

NBER WORKING PAPER SERIES

THE EFFICIENCY OF REAL-WORLD BARGAINING:
EVIDENCE FROM WHOLESALE USED-AUTO AUCTIONS

Bradley Larsen

Working Paper 20431

<http://www.nber.org/papers/w20431>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2014, Revised September 2019

I thank Panle Jia Barwick, Glenn Ellison, and Stephen Ryan for invaluable help and advice. I would also like to thank Isaiah Andrews, Gabe Carroll, Mingli Chen, Victor Chernozhukov, Denis Chetverikov, Dominic Coey, Joachim Freyberger, Ken Hendricks, Kyoo-il Kim, Steven Lang, Ariel Pakes, Christopher Palmer, Parag Pathak, Brennan Platt, Zhaonan Qu, Dan Quint, Mark Satterthwaite, Paulo Somaini, Alan Sorensen, Evan Storms, Jean Tirole, Juuso Toikka, and Anthony Zhang for helpful suggestions. This paper also benefited from conversations with many other faculty and students at various institutions. I thank several anonymous auction houses and their employees for providing data and institutional details. I acknowledge support from the National Science Foundation under grants GRFP-0645960 and SES-1530632. Earlier versions of this paper were circulated under the title, "The Efficiency of Dynamic, Post-Auction Bargaining: Evidence from Wholesale Used-Auto Auctions." The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Bradley Larsen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions
Bradley Larsen
NBER Working Paper No. 20431
August 2014, Revised September 2019
JEL No. C57,C78,D44,D47,D82,L1

ABSTRACT

This study empirically quantifies the efficiency of a real-world bargaining game with two-sided incomplete information. Myerson and Satterthwaite (1983) and Williams (1987) derived the theoretical ex-ante efficient frontier for bilateral trade under two-sided uncertainty and demonstrated that it falls short of ex-post efficiency, but little is known about how well bargaining performs in practice. Using about 265,000 sequences of a game of alternating-offer bargaining following an ascending auction in the wholesale used-car industry, this study estimates (or bounds) distributions of buyer and seller valuations and evaluates where realized bargaining outcomes lie relative to efficient outcomes. Results demonstrate that the ex-ante and ex-post efficient outcomes are close to one another, but that the real bargaining falls short of both, suggesting that the bargaining is indeed inefficient but that this inefficiency is not solely due to the information constraints highlighted in Myerson and Satterthwaite (1983). Quantitatively, findings indicate that 17–24% of negotiating pairs fail to trade even though gains from trade exist, leading an efficiency loss of 12–23% of the available gains from trade.

Bradley Larsen
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
bjlarsen@stanford.edu

Whether haggling in an open-street market, deciding upon prices between an upstream supplier and downstream producer, or negotiating a corporate takeover deal, bargaining between a buyer and seller is one of the oldest and most common ways of transacting. When both parties have incomplete information, it is known that equilibrium outcomes are difficult to characterize.¹ Myerson and Satterthwaite (1983) demonstrated that *ex-post* efficiency—trading whenever the buyer values the good more than the seller—is not possible in bargaining with two-sided incomplete information. Myerson and Satterthwaite (1983) and Williams (1987) derived the theoretical *ex-ante* efficient frontier, but, as Williams emphasized, “little is known about whether or not these limits can be achieved with ‘realistic’ bargaining procedures.”² This paper is the first attempt to bring data to this question. I develop a framework to estimate distributions of private valuations of both buyers and sellers who participate in bargaining following wholesale used-auto auctions. I then map these primitives into results from the theoretical mechanism design literature to compare real-world outcomes to efficient outcomes.

The question of whether real-world bargaining is efficient is one that cannot be addressed in a standard non-strategic framework (e.g. some form of Nash bargaining) or even a strategic alternating-offer game (e.g. Rubinstein 1982). These frameworks entail *complete information* and thus presume knowledge *a priori* that bargaining is perfectly efficient: in such a world, bargaining is never even attempted unless agreement is the efficient outcome. Treating bargaining as efficient, if it is in fact not, can result in incorrect market design recommendations or misleading calculations for welfare or pricing, or an incorrect understanding of bargaining power. The data and methodology I use in this paper allow me to study whether or not bargaining is actually efficient, rather than assuming it to be so.

Moreover, this question is indeed an empirical question—one that theory alone cannot address. Theoretical work by Chatterjee and Samuelson (1983), Satterthwaite and Williams (1989), Ausubel and Deneckere (1993), and Ausubel, Cramton, and Deneckere (2002) demonstrated that certain knife-edge or limiting cases of bargaining games may reach the theoretical *ex-ante* efficient frontier, but the limits of practical bargaining are unknown. Also, while the large theoretical literature on incomplete-information strategic bargaining has yielded valuable insights, it has done so primarily through a focus on special cases (see Appendix Table A1 for a summary of this literature). The

¹Fudenberg and Tirole (1991) stated, “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.” Fudenberg, Levine, and Tirole (1985) similarly commented “We fear that in this case [of two-sided incomplete information], few generalizations will be possible, and that even for convenient specifications of the functional form of the distribution of valuations, the problem of characterizing the equilibria will be quite difficult.” Very little work—theoretical or empirical—on bargaining with two-sided uncertainty and continuous valuations has been published before or after this time.

²In the language of Holmström and Myerson (1983), the term *ex-ante* refers to before the players learn their values and before the outcome of the bargaining is realized, and the term *ex-post* refers to after the valuations and bargaining outcomes are realized. As explained below, the *ex-ante* efficient frontier describes the limits on possible combinations of buyer and seller surplus that can be achieved under *any* bilateral bargaining mechanism in the presence of incomplete information.

general case, with alternating offers, two-sided incomplete information, and continuous valuations has received little attention because it involves complex signaling and updating by both parties. It is known to have multiple equilibria, some of which are very inefficient, but no canonical model or equilibrium characterization exists for the general setting examined in this paper.

To overcome these challenges, I take advantage of a unique, new dataset and novel empirical approach. The data consists of several hundred thousand sequences of bargaining offers between buyers and sellers at wholesale used-car auctions. It is the first bargaining dataset of this volume and detail to be analyzed in the literature, containing not only final negotiated prices on consummated deals, as most empirical bargaining datasets likely would, but also all of the back-and-forth bargaining offers between negotiating parties, and even all cases where bargaining failed to yield an agreement. The data also contains detailed information on cars and characteristics of the sale.

The empirical setting, described in Section 2, is a large market of business-to-business transactions where new- and used-car dealers buy vehicles from other dealers as well as from rental companies and banks. This industry represents the backbone of the supply side of the US used-car market, with 15 million cars annually passing through its lanes, totaling \$80 billion in sales. For each car, the auction house runs a secret-reserve-price ascending auction, followed by bargaining if the auction price falls short of the secret reserve price (which occurs more than two-thirds of the time). It is in this bargaining stage of the game that trade can fail, and thus understanding the efficiency of the bargaining is key to understanding the efficiency of the overall market. Industry wide, about 40% of sales attempts result in no trade. Why do these trades fail? These could be cases where the seller values the good more than the buyer, and hence no trade should occur even in a fully efficient world; these could be cases where gains from trade do exist, but trade fails due to the information constraints highlighted by Myerson and Satterthwaite (1983); or these may be cases where trade fails because of the particular bargaining protocol employed or the particular equilibrium played. These questions are the focus of this paper.

The bargaining I study takes place after an auction.³ This is not an unfortunate characteristic of the data, but rather a useful feature in studying what might otherwise be an intractable problem. Indeed, it is quite difficult to make any progress studying incomplete-information bargaining empirically without some kind of special lever. Specifically, given that no canonical model or even characterization of equilibria exists for such games, theory provides no obvious mapping from observables to primitives. In my setting, however, the distribution of buyer valuations can be estimated using the auction data, and bounds on the distribution of seller valuations can be obtained using sellers' responses to the first bargaining offer (the auction price). Each of these

³Many other settings similarly constitute an auction followed by bargaining, where one party collects initial bids from a number of different bidders in an auction-like stage and then selects a single bidder with whom to negotiate a final deal. Examples include business-to-business settings (e.g. procuring subcontractors), government settings (e.g. procuring services or selling government property), or private settings (e.g. selling a home). See examples in Elyakime, Laffont, Loisel, and Vuong (1997), Wang (2000), Huh and Park (2010), and An and Tang (2018).

steps imposes only minimal assumptions on the structure of the bargaining game.

I lay out a model in Section 3 and demonstrate several theoretical properties that aid in estimating model primitives. I show that the precise effect of the auction on the bilateral bargaining game is twofold: first, the auction leads to a truncation of the lower bound of the support of types who bargain. Thus, the game I study is analogous to a setting of bargaining alone where the lower bound of the support of the types in the bargaining game differs across realizations of the game in a tractable manner determined by the realization of the auction price, and the results herein average over these realizations. Second, the auction price is the first offer in the bargaining game and provides a lower bound on achievable prices in the bargaining game, similar to how a list price would provide an upper bound in many other real-world bargaining games, such as haggling over a car at a retail outlet. I also demonstrate in the model that game-level heterogeneity affects the game's outcomes in a tractable manner.

My approach to bounding the distribution of seller valuations is similar in spirit to Haile and Tamer (2003), using inequalities implied by very basic assumptions about players' rationality to learn about model primitives without imposing a complete model of the game or solving for an equilibrium. The bargaining setting is more complicated than the auction setting in Haile and Tamer (2003), however, in that it is not necessarily the case that an upper and lower bound on the valuation is observed for each individual observation in the data; instead, I obtain conditional probability statements that bound the whole distribution of valuations. This methodology is new to the empirical bargaining literature, and can be applied in alternating-offer bargaining settings, regardless of whether the bargaining follows an auction, when the econometrician observes the first offer and the response to that offer.

In order to compute expected gains from trade to measure efficiency in the real-world mechanism, it is necessary to know not only the distributions of valuations but also which player types trade and which do not. This is an equilibrium object and, as highlighted above, existing theory provides no guidance on identifying the equilibrium of games involving two-sided incomplete-information bargaining. I demonstrate, however, that even without solving explicitly for equilibrium strategies the direct-revelation mechanism corresponding to the equilibrium of the real-world game is identified in the data. This argument relies on the Revelation Principle, which has been exploited widely in the theoretical mechanism design literature. Applying this concept to my empirical setting allows me to avoid solving for or characterizing the actual equilibrium of the game and instead work with the direct mechanism corresponding to this game as implied by the data.

