid P_1^S = y)
\end{aligned}$$

To estimate this bound, we estimate $P(X_Q^{B*}(P_1^S, P_2^B) \leq x \mid P_1^S = y)$ using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel and bandwidth $n^{-1/5}$ for each value of x . We restrict all estimators to be in the interval $[0, 1]$ and rearrange them such that the estimated functions are monotone in x .

We now consider the bounds in Theorem 5. Again, these bounds only differ from the monotonicity bounds for the buyer. For the buyer, we can use arguments as in the previous paragraph to write the lower bound as

$$\int \max_{y' \geq y} P(X_Q^{B*}(y', P_2^B) \leq x \mid P_1^S = y') dF_{P_1^S}(y) = \int \max_{y' \geq y} P(X_Q^{B*}(P_1^S, P_2^B) \leq x \mid P_1^S = y') dF_{P_1^S}(y)$$

Using the estimator $\hat{P}(\hat{X}_Q^{B*}(P_1^S, P_2^B) \leq x \mid P_1^S = y')$ from above, we estimate the lower bound by

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \hat{P}(\hat{X}_Q^{B*}(P_1^S, P_2^B) \leq x \mid P_1^S = y')$$

For the bounds in Theorem 6, we estimate the seller lower bound using the following sample

analog:

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \left(\max_z \hat{m}_{AC}^S(x, y', z) \right)$$

For the buyer, define $g_Q^B(x, y, z) = \max_{z' \geq z} m_Q^B(x, y, z')$. Then we can write the lower bound as

$$\max_y \left(E[g_Q^B(x, P_1^S, P_2^B) \mid P_1^S = y] \right)$$

For each x , we estimate $E[g_Q^B(x, P_1^S, P_2^B) \mid P_1^S = y]$ using the Nadaraya-Watson kernel estimator with an Epanechnikov kernel function and bandwidth $n^{-1/5}$. Let $\hat{E}[g_Q^B(x, P_1^S, P_2^B) \mid P_1^S = y]$ denote the estimator. Our estimator is

$$\max_{Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)} \left(\hat{E}[g_Q^B(x, P_1^S, P_2^B) \mid P_1^S = y] \right)$$

where $g_Q^B(x, y, P_2^B) = \max_{z' \in \{z: z \geq P_{2,i}^B, Q_{0.05}(P_{2,i}^B) \leq z \leq Q_{0.95}(P_{2,i}^B)\} \cup \{P_{2,i}^B\}} m_Q^B(x, y, z')$.

For the bounds from Theorem 7, we estimate the seller value lower bound by

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \left(\hat{E}[g_{AC}^S(x, P_1^S, P_2^B) \mid P_1^S = y'] \right)$$

and the buyer value lower bound by

$$\frac{1}{n} \sum_{i=1}^n \max_{y' \in \{y: y \geq P_{1,i}^S, Q_{0.05}(P_{1,i}^S) \leq y \leq Q_{0.95}(P_{1,i}^S)\} \cup \{P_{1,i}^S\}} \left(\hat{E}[g_Q^B(x, P_1^S, P_2^B) \mid P_1^S = y'] \right)$$

C Monte Carlo Simulations

This section presents a Monte Carlo study of the buyer and seller marginal distribution bounds. There is naturally a great deal of flexibility in how to simulate two-sided bargaining; here we simply simulate outcome data consistent with our assumptions. We do not simulate actual equilibrium play of a two-sided bargaining game, as none of the equilibria that have been analyzed in the existing literature (Perry 1986; Grossman and Perry 1986; Cramton 1992) result in multiple offers by a given party that vary with the party's value.

C.1. Algorithm for Simulating Bargaining Data. Here we describe our algorithm for simulating bargaining data. The primary parameters we will vary in in this simulation exercise are α_b

and α_s , which we refer to as *shade factors*; the probability that a buyers and sellers accept/decline; and the means of the buyer and seller value distributions. These shade factors allow us to vary how aggressive the agents' offers are: a buyer with value b and shade factor α_b makes the same offers as a buyer with value $b + \alpha_b$ and shade factor 0, and in this sense the shade factors set a minimum level of offer shading. The probability that a buyer or seller accepts or quits (instead of making a counteroffer) allows us to investigate how the frequency of counteroffering affects the tightness of the bounds. Varying the means of the buyer and seller distributions allows us to vary the amount of potential surplus in the bargaining game. The data generating process (DGP) is as follows:

Algorithm (Bargaining Simulation Data Generating Process).