Section 4 describes each step of my estimation approach, which exploits the model's properties. After controlling for observable heterogeneity, I use a likelihood approach to deconvolve unobserved game-level heterogeneity and estimate buyer valuations using an order statistic inversion. I then estimate bounds on seller valuations, exploiting revealed preferences inequalities. I estimate the mapping between auction prices and the lower bounds of the support of buyer and seller types in

the bargaining game as well as the mapping corresponding to the direct revelation mechanism of the game. These mappings and the seller valuation bounds can each be estimated using flexible spline approximations within a constrained least squares framework.

After estimating these structural objects, I describe in Section 5 how I compute welfare under counterfactual efficient bargaining mechanisms. These counterfactual mechanisms are related to results derived in Myerson and Satterthwaite (1983) and Williams (1987), but are more complex to compute than the mechanisms they study because the distributions I estimate do not satisfy the regularity assumptions exploited by Myerson and Satterthwaite (1983) and Williams (1987) (defined below in Section 5) to simplify their analysis. I must therefore impose incentive compatibility numerically. Also, having only bounds on seller valuations, and not point estimates, I must perform a large numerical search to find bounds on efficiency measures. I ease this computational burden by deriving useful monotonicity properties that allow me to obtain bounds for some welfare measures directly using bounds on the distribution of seller valuations.

In Section 6 I then compare outcomes under efficient bargaining to those under the real bargaining to measure the relative efficiency. The first type of efficiency loss I measure is the loss due solely to incomplete information. Ideally, a buyer and seller should trade whenever the buyer values the good more than the seller (ex-post efficient trade). However, the celebrated Myerson and Satterthwaite (1983) Theorem demonstrated that, when the supports of buyer and seller types overlap, there does not exist any incentive-compatible, individually rational bargaining mechanism that is ex-post efficient and that also satisfies an ex-ante balanced budget. Williams (1987) then derived the entire ex-ante efficient frontier for any range of relative weights placed on the buyer's and seller's expected gains from trade. This frontier describes the limits on buyer and seller surplus that can be achieved by any incentive-compatible, individually rational, budget-balancing mechanism. I highlight several mechanisms along the ex-ante efficient frontier: the mechanism that places equal welfare weight on the buyer and seller surplus, which I refer to as *the second-best mechanism*; the mechanism placing all welfare weight on the seller's surplus (the seller-optimal mechanism); and the mechanism placing all welfare weight on the buyer's surplus (the buyer-optimal mechanism). The gap between the ex-ante and ex-post efficient frontiers represents an efficiency loss due to the presence of incomplete information. Using the estimated distributions, I find that incomplete information per se need not be a huge problem in this market: the second-best mechanism achieves about the same range of expected surplus as the infeasible ex-post efficient mechanism. The efficiency loss due solely to incomplete information is about \$9–59 for cars sold by dealers and \$9–77 for cars sold by fleet or lease institutions. The second-best mechanism falls short of ex-post efficiency in terms of the probability of trade by 3 to 16 percentage points, but these trades that the second-best mechanism fails to capture appear to be low-surplus trades.

The second type of efficiency loss I measure compares the real-world bargaining to the ex-ante efficient frontier. The real bargaining may fall short of this frontier for several reasons. First,

it is well known that, unlike the mechanisms discussed in Myerson and Satterthwaite (1983) and Williams (1987), real-world bargaining with two-sided uncertainty has no clear equilibrium predictions due to signaling by both parties, and many qualitatively different equilibria exist (see Ausubel and Deneckere 1993). The equilibrium play observed in the data may correspond to a particularly inefficient equilibrium. Second, it may be the case that the alternating-offer protocol used in this market is inefficient regardless of the equilibrium played; it may indeed be the case that a more efficient, practical protocol exists. Third, it may be that the real bargaining falls short of the theoretically efficient benchmark because that benchmark fails to satisfy other constraints that real-world bargaining satisfies, such as having rules that are simple for players to understand or being implementable without requiring the strong assumption that players and the market designer all have common knowledge of players' valuation distributions and beliefs (an assumption of traditional mechanism design critiqued in the influential Wilson doctrine, Wilson 1986). Because I place very little structure on the bargaining game, my analysis allows for any of these three cases to occur. Any of these cases can lead to a gap between the real outcome and that of an ex-ante efficient mechanism.

My findings indicate that the real bargaining falls short of the second-best by \$377–1,123 for cars sold by dealers and by \$223–834 for cars sold by large fleet or lease institutions. The losses of the real-world mechanism compared to the ex-post efficient frontier are similar in magnitude. These losses represent 17–23% of the ex-post gains from trade for cars sold by dealers and 12–20% for cars sold by large institutions. In terms of the probability of trade, the real-world bargaining falls short of the ex-post efficient outcome by 0.172–0.225 for cars sold by dealers and by 0.199–0.235 for cars sold by fleet and lease sellers. This implies that about 17–24% of negotiations constitute cases where the buyer indeed values the good more than the seller and yet the negotiation fails. Given that the overall rate of trade failure in the bargaining stage is about 35% in each sample, this suggests that over half of *failed* trades are cases where gains from trade exist but the parties do not trade, and the remainder of failed trades are cases where no gains from trade exist. The key takeaway of my analysis is that the real-world bargaining in this market is indeed inefficient and that this inefficiency is not solely due to the information constraints highlighted in Myerson and Satterthwaite (1983).

1 Related Literature

To my knowledge, this paper is the first to bring data to the bargaining efficiency framework of Myerson and Satterthwaite (1983). Unlike the vast structural auction literature—where researchers identify primitives to study various counterfactuals by modeling the game as one of incomplete information and strategic behavior—structural studies analyzing bargaining through a strategic,

incomplete-information lens are rare.⁴ Several exceptions that estimate models of one-sided incomplete information include Sieg (2000) and Silveira (2017), who focused on take-it-or-leave-it bargaining in trial settings, and Ambrus, Chaney, and Salitsky (2018), who studied pirate ransom negotiations and modeled bargaining following the theoretical work of Fudenberg, Levine, and Tirole (1985). Structural empirical work that highlights a role for *two-sided* uncertainty in bargaining (i.e. where both parties have private information) includes Genesove (1991), who discussed briefly the bargaining that takes place at wholesale auto auctions. Lacking detailed data on bargaining, he tested several parametric specifications for buyer and seller distributions and found that these assumptions performed poorly in explaining when bargaining occurred or when it was successful. Li and Liu (2015) studied identification of valuations in a static, two-sided incomplete-information bargaining game (a k double auction).

Another strand of the literature offers reduced-form analysis of implications of incomplete information bargaining. These studies include Merlo and Ortalo-Magne (2004), studying home sales in the U.K.; Scott Morton, Silva-Risso, and Zettelmeyer (2011), studying survey data from retail car buyers; Bagwell, Staiger, and Yurukoglu (2017), studying international trade negotiations; Backus, Blake, Larsen, and Tadelis (2018), studying alternating-offer bargaining data in an online marketplace; and Grennan and Swanson (2019), studying information disclosure in hospital-supplier bargaining.⁵ The data I analyze is new to the literature, and is particularly novel in the opportunity it presents for analyzing bargaining in detail, as it contains hundreds of thousands of observations and rich information about the characteristics of the goods sold and the actions players take during each observation of the game.

The only previous structural analysis of actual back-and-forth offers is Keniston (2011), but the setting, methodology, and focus of the two papers are quite distinct. Keniston (2011) collected several thousand observations of back-and-forth bargaining offers between riders and autorickshaw drivers in India, whereas my setting studies professionals engaging in business-to-business negotiations. The model of Keniston (2011) allowed for two-sided incomplete information, like mine, but the author embedded this model in a search-and-matching framework to model agents' outside options, whereas my paper does not explicitly model players' continuation payoffs when bargaining fails. The method of Keniston (2011) requires estimating beliefs in the bargaining subgame, relying on the assumption of a stationary equilibrium, whereas my approach does not require stationarity

⁴In a separate strand of the structural bargaining literature, a number of papers have made valuable contributions by abstracting away from incomplete information and modeling negotiated prices as arising from a Nash-bargaining surplus-splitting rule, such as Crawford and Yurukoglu (2012) and other subsequent studies of bargaining settings with externalities, and other work studying post-auction bargaining settings (Elyakime, Laffont, Loisel, and Vuong 1997; An and Tang 2018). Merlo and Tang (2012) provided identification arguments for stochastic bargaining games of complete information, and Merlo and Tang (2018) and Watanabe (2009) studied complete-information games with asymmetric priors.

⁵Merlo, Ortalo-Magné, and Rust (2015) provided a structural model of the home-sales data from Merlo and Ortalo-Magne (2004) but abstracted away from bargaining actions in order to focus on the seller's dynamic choice of list price.

assumptions or belief estimation. Keniston (2011) does not focus on the efficiency of bargaining or the Myerson-Satterthwaite Theorem, but instead compares welfare under bargaining to welfare under a fixed-price mechanism.

The approach developed in my paper can be applied to other settings with alternating-offer data to identify and estimate bounds on the distribution of valuations for the player who responds to the first offer. Larsen and Zhang (2018) presented an approach that can be used to instead obtain the distribution of valuations in bargaining games for the player who *makes* (rather than responds to) the first offer. Larsen and Zhang (2018) applied their approach to a subset of the data used in this paper to analyze the full auction-plus-bargaining mechanism rather than the bilateral bargaining studied in this paper, finding similar qualitative results regarding mechanism efficiency.

2 The Wholesale Used-Car Industry and the Data

The wholesale used-auto auction industry provides liquidity to the supply side of the US used-car market. Each year approximately 40 million used cars are sold in the United States, 15 million of which pass through a wholesale auction house. Industry wide, about 60% of these cars sell, with an average price between \$8,000 and \$9,000, totaling to over \$80 billion in revenue (NAAA 2009). The industry consists of approximately 320 auction houses scattered across the country. Throughout the industry, the majority of auction house revenue comes from fees paid by the buyer and seller when trade occurs. Buyers attending wholesale auto auctions are used-car dealers. Sellers may be car dealers (whom I will refer to as “dealers”) selling off extra inventory, or they may be large institutions, such as banks, manufacturers, or rental companies (whom I will refer to as “fleet/lease”) selling repossessed, off-lease, lease-buy-back, or old fleet vehicles.

Sellers bring their cars to the auction house, usually several days before the sale, and establish a secret reserve price. In the days preceding the sale, potential buyers may view car details and pictures online, including a condition report for cars sold by fleet/lease sellers, or may visit the auction house to inspect and test drive cars (although very few visit prior to the day of sale). The auction sale takes place in a large, warehouse-like room with 8–16 lanes running through it. In each lane there is a separate auctioneer, and lanes run simultaneously. A car is driven to the front of the lane and the auctioneer calls out bids, raising the price until only one bidder remains.

If the auction price exceeds the secret reserve price, the car is awarded to the high bidder. If the auction price is below the secret reserve price, the high bidder is given the option to enter into bargaining with the seller. If the high bidder opts to bargain, the auction house will contact the seller by phone (or in person, if the seller is present at the sale), at which point the seller can accept the auction price, end the negotiations, or propose some counteroffer higher than the auction price.⁶

⁶If the seller is not present and the auctioneer observes that the auction price and the reserve price are far enough apart that phone bargaining is very unlikely to succeed, the auctioneer may choose to reject the auction price on

If the seller counters, the auction house calls the buyer. Bargaining continues in this fashion until one party accepts or terminates negotiations (with the typical time between calls being 2-3 hours). It is this bilateral bargaining that is the focus on this paper.

The dataset used in this paper is new to the literature. The data come from six auction houses owned by one company, each maintaining a large market share in the region in which it operates. The sample period is from January 2007 to March 2010. An observation in the dataset represents a *run* of the vehicle, that is, a distinct attempt to sell the vehicle through the mechanism. For a given run, the data records the date, time, auction house location, and auction lane, as well as the seller's secret reserve price, the auction price, and, when bargaining occurs over the phone, the full sequence of buyer and seller actions (accept, quit, or counter), and the amounts of any offers/counteroffers. The data also records detailed characteristics of each car and sale. I drop a number of observations, such as those with missing variables or extreme price realizations (lying outside the lowest or highest 0.01 percentiles). I also drop car types (make-model-year-trim-age combinations) that are not offered for sale at least ten times in my sample. Appendix C.1 contains a list of all of my sample restrictions and a list of car observable characteristics that I use in estimation. In the end, I am left with 133,523 runs of cars offered for sale by used-car dealers (which I will refer to as the dealers sample), and 131,443 offered for sale by fleet/lease sellers (which I will refer to as the fleet/lease sample).

Descriptive statistics for these samples are displayed in Table 1 (with additional descriptive statistics shown in Appendix Table A3). The probability of trade is 0.705 in the dealers sample and 0.768 in the fleet/lease sample; in both of these subsamples, this trade probability is higher than the industry-wide average highlighted above (due primarily to my sample restrictions, such as focusing on certain make-model-year-trim-age combinations). In the dealers sample, the average auction price is over \$1,000 below the average reserve price and about \$600 below the average blue book price. Dealer cars are on average seven years old and have nearly 100,000 miles on the odometer. Fleet/lease cars tend to be newer (three years old and 57,000 miles), higher priced, and have a smaller gap between the reserve and auction prices. Also, unlike dealer cars, for fleet/lease cars the reserve price does not exceed the blue book price on average. All of these descriptive statistics are consistent with conversations with industry participants: dealer cars tend to be older cars with more aggressive reserve prices and tend to be less likely to sell.

Table 1 also shows information on the number of bidders participating in the auction. A precise measure of the number of bidders is difficult to obtain at these auctions, as many sales take place

behalf of the seller. If both the buyer and seller are present at the auction sale, a quick round of bargaining may sometimes take place in person immediately following the auction, but such behavior is discouraged as it delays the next sale; each auction typically takes 30–90 seconds, and inserting in-person negotiations into that procedure could drastically increase that time. Furthermore, the auction house discourages in-person interactions less they lead parties to transact off site in attempts to avoid auction house fees. Such off-site transacting is generally prevented by social norms, but in extreme cases violators could be punished through the auction house revoking access to future sales.

simultaneously in different auction lanes and bidders are not required to register for the sale of a specific car. However, for some auction sales, the company offers live video streaming and a web-based portal for remote bidding, and for these sales I can obtain a lower bound on the number of bidders from bid logs. These bid logs record each bid and the identity of the bidder if the bidder participated online. If the bidder was instead physically present on the auction house floor, the bid log only records the amount of the bid and an indicator, “floor”, rather than an identity. A lower bound on the number of distinct bidders is given by the number of distinct online identities who placed bids plus 1 if the log records any floor bids or plus 2 if the log records two consecutive floor bids (assuming no bidder bids against himself). This lower bound rarely falls below 2 (this occurs in 0.37% of observations in the dealers sample and 1.76% of observations in the fleet/lease sample). The mean of this lower bound conditional on it being at least 2 is 2.924 in the dealers sample and 2.973 in the fleet/lease sample. The distribution of this lower bound will be used in estimation in Section 4.

As data on actual back-and-forth offers is rare in the literature, I provide a period-by-period summary of this data in Table 2 (for the dealers sample). Outcomes in this table are separated by the period of the game in which the observed sequence ends. Period 1 is the auction. Observations ending in period 1 represent cases that ended with auction price exceeding the reserve price or with the auction price falling short of the reserve price and the buyer opting out of bargaining. The remaining periods are labeled with even numbers for seller turns and odd numbers for buyer turns. Table 2 demonstrates that in 10.66% of the dealers sample the game ends at the auction, and in these cases the final price when trade happens (which occurs 88.58% of the time) is naturally the auction price. The remainder of the time, the buyer opts out of bargaining. Observations ending in trade in the second period also have the final price equal to the auction price (as the auction price is the first bargaining offer). Consider now the fifth period of the game. Only 1.25% of the full sample reaches this period, but this still consists of nearly 1,700 observations. In the fifth period, when trade does occur, it occurs at an average final price of \$7,792, which is over \$600 above the average auction price (\$7,174), but still does not reach as high as the average reserve price (\$8,640). Overall, Table 2 suggests that observations ending in later periods had somewhat higher reserve prices than those ending in earlier periods. Only one buyer-seller pair in the data endured ten periods of the game, coming to agreement in the end, at a price \$2,600 above the auction price. Appendix Table A4 displays similar patterns for the fleet/lease sample. In the fleet/lease sample, the game ends at the auction 34.39% of the time. Thus, in both the dealers and fleet/lease samples, what happens *after* the auction plays a major role in the market.

3 Model

This section presents a model of the game played in wholesale used-car markets. Prior to stating the assumptions of the model, I first restate the timing of the game, which is as follows:

1. Seller sets a secret reserve price, R .
2. N bidders bid in an ascending auction.
3. If the auction price, P^A , exceeds the secret reserve price, the high bidder wins the item.
4. If the auction price does not exceed the secret reserve price, the high bidder is given the opportunity to walk away, or to enter into bargaining with the seller.
5. If the high bidder chooses to enter bargaining, the auction price becomes the first bargaining offer, and the high bidder and seller enter an alternating-offer bargaining game, mediated by the auction house.

Throughout I maintain the following assumptions:

Assumptions.

- (A1) $N \geq 2$ risk-neutral bidders participate in an ascending button auction with zero participation costs. For $i = 1, \dots, N$, each buyer i has a private valuation $\tilde{B}_i = W + B_i$, with $B_i \sim F_B$ and $W \sim F_W$, and with $(W, N, \{B_i\}_{i=1}^N)$ mutually independent.
- (A2) A risk-neutral seller has a private valuation $\tilde{S} = W + S$, with $S \sim F_S$ and with S independent of $(W, N, \{B_i\}_{i=1}^N)$.
- (A3) The bargaining lasts for up to $T < \infty$ periods; buyers incur a common bargaining cost, $c_B > 0$, for each offer made; and sellers incur a common bargaining cost, $c_S > 0$, for each offer made.
- (A4) Strategies of the bargaining subgame are continuous in the auction price.
- (A5) S has density f_S and B_i has density f_B , where f_B is positive on $[b, \bar{b}]$.

The motivation for the independent private values framework is that, according to market participants, buyers—as well as dealer-type sellers—have valuations arising primarily from their local demand and inventory needs.⁷ Also, seller valuations can depend on the value at which

⁷These buyers come from a wide geographic area, with some participants driving long distances or even flying to attend the auction sale, and thus strong correlations between local demands and inventory needs among these buyers are not likely a major concern. While there is likely some common values component to wholesale auto auctions, accounting for this in estimation would be beyond the state of the methodological literature (positive identification results do not exist for valuations at common values ascending auctions; see Athey and Haile 2007). In conversations with market participants, buyers often claim to decide upon their willingness to pay before bidding begins, sometimes having a specific retail customer lined up for a particular car, also suggesting a strong private component to valuations (see also discussions on the popular industry blog, thetruthaboutcars.com, Lang 2011). Studying similar auto auctions in Korea, Roberts (2013) and Kim and Lee (2014) provided evidence that private values models fit bidder behavior well in these settings.

the car was assessed as a trade-in; for a bank or leasing company, valuations can arise from the size of the defaulted loan.⁸ The button auction assumption simplifies the analysis of the auction, but is also not an unreasonable approximation, as it is the auctioneer in this market who raises the price and not the bidders (unlike in oral English auction) and bid increments are small. The assumption of symmetric buyers is not restrictive in this setting given that the high bidder’s identity is generally not known to the seller during bargaining and given that, in a private values ascending auction, bidders’ auction strategies will not depend on the identities of other participants. The assumption that N is independent of buyer valuations rules out endogenous entry. In Appendix C.3.2, I document some evidence supporting this assumption, following the intuition derived in Aradillas-López, Gandhi, and Quint (2016).

The form of bargaining costs in Assumption A3 is found elsewhere in the theoretical bargaining literature (e.g. Perry 1986 and Cramton 1991), and prevents players from continuing to bargain even when no surplus is to be had. The cap on the number of periods T simplifies the proofs of many of the model properties. T is assumed to be known to the players but not necessarily to the econometrician (and similarly for c_B and c_S). Assumption A4 is a technical condition required for the differentiability of the seller’s payoff, exploited in the proof of Proposition 3 to prove strict monotonicity of the seller’s secret reserve price strategy.

The assumption of positive density for B_i in Assumption A5 only plays a role in preventing division by zero when I prove strict monotonicity of the seller’s secret reserve price (Proposition 3) and when I prove identification arguments in Appendix C.4–C.6. For the support of the seller density, f_S , I will use the notation $[\underline{s}, \bar{s}]$. My results do not rely on specifying whether the supports of B_i and S are finite or infinite. In estimation and in computing welfare measures, I choose large, finite values for the support bounds of B and S . In pinning down one tail condition empirically (discussed in estimation step 4 in Section 4), I will also assume that $\underline{s} \geq \underline{b}$, given that the seller will be guaranteed a price of at least \underline{b} from the auction.