***Initialize:** Draw buyer values B from CDF F_B and seller value B from F_S . Set buyer and seller shade factors α_B and α_S , and set a cap T_{max} on the number of rounds. Set functions $p_{BQ}(k)$, $p_{BA}(k)$, $p_{SQ}(k)$, and $p_{SA}(k)$ specifying the probabilities, respectively, in round k , of the buyer quitting, buyer accepting, seller quitting, or seller accepting.*

***Round 1:** Seller offers $P_1^S = g_1(S, \alpha_s, U_1)$ where $U_1 \sim U[0, 1]$ and the function g_1 is weakly increasing in all of its arguments. We vary g_1 in our illustrations.*

***Round 2:** Buyer offers $P_2^B = U_2(B - \alpha_b)$ if $P_1^S > B$ and $P_2^B = U_2 \min\{P_1^S, B - \alpha_b\}$ if $P_1^S \leq B$, where $U_2 \in (0, 1)$ is random or fixed depending on the specific setup.*

***Round $3 \leq k < T_{max}$, k odd:** Seller responds to buyer's last offer. Two cases:*

***Case 1** $P_{k-1}^B < S$: Seller quits with probability $p_{SQ}(k)$, or else makes a counteroffer $P_k^S = U_3 P_{k-2}^S + (1 - U_3)(S + \alpha_S)$, where $U_3 \in (0, 1)$ and its distribution depends on the specific setup.*

***Case 2** $P_{k-1}^B \geq S$: Seller accepts with fixed probability $p_{SA}(k)$, or else makes a counteroffer $P_k^S = U_3 P_{k-2}^S + (1 - U_3) \max(P_{k-1}^B, S + \alpha_S)$.*

***Round $4 \leq k < T_{max}$, k even:** Buyer responds to seller's last offer P_{k-1}^S . Two cases:*

***Case 1** $P_{k-1}^S > B$: Buyer quits with probability $p_{BQ}(k)$, or else makes a counter offer $P_k^B = U_4 P_{k-2}^B + (1 - \lambda)(B - \alpha_B)$, where $U_4 \in (0, 1)$ and its distribution depends on the specific setup.*

***Case 2** $P_{k-1}^S \leq B$: Buyer accepts with probability $p_{BA}(k)$, or else makes a counter offer $P_k^B = U_4 P_{k-2}^B + (1 - U_4) \min(P_{k-1}^S, b - \alpha_B)$.*

Round T_{max} : Terminate with no trade occurring.

C.2. Results of Monte Carlo Exercise. Figure 6 illustrates several of our bounds estimated using this simulated data. For each panel, we simulate 50 replications of the DGP and then report estimated bounds averaged across these replications. In each example we set $n = 250$ and $T_{max} = 8$. We draw the values from a Beta distribution, which has support on $[0, 1]$. The Beta distribution has two parameters, α and β . We set $\alpha = 2$ and set β depending on which mean value we want to achieve. We then add the maximum shading factor to ensure that bids are always non-negative. We vary the parameters of the data generating process in each panel in order to illustrate what features will lead to bounds that are loose (panels on the left) or tight (panels on the right). We focus on three sets of bounds — the seller unconditional bounds, seller monotonicity bounds, and buyer independence bounds — to conserve space and because the intuition gained by these three cases extends to the other bounds in the paper. In each panel, lower bounds are shown with solid lines, upper bounds with dashed lines, and the true CDF is shown with a dot-dash line.

We show the seller unconditional bounds in panels A and B. For the seller unconditional bounds to be relatively tight, it must be the case that sellers quit at prices close to their values and also counter at prices close to their values. As an example, consider a setting where buyer and seller values are very highly correlated and have a similar mean. Suppose the typical play of the game is that the seller offers a price a little above her value, the buyer counters at a price a little below the seller’s value (and also below the buyer’s value, naturally), and the seller then quits. This sequence of play is consistent with the weak revealed preference assumptions that the unconditional bounds are built on (Assumption 1) and it yields the tight bounds on seller values illustrated in panel B.²⁹ We can also easily generate wide unconditional bounds. For example, consider a case where the seller typically makes offers far above her value and rarely quits. Such bounds are illustrated in panel A.³⁰ Here, the correlation structure between buyer and seller values plays no role.