The random variable W in Assumptions A1 and A2 is observed by all buyers and the seller and represents game-level heterogeneity. Conditional on W , buyers and sellers have independent private values, but unconditional on W valuations are correlated. In estimation, in Section 4, I consider W to be unobserved to the econometrician, and I incorporate an additional, additively separable game-level heterogeneity term $X'\gamma$ that is observable to both the econometrician and to the players. Incorporating this latter term, the seller’s value is $S + W + X'\gamma$ and buyer i ’s value is $B_i + W + X'\gamma$. I do not assume that B_i and S take on only positive values; this is because these random variables represent how the players value the car relative to the game-level heterogeneity component they all observe (so a negative B_i or S means that a buyer or seller values the car less than the observable

⁸These explanations for seller values are due to conversations with industry professionals. Note also that adverse selection from the seller possessing more knowledge about car quality than the buyer is likely small because of auction house information-revelation requirements and because sellers are not previous owners/drivers of the vehicles.

value of the car). Because these objects can be negative, I assume nothing that prevents the model from suggesting that players' overall valuations may be negative. However, in practice, when I estimate the pieces of my model, I find that the majority of the variation in these valuations arises from the observable heterogeneity term $X'\gamma$, and that this term has most of its mass above zero, and thus the additively separable model does not appear to be a bad approximation. I discuss this in Appendix C.2.2.

For the next several subsections, I will discuss properties of the game *conditional* on a realization of W , and thus I will omit W for notational simplicity and return to it when I incorporate game-level heterogeneity in Proposition 5. I ignore auction house fees in this analysis but discuss them in detail in Appendix D.

3.1 Payoffs

I model the game as follows. In period $t = 0$, the seller chooses her secret reserve price, $R = \rho(S)$, knowing only her type S . This choice of reserve price is not revealed to buyers, before or after the auction. In period $t = 1$, the ascending auction takes place. Let β_i denote bidder i 's auction strategy (a price at which bidder i drops out of the auction), and let the final auction price be denoted P^A . If $P^A \geq R$, the high bidder wins the car and the game ends. If $P^A < R$, the high bidder is given the opportunity to walk away (denoted $D_1^B = 1$), which ends the game, or now walk away ($D_1^B = 0$), entering into bargaining with the seller.

When the auction price is P^A , a high bidder of type B chooses to enter into bargaining rather than walk away whenever the payoff from entering into bargaining is non-negative. If the buyer chooses not to walk away, the buyer enters an alternating-offer bargaining game with the seller. In doing so the buyer immediately incurs a bargaining cost, $c_B > 0$, and this c_B will be incurred by the buyer at every offer he makes. The seller will incur a bargaining cost, $c_S > 0$, at each offer she makes. The first offer of the bargaining game is P^A . The game moves to period 2 of the game, in which the seller chooses $D_2^S \in \{A, Q, C\}$ —a choice to accept (A), quit (Q), or counter (C). If the seller chooses Q or A the game ends. If the seller chooses C , the seller specifies a counteroffer P_2^S , and play continues to period 3, with the buyer choosing $D_3^B \in \{A, Q, C\}$, and so on up to period T . If period T is reached, the player whose turn it is can only choose to accept or quit.

Throughout the game, bargaining offers must be weakly greater than the auction price. In practice, this is an understood norm at the auction house, and it is supported in the data (bargaining prices lie above the auction price nearly 100% of the time). This feature means that the auction price plays a similar role for the seller that a list price would play for a buyer in many other real-world haggling scenarios.⁹ In the auto auction setting, this ability of the seller to accept the auction

⁹In haggling in the presence of posted list price, a seller and buyer may negotiate over prices in a range *below* the list price but the buyer may at any point choose to end the game by returning to the list price and accepting it. In such a haggling setting, the list price can be thought of the first bargaining offer, just as the auction price is here.

price can either be modeled as an additional action available to the seller at any of her turns, or can be modeled as a rule enforced by the auction house that all bargaining offers must lie weakly above the auction price. I follow the latter approach.

The payoffs in the game are as follows. If a buyer of type B and a seller of type S agree to trade at a price P , the buyer's payoff is $B - P$ less the per-offer bargaining costs the buyer has incurred up to that point. If trade occurs in round 1 of the game (i.e. at the auction), the buyer's payoff will be $B - P$, with $P = P^A$, the auction price. Similarly, if the buyer and seller agree to trade at a price of P , the seller's payoff is P , less any bargaining costs incurred by the seller up to that point.

When disagreement occurs, the buyer receives a payoff of zero and the seller a payoff of S (less any incurred costs). This modeling choice is one of the abstractions (and limitations) of the model. In practice, a buyer who fails to acquire a car may choose to later re-enter the market to bid on a similar car. The approach I adopt—treating buyers' outside option as a 0 payoff—means that the object I model as the buyer's valuation is actually the buyer's full willingness to pay minus a discounted continuation value of re-entering the market. Similarly, what I model as the seller's value S is in practice the seller's discounted continuation value of re-entering the market to attempt to sell the car again at the auction house, at a competing wholesale outlet, or at her own lot. These abstractions are appropriate under the following interpretation of my counterfactual exercises: For a given buyer and seller pair who meet in bargaining today, holding fixed their continuation values of re-entering the market, how would their expected gains from trade improve if today's bargaining game were efficient? Where these abstractions become a limitation is that they do not allow me to model how players' continuation values might change if the bargaining mechanism were to change *permanently*. In Appendix B.4, I demonstrate that my qualitative and quantitative findings are similar in several analyses that cut the data based on variables related to players' continuation values. These analyses do not alleviate all concerns associated with ignoring these continuation-game dynamics. I ignore these dynamics across games in order to focus on dynamics within the game; studying instead the dynamics across games would be an interesting avenue for future research.¹⁰

3.2 Equilibrium Concept

In what follows, I will focus on pure strategy Bayesian Nash equilibria (BNE). A BNE of the game is as follows. Let H_t represent the history of offers, including the auction price, up through period $t - 1$ of the bargaining game. The strategy of a buyer of type b_i is a history-contingent set of actions $\sigma^B(b_i) = \{\beta_i, \{D_t^B|H_t\}, \{P_t^B|H_t\}\}$, where the decisions D_t^B and offers P_t^B included are

¹⁰Note that players' continuation values *within* a given game are addressed in the model; see Appendix A; it is only players' continuation values *across* instances of the game that I abstract away from.

those for periods in which it is the buyer's turn.¹¹ The strategy of a seller of type s is a history-contingent set of actions $\sigma^S(s) = \{\rho, \{D_t^S|H_t\}, \{P_t^S|H_t\}\}$, where the decisions and offers are those for periods in which it is the seller's turn. A set of strategies $\sigma^{B^*}(b_i)$ for all buyers and $\sigma^{S^*}(s)$ for the seller constitutes a BNE of this game if, for each player, his or her strategy is a best response to opponents' strategies and players update their beliefs about opponent valuations using Bayes rule at each history of the game that is reached with positive probability.¹²

It is simple to derive a multiplicity of equilibria of the game, such as the following three examples (none of which need violate Assumption A4):

Three Examples of Equilibria of the Bargaining Subgame:

1. *Sellers only accept or quit at $t = 2$, and buyers reject all (off-equilibrium) offers at $t = 3$.*
2. *Sellers make uninformative offers (equal to \bar{s} , say) at $t = 2$, buyers counter at $t = 3$, and sellers only accept or quit at $t = 4$. Buyers reject all off-equilibrium offers at $t = 3$ or $t = 5$.*
3. *All offers and counteroffers must lie within a particular set of possible values, and in the (off-equilibrium) case in which any player deviates from these offers, the opponent responds by quitting.*

Ausubel and Deneckere (1993) provided a discussion of other partial-pooling equilibria for a similar bargaining game but with one-sided offers, and Ausubel, Cramton, and Deneckere (2002) suggested that such arguments can be extended to two-sided offer games as well.

3.3 Mechanism Design Framework for Evaluating Bargaining Efficiency

Prior to deriving the properties of BNE of this game, I describe the mechanism design framework I use to assess efficiency of bargaining, as it is the motivation for deriving some of the game's properties. By the Revelation Principle (Myerson 1979), any BNE of an incomplete-information trading game has a corresponding, payoff-equivalent, direct-revelation mechanism. In a direct mechanism, a buyer of type b and seller of type s report their true types to the mechanism designer and then trade occurs with probability $x(s, b)$ (the allocation function), where this allocation function is determined so that players receive the same expected outcomes as in the original game.

¹¹This discussion ignores the possibility of buyers conditioning their auction strategies on information observed during the auction (such as the points at which opponents drop out); as shown in Proposition 1 below and discussed in Appendix B.2, bidders would not gain from conditioning on such information.

¹²Note that Perfect Bayes Equilibrium (PBE) is a refinement of BNE (and thus, every PBE is also a BNE) requiring that the researcher also specify how beliefs are updated at histories of the game that are never reached in equilibrium. I focus on the broader equilibrium concept, BNE, because the PBE concept does not meaningfully narrow down the set of equilibria in sequential bargaining games of incomplete information (see discussion in Gul and Sonnenschein 1988) and because none of my identification or estimation arguments rely on specifying how beliefs are updated after zero-probability events. See Appendix B.1 for more discussion of the equilibrium concept.

The allocation function corresponding to ex-post efficient trade is simply $x^*(s, b) \equiv 1\{s \leq b\}$. The allocation function corresponding to a given point along the ex-ante efficient frontier, on the other hand, will maximize a convex combination of the buyer’s and seller’s ex-ante expected gains from trade, with weight η given to the seller’s gains and weight $1 - \eta$ given to the buyer’s. I will use the notation $x^\eta(\cdot)$, for a given $\eta \in [0, 1]$, to denote the allocation function corresponding to a point on the ex-ante efficient frontier. Computing $x^\eta(\cdot)$ boils down to solving a linear programming problem, described in Section 5 and Appendix C.7. The direct mechanism corresponding to the real-world bargaining, which I denote $x^{RW}(\cdot)$, can be estimated directly from the data, as described in Section 4. Computing each of these allocation functions requires estimates of F_B and F_S , the distributions of buyer and seller valuations. Thus, a key focus of this paper is the estimation of these distributions without imposing *a priori* any restrictions on how efficient the real-world bargaining is relative to these counterfactual benchmarks.

3.4 Model Properties

I now describe a number of properties of this game that will hold in any equilibrium.¹³ These properties will then be exploited in Section 4 to estimate the distributions of buyer and seller valuations, the support of types who enter the bargaining game, and the allocation function corresponding to the real-world mechanism.

Bidding Behavior. The first property concerns a bidder’s auction strategy, which is a price at which he will stop bidding:

Proposition 1. *If Assumption A1 holds then truthtelling is a weakly dominant action in the auction regardless of other players’ strategies in the auction or players’ strategies in the BNE of the continuation game.*

Proposition 1 states that bidders receive no positive benefit from deviating from a strategy involving dropping out of the bidding when, and only when, the auction price reaches their valuations (i.e. *truthtelling*). The intuition behind the proposition is that, as in a standard ascending button auction, a bidder will not find it optimal to drop out before the current price reaches his value because doing so would make the bidder miss out on a chance to win the auction. A bidder will also not find it optimal to remain in the auction once the current price passes his value because doing so will yield a negative payoff if the bidder does end up winning. One implication of the proposition is that a bidder will not gain from conditioning his auction strategy on any information revealed during the auction, such as other bidders’ drop-out points. Appendix B.2 expounds on this result and proves that bidders bidding above their valuations and then attempting to bargain

¹³Appendix B.5 discusses an extension of this model in which sellers have some uncertainty about the distribution of buyer valuations when choosing the reserve price, which addresses explicitly why sellers may accept offers below their secret reserve price. Appendix B.5 also provides several other explanations of this phenomenon.

to a lower final price later could not occur in equilibrium even if bargained prices below the auction price were allowed by the auction house.

There can exist BNE of this game in which bidders are indifferent between bidding truthfully and not. For example, one such BNE would consist of the seller setting a very high reserve price, all bidders dropping out at zero, and, in bargaining, the seller immediately rejecting any (off-equilibrium) positive auction price; in this equilibrium, bidders would receive a payoff of zero, but would receive no less by bidding truthfully. To rule out such cases, and motivated by Proposition 1, I make the following assumption:

Assumption. (A6) *All bidders follow the weakly dominant strategy of bidding truthfully.*

In practice, for identification and estimation, it is only the highest bid that plays any role.

Seller’s Choice to Accept the Auction Price or Quit. I now demonstrate that bounds on the distribution of seller valuations can be achieved by an argument similar to the Haile and Tamer (2003) bounds in English auction settings. The argument differs from Haile and Tamer (2003), however, in that it is not possible here to construct both an upper and lower bound on the seller’s value for each individual realization of the game. This is because, as shown below, a lower bound on a seller’s valuation is only observed when the seller chooses to quit. Therefore, rather than observation-level bounds, I will obtain bounds on the *distribution* of seller values relying on probability statements formed from observations of many sellers’ decisions to accept or walk away from an offer on the table.

Let $D^S = A$, without a t subscript (to distinguish this from the period-specific action described in Section 3.1), represent the event in which the seller takes an action in period 1 or 2 that results in the game ending in agreement at the auction price. This event occurs either when 1) the auction price exceeds the reserve price or 2) the auction price fails to meet the reserve price, the high bidder does not opt out of bargaining, and the seller accepts the auction price on her first bargaining turn. Similarly, let $D^S = Q$ represent the event in which the seller takes an action in period 2 that results in the game ending in disagreement at the auction price. This event happens when the auction price fails to meet the reserve price, the high bidder does not opt out of bargaining, and the seller quits on her first bargaining turn rather than accepting the auction price or making a counteroffer.¹⁴

I exploit the following assumption:

Assumption. (A7) *The seller never (i) accepts an auction price below her value or (ii) walks away from (quits at) an auction price above her value.*

The conditions in Assumption A7 will be satisfied in any BNE of the game.¹⁵ These conditions imply that, if the realized auction price is p^A and the seller accepts, it must be the case that the

¹⁴Note that the events $D^S = A$ and $D^S = Q$ as defined above are observable to the econometrician for every instance of the game recorded in the data, not just those in which bargaining occurs.

¹⁵To see this, suppose to the contrary that a seller’s strategy involves accepting an auction price $p^A < s$. A

seller values the good *less* than p^A . Similarly, if the seller quits when the auction price is p^A , it must be the case that the seller values keeping the car herself *more* than p^A . These conditions imply bounds on distribution of S :

$$\begin{aligned}\Pr(D^S = A|P^A = p^A) &\leq \Pr(S \leq p^A) = F_S(p^A) \\ \Pr(D^S = Q|P^A = p^A) &\leq \Pr(S \geq p^A) = 1 - F_S(p^A) \Rightarrow \Pr(D^S \neq Q|P^A = p^A) \geq F_S(p^A)\end{aligned}$$

I state these bounds as the following proposition, where \mathcal{F} represents the space of all possible CDFs (i.e. right-continuous, weakly increasing functions approaching 0 to the left and 1 to the right):

Proposition 2. *Under Assumptions A2 and A7, for any $v \in [\underline{s}, \bar{s}]$, any CDF of seller valuations $F_S \in \mathcal{F}$ must satisfy $F_S(v) \in [\Pr(D^S = A|P^A = v), \Pr(D^S \neq Q|P^A = v)]$.*

I now highlight several interesting features of these bounds. First, the bounds do not cross, because $D^S = A \Rightarrow D^S \neq Q$, and therefore $\Pr(D^S = A|P^A = v) \leq \Pr(D^S \neq Q|P^A = v)$. A monotonized version of these bounds may cross, however, and such a crossing would indicate a violation of Assumption A7; see discussion in Appendix B.3. Second, these bounds rely only on Assumption A2 (that is, that buyer and seller valuations are independent) and the conditions in Assumption A7. Under these assumptions alone, the bounds are sharp. However, under the additional assumptions imposed elsewhere in the paper (namely, that actions correspond to some BNE), the bounds are not necessarily sharp, but are conservative. Appendix B.3 discusses sharpness formally.

Third, the width of these bounds will be determined by the frequency with which (i) the seller chooses to make a counteroffer in response to the first offer or (ii) the buyer opts out of bargaining. Specifically, the bounds can be re-written

$$\Pr(D^S = A|P^A = p^A) \leq F_S(p^A) \leq \Pr(D^S = A|P^A = p^A) + \underbrace{\Pr(D^S \neq A \cap D^S \neq Q|P^A = p^A)}_{\text{Prob. seller counters or buyer opts out}}$$

The object $\Pr(D^S \neq A \cap D^S \neq Q|P^A = p^A)$ is the probability that either the seller makes a counteroffer in response to the first offer or the buyer opts out of bargaining. If, at a given $P^A = p^A$, the buyer doesn't opt out and the seller only accepts or quits (doesn't counter), this probability will be zero, and the bounds will collapse to a point equal to the probability of acceptance at that p^A , $\Pr(D^S = A|P^A = p^A)$.

These bounds can be applied to other alternating-offer bargaining settings, independent of profitable deviation from this strategy would be to quit instead, as this would yield a payoff of s . Now suppose a seller's strategy involves quitting when facing an auction price $p^A > s$. A profitable deviation from this strategy would be to accept instead, as this would yield a payoff of p^A . Thus, any BNE cannot involve violations of A7. Note that bargaining costs from Assumption A3 do not play a role in these bounds as those costs are only incurred by a player when making an offer, not when accepting or quitting.

whether the bargaining follows an auction. In such cases, the decisions $D^S = A$ or $D^S = Q$ would represent the first action taken by the player in the bargaining game who responds to the first offer.

The Lower Support of Buyer and Seller Types Who Bargain. One advantage of studying bargaining following an ascending auction is that the auction outcome affects the bargaining game in a tractable manner, allowing me to isolate the bargaining game from the auction. Let $\pi^B(p^A, b)$ represent the buyer's expected payoff from entering into bargaining conditional on his value b and the realization of the auction price. Let $\chi(b)$ be defined by $\pi^B(\chi(b), b) = 0$. The high bidder will end up in bargaining when $P^A < R$ and when $\pi^B(P^A, b) \geq 0$. The object $\chi^{-1}(p^A)$ then represents the buyer type that would be indifferent between bargaining and not bargaining when the realized auction price is p^A . As above, $\rho(\cdot)$ is the seller's secret reserve price strategy; that is, $R = \rho(S)$.

Proposition 3. *If Assumptions A1–A6 hold, then in any BNE satisfying Assumption A4, conditional on an auction price $P^A = p^A$ and conditional on bargaining occurring, the support of seller types in the bargaining game is $[\underline{s}(p^A), \bar{s}]$ and the support of buyer types is $[\underline{b}(p^A), \bar{b}]$, where $\underline{s}(\cdot) \equiv \rho^{-1}(\cdot)$ and $\underline{b}(\cdot) \equiv \chi^{-1}(\cdot)$. Moreover, $\rho(\cdot)$ and $\chi(\cdot)$ are strictly increasing, with $\rho(s) \geq s$ and $\chi^{-1}(p^A) > p^A$.*

The intuition behind this result is as follows. When the auction price is p^A and bargaining occurs, it will be common knowledge among the two bargaining parties that the seller's type s satisfies $\rho(s) \geq p^A$ (i.e. the reserve price is above the auction price), implying $s \in [\rho^{-1}(p^A), \bar{s}]$. Similarly, bargaining occurring means the buyer did not opt out, so $\chi(b) \geq p^A$, implying $b \in [\chi^{-1}(p^A), \bar{b}]$. Thus, the game I study is analogous to a setting of bargaining alone where the lower bound of the support of the types in the bargaining game differs across realizations of the game as determined by the realization of the auction price, and when I present welfare results later they will average over these realizations. The clean relationship between the auction and the bargaining game obtained in Propositions 1 and 3 would not exist if the pre-bargaining stage were a first-price auction rather than an ascending auction; the first-price auction would affect the bargaining (and vice versa) in an intractable manner.¹⁶

The proof of Proposition 3 also addresses the seller's choice of reserve price, demonstrating that $\rho(\cdot)$ is strictly monotone using a monotone comparative statics result from Edlin and Shannon (1998), a special case of Topkis's Theorem. Assumption A4, continuity of the equilibrium of the bargaining subgame in the auction price, is required to prove differentiability of the seller's payoff in order to apply the Edlin and Shannon (1998) result.

Proposition 3 implies that, given an auction price p^A , the distributions of buyer and seller types in bargaining are given by $\frac{F_B(b)}{1-F_B(\chi^{-1}(p^A))}$ and $\frac{F_S(s)}{1-F_S(\rho^{-1}(p^A))}$, respectively. These distributions corre-

¹⁶Elyakime, Laffont, Loisel, and Vuong (1997) discussed this issue and adopted a model in which a first-price auction takes place under incomplete information and post-auction bargaining takes place under complete information (Nash bargaining).

spond precisely to on-equilibrium-path Bayes updating of the buyer’s and seller’s beliefs about their opponents’ type given the actions occurring prior to the bargaining, as highlighted in Section 3.2. Also, the seller’s beliefs in the bargaining game do not condition on N , the number of bidders. This is due to a convenient property of the symmetric independent private values button auction: the distribution of the maximum order statistic (here, the valuation of the buyer entering bargaining) conditional on a lower order statistic (here, the auction price), does not depend on N , a result first shown in Song (2004) and extended to the unobserved heterogeneity case in Freyberger and Larsen (2017). Thus, the number of bidders does not enter into the seller’s beliefs about the density of buyer valuations she faces in bargaining once she knows the realization of P^A .