We illustrate the seller monotonicity bounds in panels C and D. The monotonicity bounds will improve upon the unconditional bounds when there is some probability that sellers who start with relatively low first offers end the game at relatively high final accept/counter or quit prices. This can occur due to randomness in the value of the buyer to whom the seller is matched and



²⁹The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0.999. We set $p_{BQ}(k) = p_{BA}(k) = p_{SA}(k) = 0$, $p_{SQ}(k) = 0.95$, $g_1(S, \alpha_s, U_1) = 1.1(S + \alpha_s)$, and $U_2 = 0.9$, $U_3 \sim U[0, 0.5]$, $U_4 \sim U[0, 0.5]$.

³⁰The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0. We set $p_{BQ}(k) = p_{BA}(k) = 1$, $p_{SA}(k) = 0$, $p_{SQ}(k) = 0.25$, $g_1(S, \alpha_s, U_1) = 1.5(S + \alpha_s) + 0.5$, and $U_2 = 0.9$, $U_3 \sim U[0, 0.5]$, $U_4 \sim U[0, 0.5]$.

Figure 6: Simulation Results



(A) Unconditional Bounds, Seller (Wide)

(B) Unconditional Bounds, Seller (Tight)



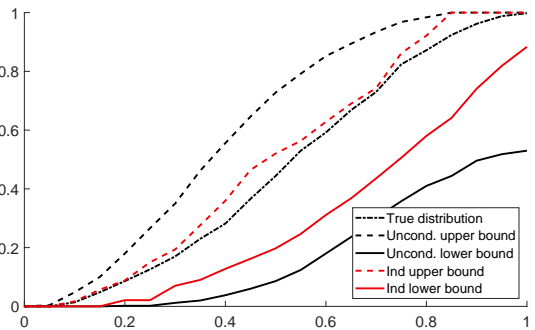

(C) Monotonicity Bounds, Seller (Wide)

(D) Monotonicity Bounds, Seller (Tight)



(E) Independence Bounds, Buyer (Wide)

(F) Independence Bounds, Buyer (Tight)



Notes: Table shows bounds estimated from simulated data under cases where bounds are wide (on left) vs. narrow (on right). Panels A and B show unconditional seller bounds. Panels C and D show seller monotonicity bounds. Panels E and F show buyer independence bounds. Lower bounds are shown with solid lines, upper bounds with dashed lines, and true CDF with a dot-dash line. Panels C–F also show unconditional bounds for comparison.

due to features of bargaining at later rounds of the game. We illustrate such a case in panel D.³¹ If, however, the final accept/counter and quit prices of a seller are, like P_1^S , deterministically monotononic in the seller's value, the monotonicity assumption will do nothing to improve upon the unconditional bounds (because $X_{AC}^{S*} = X_{AC}^S$ and $X_Q^{S*} = X_Q^S$ in that case, and the unconditional bounds will equal the monotonicity bounds). We illustrate this situation in panel C, where the monotonicity bounds are equally as wide as the unconditional bounds.³²

Finally, we illustrate the buyer independence bounds in panels E and F. Recall that these bounds are obtained by combining $P(B \leq x | P_1^S = y) = P(B \leq x)$ (buyer independence) with weak rationality on the part of the buyer ($X_{AC}^B \leq B$ for the buyer upper bound). The buyer independence assumption will therefore yield no improvement over the buyer unconditional upper bounds if X_{AC}^B and X_Q^B are, like B , independent of P_1^S . This case is illustrated in panel E.³³ It is easy to generate a case in which the maximum accept/counter price of the buyer *does* depend on P_1^S , and this yields a much tighter upper bound. To do so, we generate data such that B is independent of P_1^S , but X_{AC}^B and X_Q^B are not because bids in later rounds directly depend on P_1^S .³⁴

D Related Extensive-Form Models

In this section we consider three extensive-form bargaining models: Cramton (1992), Perry (1986), and Grossman and Perry (1986). In each case, when we discuss unobserved heterogeneity, we differ slightly from the notation in the body of the paper in that we write a seller's value with additively separable unobserved heterogeneity included as $\tilde{S} = S + W$ (whereas in the body of the body of the paper we instead write $S = \tilde{S} + W$). Similarly, for the multiplicative case, we write $\tilde{S} = SW$. We apply this notation to the buyer's value and to buyer and seller offers as well. We adopt this change in notation here so that variables without ($\tilde{\cdot}$) always represent those absent unobserved heterogeneity.