The Real-World Mechanism. The allocation function corresponding to the real-world mechanism, x^{RW} , satisfies the following property:

Proposition 4. *Under Assumptions A1–A6, in any BNE satisfying Assumption A4, the allocation function x^{RW} can be written as*

$$x^{RW}(r, b; p^A) \equiv 1 \{b \geq g(r, p^A)\} \quad (1)$$

where $g(r, p^A)$ is an unknown function that is weakly increasing in r .

Proposition 4 demonstrates that x^{RW} depends on a cutoff function defining the boundary between those types who trade and those who do not. Ausubel and Deneckere (1993) referred to this property as the “Northwestern Criterion” as it implies that trade occurs if and only if players’ types lie northwest of a boundary defined by g . The proof of Proposition 4 relies directly on an argument presented in Storms (2015), and also exploits the strict monotonicity of $\rho(\cdot)$ proved in Proposition 3, which makes it possible to model the allocation conditional on a realization of the reserve price, $R = r$, rather than conditional on the seller’s type. This is particularly useful in that it allows me to evaluate the allocation function for the real-world bargaining without knowing where the true distribution of seller valuations lies within the bounds from Proposition 2.

Game-level Heterogeneity. The above results are derived conditional on a given realization of game-level heterogeneity. I now consider the additively separable structure of buyer and seller valuations in the common component W .

Proposition 5. *Suppose, when $W = 0$, the equilibrium is such that the reserve price is r ; the auction price is p_A ; the lowest buyer type who would choose to bargain is $\chi^{-1}(p^A)$; and, for each period t at which the game arrives, the offer is given by $P_t = p_t$ and the decision to accept, quit, or counter is given by $D_t = d_t$. Then, under Assumptions A1–A6, when $W = w$, the equilibrium will be such that the reserve price is $\tilde{r} = r + w$; the auction price is $\tilde{p}_A = p_A + w$; the lowest buyer type*

who would choose to bargain is $\chi^{-1}(\tilde{p}^A - w) + w$; the period t decision is d_t ; and, for any period t offer that is accepted with positive probability, the period t offer is $p_t + w$.

Proposition 5 is similar to results used elsewhere in the empirical auctions literature (Haile, Hong, and Shum 2003; Asker 2010) but is a generalization specific to this setting. It implies that continuous actions of the game (reserve prices, auction prices, and bargaining offers) will be additively separable in W ; choice probabilities for discrete actions (accepting, declining, or countering in response to an offer) will be unaffected by the value of W . An immediate implication of Proposition 5 is that the allocation function is invariant to game-level heterogeneity; that is, $x^{RW}(r + w, b + w; p^A + w) = x^{RW}(r, b; p^A)$.

4 Estimating Valuations and the Bargaining Mechanism

In this section, I exploit the model properties derived above in order to estimate the distribution of buyer and seller valuations and the bargaining mechanism. Identification and estimation require the following additional assumptions on the data. Below, let F_R , F_{P^A} , and F_W represent the cumulative distribution functions of R , P^A , and W .

Assumptions.

(A8) F_R , F_{P^A} , and F_W have densities f_R , f_{P^A} , and f_W satisfying the following: (i) the characteristic functions of f_R and f_W have only isolated real zeros; (ii) the real zeros of the characteristic function of f_{P^A} and its derivative are disjoint; and (iii) $E[W]=0$.

(A9) The supports of S and B satisfy $\underline{s} \geq \underline{b}$.

(A10) Observations of random variables (S, B_i, W, N) across instances of the game are identically and independently distributed.

(A11) All observations in the data are generated by the same equilibrium.

Assumption A8 lists the sufficient conditions from Evdokimov and White (2012) for proving identification of f_R , f_{P^A} , and f_W .¹⁷ I use Assumption A9 in pinning down the left tail of the upper bound on the seller valuation CDF. Motivation for this assumption ($\underline{s} \geq \underline{b}$) is that any seller is guaranteed a price of at least \underline{b} from participating.

Assumption A10 is common in the empirical games literature, and it abstracts away from dynamics across instances of the game.¹⁸ Assumption A11 is not required for steps 1–4 below

¹⁷Evdokimov and White (2012) demonstrated that these are weaker conditions than those used previously in the empirical auctions literature in settings relying on convolution arguments (Li and Vuong 1998; Krasnokutskaya 2011). The assumption that $E[W] = 0$ is a location normalization, and this normalization could alternatively be placed on R or P^A without loss of generality.

¹⁸Appendix C.1 and Section B.4 discuss some simple ways in which I do analyze inter-game dynamics.

but is required for steps 5–6. For example, even if different equilibria of the bargaining subgame are played in different observations of the data, the distribution of buyer valuations can still be estimated (step 3) using the distribution of auction prices, as described below. Similarly, the revealed preference arguments used to bound the distribution of seller valuations (step 4) will still hold even if A11 fails. Steps 5–6, however, require inverting policy functions that will depend on the equilibrium of the game. Fortunately, none of the steps below, including 5–6, require fully specifying or solving for the equilibrium. Like Assumption A10, Assumption A11 is also common in the structural literature. The typical approach in the literature to handling cases where particular subsamples of the data are believed to have been generated by different equilibria is to estimate the model separately in these subsamples. In line with this, throughout the estimation, I treat the dealers and fleet/lease samples separately because, according to conversations with industry professionals, this is likely the most important division of the data in which behavior may differ (although, as highlighted in Section 6, I find very similar results between the two samples). I perform additional subsample analyses in Appendix B.4 and C.1.3.

I now provide an overview of each estimation step. I do not describe all of the technical details for each step here, but include them in Appendix C. Appendix C also contains nonparametric identification arguments, arguments for consistency of the estimates, and evidence of goodness of fit for each estimation step.

Step 1) Accounting for Observed Heterogeneity Empirically. To account for game-level characteristics that are observed to the econometrician as well as the players, I apply Proposition 5. Let R^{raw} and $P^{A,raw}$ be random variables representing the reserve price and auction price in the raw data, prior to any adjustments for heterogeneity. As above, let W be a random variable representing unobserved game-level heterogeneity. Let X be a random variable representing game-level heterogeneity that is instead observed (by the econometrician as well as the players), with X independent of W , S , B , N . Let realizations of R^{raw} , $P^{A,raw}$, X , and W for game j be denoted by lower case letters with subscript j .

I specify the total game-level heterogeneity (observed plus unobserved) for observation j to be $x'_j\gamma + w_j$, where γ is a vector of parameters to be estimated. Proposition 5 implies that auction prices and reserve prices can be “homogenized” (Haile, Hong, and Shum 2003) by estimating the following joint regression of reserve prices and auction prices on observables:

$$\begin{bmatrix} r_j^{raw} \\ p_j^{A,h} \end{bmatrix} = \begin{bmatrix} x'_j\gamma \\ x'_j\gamma \end{bmatrix} + \begin{bmatrix} \tilde{r}_j \\ \tilde{p}_j^A \end{bmatrix},$$

where $\tilde{r}_j = r_j + w_j$, $\tilde{p}_j^A = p_j^A + w_j$. In the vector x_j I include a rich vector of controls, including flexible mileage terms, dummies for each make-model-year-trim-age combination, and a number of other factors described in detail in Appendix C.1. An estimate of \tilde{r}_j is then given by subtracting $x'_j\hat{\gamma}$

from r_j^{raw} , and similarly for \tilde{p}_j^A . Variation in these two quantities is then attributed to unobserved game-level heterogeneity and to players' private valuations, as detailed below.

Step 2) Accounting for Unobserved Heterogeneity Empirically. To account for heterogeneity W in the game that is observed by the players but not by the econometrician, I apply a result due to Kotlarski (1967), which implies that observations of $\tilde{R} = R + W$ and $\tilde{P}^A = P^A + W$ (which are additively separable in W by Proposition 5) are sufficient to recover the densities f_W , f_R , and f_{PA} . This result has been applied elsewhere in first-price auction work (e.g. Li, Perrigne, and Vuong 2000; Krasnokutskaya 2011); my application of this deconvolution argument using instead an *ascending auction bid* and a *reserve price* to identify unobserved heterogeneity parallels Decarolis (2018) and Freyberger and Larsen (2017). I estimate these densities using a flexible maximum likelihood approach, where the likelihood of the joint density of (\tilde{R}, \tilde{P}^A) is given by

$$\mathcal{L}(f_{PA}, f_R, f_W) = \prod_j \left[\int f_{PA}(\tilde{p}_j^A - w) f_R(\tilde{r}_j - w) f_W(w) dw \right] \quad (2)$$

I approximate each of the densities f_{PA} , f_R , and f_W as Hermite polynomials, as suggested by Gallant and Nychka (1987) (I use fifth-order polynomials). This also yields estimates of the CDFs F_W , F_R , and F_{PA} . Appendix C.2 describes technical details and nonparametric identification.

Step 3) Estimating the Distribution of Buyer Valuations. I recover the distribution of buyer valuations, F_B , from the distribution of auction prices, F_{PA} , which, by Proposition 1, will coincide with the distribution of the second order statistic of buyer valuations. The relationship of F_{PA} and $\Pr(N = n)$ (the distribution of the number of bidders) to F_B is as follows:

$$F_{PA}(v) = \sum_n \Pr(N = n) [nF_B(v)^{n-1} - (n-1)F_B(v)^n] \quad (3)$$

The right-hand side of (3) is strictly monotonic in $F_B(\cdot)$, and thus F_B is nonparametrically identified by $\Pr(N = n)$ and F_{PA} (see, for example, Athey and Haile 2007). I estimate the object F_B by solving (3) numerically on a grid of values for v , pluggin in an estimate of $\widehat{\Pr}(N = n)$ and the maximum likelihood estimate $\hat{F}_{PA}(v)$ from (2).