³¹The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0.999. We set $p_{BQ}(k) = p_{BA}(k) = p_{SA}(k) = 0$, $p_{SQ}(k) = 0.5$, $g_1(S, \alpha_s, U_1) = 1.5(S + \alpha_s)$, and $U_2 = 0.5$, $U_2 \sim U[0, 0.5]$, $U_3 \sim U[0, 0.5]$.

³²The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0.999. We set $p_{BQ}(k) = p_{SQ}(k) = 1$, $p_{BA}(k) = p_{SA}(k) = 0$, $g_1(S, \alpha_s, U_1) = 1.5(S + \alpha_s)$, and $U_2 = 0.5$, $U_3 \sim U[0, 0.5]$, $U_4 \sim U[0, 0.5]$.

³³The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0. We set $p_{BQ}(k) = 0.95$, $p_{SQ}(k) = p_{BA}(k) = p_{SA}(k) = 0$, $g_1(S, \alpha_s, U_1) = U_1$ with $U_1 \sim [1, 1.5]$, and $U_2 \sim U[0.75, 1]$, and $U_3 = U_4 = U_1$.

³⁴The specification for this case is as follows: both the buyer and the seller have mean values of 0.5 and their correlation is 0. We set $p_{BQ}(k) = 0.95$, $p_{SQ}(k) = p_{BA}(k) = p_{SA}(k) = 0$, $g_1(S, \alpha_s, U_1) = U_1(S + \alpha_s)$ with $U_1 \sim [1, 1.5]$, and $U_2 \sim U[0.75, 1]$, and $U_3 = U_4 = \max\{1 - P_1^S, 0\}$.

D.1. Cramton (1992). This model studies a setting similar to ours, where a seller and buyer with independent private values engage in bargaining. One possible outcome in the Cramton (1992) equilibrium is for the seller to make the first offer, P_1^S , which completely reveals the seller's value S . The buyer then either accepts, quits, or makes a counteroffer P_2^B that completely reveals her value B .³⁵ These first two offers are

$$P_1^S = \frac{\delta S + \gamma(S)}{1 + \delta} \quad (15)$$

$$P_2^B = \frac{\delta B + S}{1 + \delta} \quad (16)$$

where δ is a discount factor and the object $\gamma(S)$ is the buyer type who is indifferent between accepting and rejecting the seller's offer of P_1^S given that the seller has revealed her type to be S . For our arguments here, we assume $\gamma(s)$ is differentiable with $\gamma'(s) \in (-\delta, 0)$. It is possible to show that $\gamma'(s) \in (-\delta, 0)$ is satisfied with plenty of slack in the uniform distribution case, which we state as the following lemma:³⁶

Lemma 1. *In the Cramton (1992) model, if buyer and seller values are uniformly distributed, the function $\gamma(\cdot)$ satisfies, with slack, $\gamma'(\cdot) \in (-\delta, 0)$.*

Proof. In the equilibrium studied in Cramton (1992), the function $\gamma(\cdot)$ is quite complex, and depends on the distribution of buyer values, F_B , with density f_B . Cramton (1992) shows that $\gamma(s)$ is the solution to the following, when the seller's type is s and the seller believes the buyer's value is bounded above by some value \bar{b}

$$F_B(\bar{b}) - F_B(\gamma) - (1 - \delta^2)(\gamma - s)f_B(\gamma) = \int_s^\gamma \delta^3 \left(\frac{b - s}{\gamma - s} \right)^{1+\delta} dF_B(b) \quad (17)$$

In the uniform case, $\gamma(s)$ has a closed-form solution given by $\gamma(s) = \alpha - (2\alpha - 1)s$, where α is defined by $1 - 2\alpha = \frac{-\delta}{2+\delta-\delta^2}$. Note that $\alpha \in (\frac{1}{2}, \frac{3}{4})$ for $\delta \in (0, 1]$, and thus $\gamma'(s) \in [-.5, 0)$ for $\delta \in (0, 1]$, and thus $\gamma'(s) < 0$ is satisfied with slack in the uniform case.

³⁵If the buyer chooses to make a counteroffer, P^B , the buyer exploits this knowledge of the seller's type and makes an offer that corresponds to the Rubinstein (1982) equilibrium offer for the case where the buyer and seller know each others' values. Note that, in the Cramton (1992) equilibrium, the timing of these offers is also important in revealing an agent's value, but focusing on the level of the offers is sufficient for our point here.

³⁶It is also possible to derive sufficient conditions for these properties to hold outside of the uniform case; these would be similar to the assumption referred to as “(F δ)” in Cramton (1992). These conditions are quite cumbersome. Like Cramton, therefore, we instead show that these are satisfied for the case where buyer values are uniformly distributed on $[0, 1]$, while still having plenty of slack, and thus they do not appear to be overly restrictive.

Note that $\gamma'(s) = 1 - 2\alpha$. Setting $\gamma'(s)$ to be weakly greater than $-\delta$ yields

$$\frac{-\delta}{2 + \delta - \delta^2} \geq -\delta \iff \delta^2 \leq 2,$$

and thus $\gamma'(s) > -\delta$ is satisfied with slack in the uniform case. \square

An immediate result of this property is that the equilibrium offers (15) and (16) satisfy Assumption 2 (strictly, in fact): P_1^S is strictly monotone in S because $\gamma'(s) > -\delta$, and hence P_2^B is also strictly monotone in B conditional on P_1^S .

Now consider a modification of the Cramton setting in which a buyer and seller play the equilibrium of Cramton (1992), but in a given realization of the game buyer and seller values are both shifted additively by a common amount, W , that is independent of B and S . Specifically, a buyer's value is given by $B + W$ and a seller's by $S + W$, where $W = w$ is known to both agents but not to the econometrician. Cramton's model assumes, without loss of generality, that buyer values are distributed on $[0, 1]$. In our modification, we instead have values distributed on $[w, 1 + w]$. In this environment, the equilibrium offers simply shift additively by W as well, becoming $P_1^S + W$ and $P_2^B + W$, as we demonstrate in the following lemma:

Lemma 2. *Suppose seller and buyer values in the Cramton (1992) setting are given by $S + W$ and $B + W$. If, when $W = 0$, the first two offers are given by $P_1^S = p_1^S$ and $P_2^B = p_2^B$, then, when $W = w$, these offers are given by $p_1^S + w$ and $p_2^B + w$.*

Proof. We first prove the following claim: The function $\gamma(\cdot)$ satisfies additive separability. Let $\tilde{\gamma}(s, w)$ represent the value of γ in a game in which $W = w$; thus $\gamma(s) = \tilde{\gamma}(s, 0)$. We will show that $\tilde{\gamma}(s, w) = \gamma(s) + w$. To see this, let $\tilde{b} = \bar{b} + w$, and $\tilde{s} = s + w$, and let $\tilde{F}_{\tilde{B}}$ and $\tilde{f}_{\tilde{B}}$ be the distribution and density of \tilde{B} .

The condition defining $\tilde{\gamma}$ is given by modifying (17) as follows:

$$\tilde{F}_{\tilde{B}}(\tilde{b}) - \tilde{F}_{\tilde{B}}(\tilde{\gamma}) - (1 - \delta^2)(\tilde{\gamma} - \tilde{s})\tilde{f}_{\tilde{B}}(\tilde{\gamma}) = \int_{\tilde{s}}^{\tilde{\gamma}} \delta^3 \left(\frac{x - \tilde{s}}{\tilde{\gamma} - \tilde{s}} \right)^{1+\delta} d\tilde{F}_{\tilde{B}}(x)$$

Note that, for any number x , $\tilde{F}_{\tilde{B}}(\tilde{x}) = F_B(\tilde{x} - w)$ and $\tilde{f}_{\tilde{B}}(\tilde{x}) = f_B(\tilde{x} - w)$. We now apply a change of variables from x to $y = x - w$ in the integral, yielding

$$\int_s^{\tilde{\gamma}-w} \delta^3 \left(\frac{y + w - \tilde{s}}{\tilde{\gamma} - \tilde{s}} \right)^{1+\delta} dF_B(y)$$

Combining these results yields

$$F_B(b) - F_B(\tilde{\gamma} - w) - (1 - \delta^2)(\tilde{\gamma} - \tilde{s})f_B(\tilde{\gamma} - w) = \int_s^{\tilde{\gamma}-w} \delta^3 \left(\frac{y-s}{\tilde{\gamma}-\tilde{s}} \right)^{1+\delta} dF_B(y) \quad (18)$$

Comparing (17) to (18) demonstrates that, if γ is the solution to the former then $\gamma + w$ is the solution to the latter, completing the proof of the claim.