To estimate $\widehat{\Pr}(N = n)$, I use the subsample of the data for which bid logs are available, in which I observe a lower bound on N that varies from auction to auction (see discussion in Section 2). I set $\widehat{\Pr}(N = n)$ equal to the empirical frequency with which this lower bound equals n . This treats the distribution of the lower bound as though it is the true distribution of the number of bidders. I gathered some additional independent data supporting this choice by physically attending over 200 auction sales and recording the number of bidders (see Appendix C.3.1). It turns out, however, that the choice of $\Pr(N = n)$ is, perhaps surprisingly, not critical to the welfare estimates of this paper. Specifically, the choice of $\Pr(N = n)$ affects the estimate of the full underlying buyer distribution,

F_B , but has a negligible effect on the transformation of F_B used in evaluating welfare, which is the distribution of the high bidder's value conditional on the auction price integrated against the auction price density (i.e. the maximum order statistic distribution). In Appendix C.3.1, I demonstrate numerically that this latter object is not sensitive to the choice of $\Pr(N = n)$. I also provide a new mathematical result, stated as Proposition 10, proving that, for a class of possible choices of $\Pr(N = n)$ (Poisson distributions with mean λ), conditional on \hat{F}_{PA} , the inferred maximum order statistic distribution has a derivative that is identically zero with respect to the choice of λ .¹⁹

Step 4) Estimating Bounds on the Distribution of Seller Valuations. Proposition 2 demonstrates that, absent unobserved game-level heterogeneity, the objects $\Pr(D^S = A|P^A = v)$ and $\Pr(D^S \neq Q|P^A = v)$ will provide bounds on $F_S(v)$. The same revealed preference arguments used to derive that result—Assumption A7—extend to the case of unobserved heterogeneity, providing bounds on the distribution of $\tilde{S} \equiv S + W$ conditional on $\tilde{P}^A \equiv P^A + W$. Here I also incorporate the bound provided by secret reserve prices themselves: $R \geq S \Rightarrow F_S(v) \geq F_R(v)$.

To describe these bounds, let $q(v; F_S) = \int F_S(v - w) \frac{M_S(v, w)}{\int M_S(v, z) dz} dw$, where $M_S(v, w) \equiv f_{PA}(v - w)f_W(w)$ is the joint density P^A and W . Bounds on F_S , which I denote $[F_S^L(v), F_S^U(v)]$, are given by the solution to the following minimization problem:

$$\min_{(F_S^L, F_S^U) \in \Phi} \left\{ \left| \Pr(D^S = A|\tilde{P}^A = v) - q(v; F_S^L) \right|^2 + \left| \Pr(D^S \neq Q|\tilde{P}^A = v) - q(v; F_S^U) \right|^2 \right\} \quad (4)$$

where $|\cdot|^2$ represents the L^2 -norm, and the set Φ is the set of feasible pairs of CDFs that can be bounds on F_S ; that is, $\Phi = \{F_S^L \in \mathcal{F}, F_S^U \in \mathcal{F} : F_S^U(v) \geq F_S^L(v) \geq F_R(v) \forall v \in [\underline{s}, \bar{s}]\}$, where \mathcal{F} is the set of possible CDFs, as defined in Section 3.4. Appendix C.4 demonstrates that these bounds are nonparametrically identified.

The objects F_S^L and F_S^U , as well as the functions χ^{-1} and g in steps 5 and 6 below, can be estimated using a minimum-distance, constrained least squares procedure. I will describe this approach in slightly more detail in this step and be more brief in my description in steps 5 and 6. Additional technical details for each of these steps are found in Appendices C.4–C.6. To estimate the functions F_S^L and F_S^U , I first parameterize each as a very flexible piecewise linear spline; I denote these approximations $F_S^L(\cdot, \theta^{S,L})$ and $F_S^U(\cdot, \theta^{S,U})$. Denote the fixed vector of spline knots $\{v_k^S\}_{k=1}^{K_S}$. I choose $K_S = 200$; as discussed in Appendix C.4, the estimates are not sensitive to this choice. I estimate the parameter vectors $\theta^{S,L}$ and $\theta^{S,U}$ using the following objective function:

$$\min_{\theta^{S,L}, \theta^{S,U}} \sum_{k=1}^{K_S} \left\{ \left[\widehat{\Pr}(D^S = A|\tilde{P}^A = v_k^S) \left(\int \hat{M}_S(v_k^S, z) dz \right) - \int F_S(v_k^S - w; \theta^{S,L}) \hat{M}_S(v_k^S, w) dw \right]^2 \right.$$

¹⁹This result is new to the literature and sheds some light on why many estimated welfare results from order-statistic-inversion estimates can be insensitive to the choice of $\Pr(N = n)$ used in the inversion, as I find to be the case in this study.

$$+ \left[\widehat{\Pr}(D^S \neq Q | \tilde{P}^A = v_k^S) \left(\int \hat{M}_S(v_k^S, z) dz \right) - \int F_S(v_k^S - w; \theta^{S,U}) \widehat{M}_S(v_k^S, w) dw \right]^2 \} \quad (5)$$

This approach searches for the lowest and highest possible values of F_S that can rationalize the observed behavior in the data described by the conditional probabilities $\Pr(D^S = A | \tilde{P}^A = v)$ and $\Pr(D^S \neq Q | \tilde{P}^A = v)$. I impose several constraints on the minimum distance problem in (5): (i) F_S^L lies graphically above F_R and graphically below F_S^U ; (ii) F_S^L and F_S^U lie in $[0, 1]$; (iii) F_S^L and F_S^U are weakly increasing; and (iv) $F_S^L(v)$ and $F_S^U(v)$ are equal to 0 for any $v < v_1^S$ and equal to 1 for any $v > v_{K_S}^S$. These last three constraints ensure that F_S^L and F_S^U will correspond to proper distribution functions. The only constraint of (iv) that binds in practice is that of the left tail of F_S^U . My assumption essentially bounds that left tail below by \underline{b} , as stated in Assumption A9. This is discussed in more detail in Appendix C.4.

Computing (5) requires first-step estimates of several other objects, including \hat{F}_R , \hat{f}_{P^A} , and \hat{f}_W , which come from the maximum likelihood procedure in (2). The procedure also requires the objects $\widehat{\Pr}(D^S = A | \tilde{P}^A = \tilde{p}^A)$, and $\widehat{\Pr}(D^S \neq Q | \tilde{P}^A = \tilde{p}^A)$, which I estimate using local linear regressions.

Step 5) Estimating the Lower Support of Bargaining Types. By Proposition 3, both $\underline{b}(\cdot) \equiv \chi^{-1}(\cdot)$ and $\underline{s}(\cdot) \equiv \rho^{-1}(\cdot)$ are increasing functions, and correspond to the lower support of buyer and seller types who enter the bargaining game. For any function $F_S(\cdot)$ lying in the estimated bounds $[\hat{F}_S^L(\cdot), \hat{F}_S^U(\cdot)]$, the function $\rho(s)$ can be constructed as $\rho(s) = F_R^{-1}(F_S(s))$, with F_R replaced with the estimated \hat{F}_R from (2). Similarly, $\rho^{-1}(r)$ can be constructed as $\rho^{-1}(r) = F_S^{-1}(F_R(r))$.

To describe the identification and estimation of $\chi^{-1}(\cdot)$, let $D_1^B = 0$ represent the buyer's decision to *not* walk away (and let $D_1^B = 1$ represent walking away) when informed that the high bid does not meet the reserve price, which occurs with the following conditional probability:

$$\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R}) = \int \frac{1 - F_B(\chi^{-1}(\tilde{p}^A - w))}{1 - F_B(\tilde{p}^A - w)} \left(\frac{M_\chi(\tilde{p}^A, w)}{\int M_\chi(\tilde{p}^A, z) dz} \right) dw \quad (6)$$

where $M_\chi(\tilde{p}^A, w) \equiv f_{P^A}(\tilde{p}^A - w)(1 - F_R(\tilde{p}^A - w))f_W(w)$ is the likelihood of the event $(P^A = \tilde{p}^A - w, \tilde{P}^A < \tilde{R}, W = w)$. Appendix C.5 demonstrates that $\chi^{-1}(\cdot)$ is nonparametrically identified. For estimation, I approximate $h_\chi(\cdot) \equiv 1 - F_B(\chi^{-1}(\cdot))$ as a flexible piecewise linear spline parameterized by θ^χ . Like the bounds on seller valuations, these parameters can be estimated using constrained least squares. I do so by evaluating the left-hand side and right-hand side of (6) on a fixed grid of points for the auction price \tilde{p}^A and search for the value of the parameter vector θ^χ that minimizes the distance between the left- and right-hand sides. This procedure requires estimates of densities and CDFs from above, as well an estimate of the conditional probability of not walking away, $\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R})$ which I estimate using a local linear regression. Technical details are found in Appendix C.5.

Step 6) Estimating the Direct Mechanism Corresponding to Real-World Bargaining.

Proposition 4 demonstrates that the allocation function corresponding to the real-world mechanism can be written as $x^{RW}(r, b; p^A) \equiv 1 \{b \geq g(r, p^A)\}$ for some unknown function $g(\cdot)$. The empirical object that can be used to identify this function $g(\cdot)$ is the probability of trade conditional on a realization of \tilde{R} and \tilde{P}^A . Let $\mathcal{A} \in \{0, 1\}$ be a random variable indicating whether or not trade occurs in a given instance of the game. The conditional probability of trade is given by

$$\Pr(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A) = \int \frac{1 - F_B(g(\tilde{r} - w, \tilde{p}^A - w))}{1 - F_B(\tilde{p}^A - w)} \left(\frac{M_g(\tilde{r}, \tilde{p}^A, w)}{\int M_g(\tilde{r}, \tilde{p}^A, z) dz} \right) dw \quad (7)$$

where $M_g(\tilde{r}, \tilde{p}^A, w) \equiv f_R(\tilde{r} - w) f_{P^A}(\tilde{p}^A - w) f_W(w)$ is the joint density of (R, P^A, W) . Appendix C.6 demonstrates that $g(\cdot)$ is nonparametrically identified. For estimation, I approximate $h_g(r, p_A) \equiv \frac{1 - F_B(g(r, p^A))}{1 - F_B(p^A)}$ using a flexible bilinear spline parameterized by θ_g . As with the estimation of the seller CDF bounds and the estimation of $\chi^{-1}(\cdot)$, I obtain an estimate of θ_g using constrained least squares. I do so by evaluating the left-hand side and right-hand side of (7) on a fixed grid of points and searching for the parameters θ_g to minimize the distance between the left- and right-hand sides. As with preceding steps, I estimate the conditional probability $\Pr(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A)$ in a first step; for this two-dimensional conditional probability, I use a tensor product of cubic b-spline functions. The other objects in (7) consist of densities and CDFs estimated above. Technical details are found in Appendix C.6.