Now consider the equilibrium of this game conditional on a realization of W . Offers will be given by

$$\begin{aligned} \tilde{p}_1^S &= \frac{\delta \tilde{s} + \tilde{\gamma}(s, w)}{1 + \delta} = \frac{\delta s + \gamma(s)}{1 + \delta} + w \\ \tilde{p}_2^B &= \frac{\delta \tilde{b} + \tilde{s}}{1 + \delta} = \frac{\delta b + s}{1 + \delta} + w \end{aligned}$$

Thus, additive separability of the offers is satisfied. \square

We now demonstrate that the presence of such unobserved heterogeneity can lead to a violation of the monotonicity assumption, even though stochastic monotonicity is still satisfied. We show that monotonicity of the seller's first offer \tilde{P}^S in the seller's value \tilde{S} is violated in this setting, and we prove an analogous result for the buyer. Note that here we are considering what this setting would look like to us as econometricians, where we see observations of different instances of the game and where realizations of W may vary across these observations.

Lemma 3. *The Cramton (1992) equilibrium offers in a game with additive unobserved heterogeneity can violate Assumption 2, but Assumption 3 is still satisfied.*

Proof. Suppose s increases by 1 and w decreases by $\eta < 1$ (so \tilde{s} increases overall). Because $\gamma(s)' \in (-\delta, 0)$, this means that the change in \tilde{p}_1^S due to the change in s is at most an increase of $\frac{\delta}{1+\delta}$, and the change in \tilde{p}_1^S due to the change in w is a decrease of η . For any $\eta \in \left(\frac{\delta}{1+\delta}, 1\right)$, \tilde{p}_1^S decreases even though \tilde{s} increases, violating seller monotonicity.

For buyer monotonicity, suppose that s increases by η_s and w decreases by η_w such that \tilde{p}_1^S does not change. It then holds that $0 < \eta_w < \eta_s$. Next suppose that b increases by $\eta_b \in (\eta_w < \eta_s)$. Then $\tilde{b} = b + w$ decreases, but since

$$\frac{\delta(b + \eta_b) + (s + \eta_s)}{1 + \delta} + w - \eta_w > \frac{\delta b + s}{1 + \delta} + w$$

\tilde{p}_2^B increases.

To see that stochastic monotonicity is satisfied for the seller, let $g(s) \equiv \frac{s-\gamma(s)}{1+\delta}$, which is strictly increasing assuming $\gamma(\cdot)$ is strictly decreasing. Then we have

$$\begin{aligned}
P(\tilde{S} \leq x | \tilde{P}_1^S = y) &= P(S \leq x - W | P_1^S + W = y) \\
&= \int P(S \leq x - y + f(S) | w = y - f(S), W = w) f_{W|w=y-f(S)}(w) dw \\
&= \int P\left(S \leq x - y + \frac{\delta S + \gamma(S)}{1 + \delta} \middle| w = y - f(S), W = w\right) f_W(w) dw \\
&= \int P(g(S) \leq x - y | f(S) = y - w, W = w) f_W(w) dw \\
&= \int P(g(S) \leq x - y | f(S) = y - w) f_W(w) dw \\
&= \int P(g(f^{-1}(y - w)) \leq x - y) f_W(w) dw \\
&= \int \mathbf{1}(g(f^{-1}(y - w)) \leq x - y) f_W(w) dw
\end{aligned}$$

Since $g(f^{-1}(\cdot))$ is a strictly increasing function, $P(\tilde{S} \leq x | \tilde{P}_1^S = y)$ is strictly decreasing in y . In the third and fifth line, we use that W and S are independent.

Using similar arguments, we can also show that the stochastic monotonicity condition of the buyer holds. To do so, write

$$\begin{aligned}
P(\tilde{B} \leq x | \tilde{P}_1^S = y, \tilde{P}_2^B = z) &= P\left(B \leq x - W \middle| f(S) + W = y, \frac{\delta B + S}{1 + \delta} + W = z\right) \\
&= \int P\left(B \leq x - w \middle| f(S) + w = y, \frac{\delta B + S}{1 + \delta} + w = z\right) f_W(w) dw \\
&= \int P\left(B \leq x - w \middle| S = f^{-1}(y - w), \frac{\delta B + f^{-1}(y - w)}{1 + \delta} = z - w\right) f_W(w) dw \\
&= \int P\left(B \leq x - w \middle| S = f^{-1}(y - w), B = \frac{1}{\delta}((1 + \delta)(z - w) - f^{-1}(y - w))\right) f_W(w) dw \\
&= \int \mathbf{1}\left(\frac{1}{\delta}((1 + \delta)(z - w) - f^{-1}(y - w)) \leq x - w\right) f_W(w) dw
\end{aligned}$$

which is decreasing in z . In the third line, we used that W is independent of (S, B) . \square

The Cramton model assumes independence of B and S , and this immediately yields the result that the independence assumption for the seller is satisfied in his model: S is independent of P_2^B

conditional on P_1^S because P_1^S completely reveals S to the buyer, and hence, conditional on P_1^S , there is no variation left in S . However, even maintaining the independence of the components B and S , if additive unobserved heterogeneity is introduced into the game, then $B + W$ will be correlated with $P_1^S + W$, violating our buyer independence assumption. The proof of Lemma 4, focusing on the uniform distribution case, demonstrates that seller independence can also be violated without violating positive correlation.

Lemma 4. *The Cramton (1992) equilibrium offers in a game with additive unobserved heterogeneity can violate Assumption 4.i for the seller and Assumption 4.ii for the buyer.*

Proof. In the Cramton model *with unobserved heterogeneity*, clearly \tilde{B} is correlated with \tilde{P}_1^S through W , so buyer independence (Assumption 4.ii) is violated. For seller independence (Assumption 4.i), note from Lemma 3 that \tilde{P}_2^B can be written as follows, where \tilde{P}_1^S is fixed at y :

$$\tilde{P}_2^B = \frac{\delta B + f^{-1}(y - W)}{1 + \delta} + W \quad (19)$$

Now consider a change in \tilde{S} . Holding \tilde{P}_1^S fixed at y , this change in \tilde{S} must also correspond to a change in W (or else \tilde{P}_1^S could not remain constant).

This change in W will necessarily affect \tilde{P}_2^B unless the terms in (19) depending on W completely offset on another; that is, unless $\frac{d}{dw} \left(\frac{f^{-1}(y-w)}{1+\delta} + w \right) = 0$. To see that this is not the case, note that $\gamma' \in (-\delta, 0)$ implies $f' \in (0, \frac{\delta}{1+\delta})$, and, by the inverse function theorem, $f^{-1'} \in (\frac{1+\delta}{\delta}, \infty)$. This implies that

$$\frac{d}{dw} \frac{f^{-1}(y-w)}{1+\delta} + w \in (-\infty, -1/\delta + 1).$$

For any $\delta < 1$, this derivative is non-zero, and thus variation in W also leads to variation in \tilde{P}_2^B , violating seller independence. \square

D.2. Perry (1986). The Perry model has no discounting, and instead agents face a per-offer additive cost of bargaining, c_S for the seller and c_B for the buyer. Both the buyer and seller have private values, and they alternate offers. The equilibrium that Perry focuses on has the property that only one player makes an offer and the other player accepts or rejects (and never makes a counteroffer). One outcome in the Perry equilibrium is for the seller to make the first offer, P_1^S , with this offer given by

$$P_1^S = \frac{1 - F_B(P_1^S)}{f_B(P_1^S)} + S \quad (20)$$

where f_B is the density of buyer values. In this equilibrium, the seller's first offer, P_1^S , clearly satisfies monotonicity (Assumption 2.i), and hence also satisfies the weaker condition of stochastic monotonicity (Assumption 3.i).

In an additively separable unobserved heterogeneity version of this model, the seller's offer will also be additively separable in the unobserved heterogeneity. Specifically,

$$\tilde{P}_1^S = \frac{1 - F_{\tilde{B}}(\tilde{P}_1^S)}{f_{\tilde{B}}(\tilde{P}_1^S)} + \tilde{S} = \frac{1 - F_B(\tilde{P}_1^S - W)}{f_B(\tilde{P}_1^S - W)} + S + W = P_1^S + W$$

Thus, $\tilde{P}_1^S = P_1^S + W$.

In this additively separable version of the model, seller monotonicity (Assumption 2.i) can be violated. To show this, we re-write (20) as $\phi(p_1^S) = s$, where $\phi(p_1^S) \equiv p_1^S - \frac{1 - F_B(p_1^S)}{f_B(p_1^S)}$ is the buyer's *virtual value* function. Implicit differentiation of $\phi(p_1^S) = s$ with respect to s yields $\frac{dp_1^S}{ds} = \frac{1}{\phi'(p_1^S)}$. Consider now a case where s increases by 1 and w decreases by $\eta < 1$, and hence \tilde{s} increases overall. The object \tilde{p}_1^S will increase by $\frac{1}{\phi'(p_1^S)} - \eta$. For any distribution F_B with $\phi'(\cdot) > 1$, there exists an $\eta < 1$ such that p_1^S will increase by less than η when s increases by 1, and, in such a case, \tilde{p}_1^S will decrease overall. The uniform distribution on $[0,1]$ is one such example, where this condition is satisfied with slack, with $\phi'(\cdot) = 2$.

Consider now a case in which agents play the equilibrium of Perry (1986), but in a given realization of the game buyer and seller values are both scaled *multiplicatively* (rather than shifted additively) by W (again, that is common knowledge to both agents). Thus, $\tilde{S} = WS$ and $\tilde{B} = WB$. In this case, it can be shown that

$$\tilde{P}_1^S = WP_1^S = W \frac{1 - F_B(\tilde{P}_1^S/W)}{f_B(\tilde{P}_1^S/W)} + WS \quad (21)$$

where P_1^S is the offer the seller would make if the realization of W were 1. The presence of this multiplicative heterogeneity can lead to violations of weak monotonicity across instances of the game. Consider, for simplicity, the case where $B \sim U[0,1]$. In this case, (21) simplifies to $2p_1^S = w + ws$. Suppose s increases by 1 and w decreases, scaling down by some factor $\eta \in (\frac{1}{2}, 1)$; overall, \tilde{s} increases by a factor of $2\eta > 1$. However, (21) then implies that $\tilde{p}_1^S = \eta p_1^S$, and thus \tilde{p}_1^S decreases.

Regarding independence, the independence of the buyer's value from the seller's first offer (Assumption 4.ii) is clearly satisfied in this model absent unobserved heterogeneity. However,

once unobserved heterogeneity is included, the seller's first offer \tilde{P}_1^S and buyer's value \tilde{B} will clearly be correlated through W in both the additively or multiplicatively separable unobserved heterogeneity models. The Perry model cannot serve for studying monotonicity of the buyer's first offer or independence of the buyer's offer from the seller's value because only one offer occurs in equilibrium.

D.3. Grossman and Perry (1986). In this model, the buyer has a private value but the seller does not (she has a value of 0). The parties alternate offers, with the seller moving first. Grossman and Perry first focus on a two-period game, and then extend their results to an infinite-horizon game, but the equilibrium they focus on is the same in both cases: the seller makes an offer to the buyer and the buyer either accepts or makes a counteroffer at a price that the seller immediately accepts (similar to the Cramton 1992 equilibrium). All sellers make the same first offer (because sellers have no private value) and all buyers who make a counteroffer make the same counteroffer. Because these offers do not depend on agents' values, our buyer monotonicity and independence assumptions are trivially satisfied by this model. Seller monotonicity and independence are also trivially satisfied because the seller has no private value.

If the analysis were to be modified to include unobserved heterogeneity that varies across instances of the game, seller and buyer monotonicity, as well as seller independence, would still be satisfied because any variation in offers across games would necessarily be driven entirely by variation in unobserved heterogeneity. For example, conditioning on the first offer \tilde{P}_1^S would effectively condition on the realization of unobserved heterogeneity W , and, conditional on W , the buyer's offer \tilde{P}_2^B is trivially weakly monotonic in the buyer's value B (because it is constant). Buyer independence would be violated, because \tilde{B} would be correlated with the seller's first offer through W .