Summary of Identification. Appendices C.2–C.6 provide nonparametric identification proofs for each of the objects I estimate. Here I provide a brief summary of the identification. The step 1 regression controlling for observable heterogeneity identifies the joint distribution of (\tilde{R}, \tilde{P}^A) (the residuals). This joint distribution identifies the marginal distributions F_R , F_{P^A} , and F_W , estimated in step 2. The underlying buyer distribution, F_B , estimated in step 3, is identified by the probability mass function, $\Pr(N = n)$, and by the marginal distribution of auction prices, F_{P^A} . Bounds on the distribution of seller valuations, estimated in step 4, are identified by F_R , f_{P^A} , and f_W , and by the conditional probabilities of sellers accepting and not quitting, $\Pr(D^S = A | \tilde{P}^A = \tilde{p}^A)$ and $\Pr(D^S \neq Q | \tilde{P}^A = \tilde{p}^A)$. The object $\chi^{-1}(\cdot)$, estimated in step 5, is identified by F_R , F_B , f^{P^A} , and f_W , and by the conditional probability of buyers not walking away from bargaining, $\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R})$. The object $\rho^{-1}(\cdot)$ is identified by F_R for any F_S lying in the seller CDF bounds. Finally, the object $g(\cdot)$, estimated in step 6, is identified by F_B , f_R , f_{P^A} , and f_W , and by the conditional probability of trade, $\Pr(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A)$.

5 Computing Bargaining Efficiency

To evaluate efficiency, I consider several welfare measures and compute these measures for the real-world bargaining and for ex-ante and ex-post efficient mechanisms. Each welfare measure depends on the estimated densities, CDFs, and lower support functions ($\underline{b}(\cdot) \equiv \chi^{-1}(\cdot)$ and $\underline{s}(\cdot) \equiv \rho^{-1}(\cdot)$) obtained in the estimation steps 2–5 above. Each welfare measure will also depend on an allocation function, x . As discussed in Section 3.3, the allocation function corresponding to the real-world bargaining is x^{RW} (estimated in step 6); the ex-post efficient allocation function is $x^*(s, b; p^A) \equiv 1\{s \leq b\}$; and the ex-ante efficient frontier consists of allocation functions, which I denote x^η , that place welfare weight of η on the seller’s expected utility and $1 - \eta$ on the buyer’s. Several points η along the ex-ante efficient frontier are of particular interest: x^1 , the allocation function corresponding to a take-it-or-leave-it offer by the seller; x^0 , the allocation function corresponding to a take-it-or-leave-it offer by the buyer; and $x^{\frac{1}{2}}$, the equal-weighted ex-ante efficient mechanism. In discussing results below I will refer to $x^{\frac{1}{2}}$ as *the second-best mechanism*, x^1 as the *seller-optimal mechanism*, and x^0 as the *buyer-optimal mechanism*.

The first welfare measure I consider is the overall expected gains from trade. For a given allocation function x and densities f_S and f_B , this is given by

$$\int_{\underline{b}}^{\bar{b}} \left[\int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} (b - s)x(s, b; p^A) f_S(s|p^A) f_B(b|p^A) ds db \right] f_{p^A}(p^A) dp^A \quad (8)$$

where $f_S(s|p^A) = \frac{f_S(s)}{1 - F_S(\underline{s}(p^A))}$ and $f_B(b|p^A) = \frac{f_B(b)}{1 - F_B(\underline{b}(p^A))}$ are the Bayes-updated beliefs of agents about their opponents’ types when bargaining starts. This welfare measure, along with the others I consider, is integrated over realizations of the lower bound of the support of buyer and seller types (i.e. integrated over the realized auction price p^A). I also evaluate several other welfare measures that are related to (8): the buyer’s or seller’s gains from trade (constructed by replacing $b - s$ in (8) with just b or s) and the probability of trade (constructed by replacing $b - s$ with 1).

The efficiency loss due to incomplete information—the loss highlighted in Myerson and Satterthwaite (1983)—can be estimated by evaluating the gains from trade using (8) for the ex-post efficient allocation function (i.e. replacing $x(s, b; p^A)$ in (8) with $x^*(s, b; p^A)$) and comparing this to (8) evaluated using the second-best allocation function ($x^{\frac{1}{2}}$). The efficiency loss due to other sources beyond those highlighted in Myerson and Satterthwaite (1983) can be estimated by comparing this second-best efficient outcome to (8) evaluated at the real-world mechanism (x^{RW}).²⁰

Evaluating the efficiency of bargaining at a given allocation function x is therefore easy once the densities, CDFs, and lower support of the bargaining types are known; it simply involves numerically evaluating integrals like (8). Computing the ex-ante efficient allocation functions themselves (x^η),

²⁰For the expected gains from trade in the real bargaining, I also incorporate an upper bound on the amount of bargaining costs incurred in the real-world mechanism (see Appendix C.7.3).

however, is extremely computationally involved in my setting. This is because it must be done at each realization of the lower bound of the support (each p^A) and because the type distributions I estimate are *irregular*, in the sense of Myerson (1981). F_B and F_S are referred to as regular if $b - \frac{1-F_B(b)}{f_B(b)}$ and $s + \frac{F_S(s)}{f_S(s)}$, the *virtual valuations* of buyers and sellers, are increasing. Myerson and Satterthwaite (1983) and Williams (1987) derived convenient solutions for the mechanisms along the ex-ante efficient frontier under the assumption of regularity. Without regularity, I am forced to numerically enforce a large number of incentive compatibility constraints (see Appendix C.7). Furthermore, although I have point estimates of F_B , I only have bounds on F_S , and, without further theory, obtaining bounds on welfare measures would require evaluating each mechanism at all possible distributions F_S within the estimated bounds. Fortunately, I am able to derive a number of useful monotonicity results that simplify this computation entirely. These results are summarized in Table 3 and in the following proposition:

Proposition 6. *A first order stochastically dominating change in F_S will lead to the monotonic changes in welfare measures described in Table 3.*

Each cell in Table 3 marked with a \downarrow signifies that the specified measure will decrease given a first-order stochastically dominating shift in the distribution of seller valuations (e.g. a shift from F_S^U to F_S^S). Each cell marked with an asterisk indicates that I have no proven monotonicity result for that welfare measure. For these latter welfare measures, I am forced to obtain bounds numerically through a massive grid search. This still yields valid bounds on welfare, but is very computationally expensive.²¹ Appendix C.7 contains technical details on this numerical procedure and the procedure for computing the ex-ante efficient allocation functions, x^n , which builds on results from Myerson and Satterthwaite (1983) and Williams (1987).

6 Putting It All Together: How Efficient Is Bargaining?

6.1 Distribution Estimates

This section presents the distributions of buyer and seller valuations estimated using the procedures described in Section 4. In each figure that follows, monetary values are denoted in units of \$1,000. In Figure 1, panels A and B display, for the dealers and fleet/lease samples respectively, the distribution of the auction price net of unobserved heterogeneity, F_{PA} (the dashed line), and the estimated underlying distribution of buyer valuations, F_B . The distribution of auction prices does not entirely dominate that of the underlying buyer valuations in a first order stochastic dominance

²¹To give a rough idea of the computational burden, even with the techniques I introduce here to reduce this burden, it takes about one year of computation time for a single machine to compute all of the estimates and confidence intervals reported in the paper. I parallelize these computations on a high-performance computing cluster to reduce this time to less than one week.

sense. This is due to the distribution of the number of bidders, $\Pr(N = n)$, having much of its mass at two or three bidders.

Panels C and D of Figure 1 show the distribution of secret reserve prices net of unobserved heterogeneity, F_R (the dashed line), and the estimated upper and lower bounds on the distribution of seller valuations, F_S^L and F_S^U , in solid lines. These bounds suggest that, for dealer cars (panel C), when the first bargaining offer (the auction price) is about -\$1,000 (i.e. \$1,000 lower than would be predicted based on car-level heterogeneity), sellers choose to accept this offer or walk away from it with frequencies that imply that the probability that S is less than -\$1,000 is in the range $[0.56, 0.80]$. For the fleet/lease sample (panel D) the corresponding probability inferred from sellers accepting or walking away from an offer of this same magnitude is in the range $[0.61, 0.80]$.

Comparing the top panels of Figure 1 to the bottom panels, it is clear that there is overlap in the support of buyer valuations and seller valuations. This feature illustrates what is referred to in the theoretical bargaining literature as the “no gap” case (i.e. there is no gap between the upper bound of the support of seller valuations and the lower bound of the support of buyer valuations, and hence there is uncertainty as to whether gains from trade actually exist), and is the case motivating Myerson and Satterthwaite (1983) (see Fudenberg and Tirole 1991). However, the actual overlap in terms of mass appears to be small, as most seller values (at least 80% in each sample) lie below zero—in some cases, far below zero—while buyer values are centered about zero and are much less dispersed. This implies that the actual efficiency loss due to incomplete information may be small in this setting. A more precise quantitative analysis of the overlap in buyer and seller mass will be discussed below, taking into account the support of the types who actually end up in the bargaining game.

6.2 Graphical Analysis of Bargaining Efficiency

Using the approach described in Section 5, I compute buyer gains and seller gains in the real-world bargaining mechanism as well as the ex-post and ex-ante efficient frontiers. The performance of the real-world bargaining relative to these theoretical frontiers is displayed in Figure 2. In panels A–D, the dashed line displays the ex-post efficient frontier in the space of buyer gains (the vertical axis) and seller gains (the horizontal axis). The solid line displays the ex-ante efficient frontier. The solid dot indicates the expected gains in the real-world mechanism. Panels A and B use the seller CDF lower bound and panels C and D use the upper bound.

Comparing the ex-ante efficient frontier to the ex-post efficient frontier provides an indication of the size of efficiency loss due strictly to incomplete information. In each case in Figure 2, the ex-ante efficient frontier lies close to the ex-post efficient frontier. This suggests that, in this market, incomplete information per se may not be leading to large inefficiencies, likely due to the limited overlap in buyer and seller distributions suggested by Figure 1. Comparing the real-world

bargaining outcome to the ex-ante efficient frontier, on the other hand, provides an indication of the size of efficiency loss due to the fact that the mechanism used in practice (alternating-offer bargaining) or its equilibrium may be inefficient. Panels A and B demonstrate that the real-world mechanism lies not far from the ex-ante efficient frontier when evaluated at the seller valuation CDF lower bound, but panels C and D demonstrate a clear shortfall in the efficiency of the real-world mechanism when evaluated at the seller valuation CDF upper bound, where the outcome lies in the interior of the ex-ante efficient frontier.

Panels E and F of Figure 2 display the probability of trade along the ex-ante efficient frontier evaluated at the seller valuation CDF upper and lower bounds. Each point on the horizontal axis represents the weight $\eta \in [0, 1]$ given to the seller’s valuation in evaluating ex-ante efficiency. The dashed line represents the probability of trade in the real bargaining (conditional on bargaining occurring). These figures demonstrate that the probability of trade is higher when η is closer to 0.5, and decreases slightly as η approaches 1. The probability of trade decreases dramatically as the seller’s welfare weight goes to 0, and in this range of η the real-world bargaining outperforms the ex-ante efficient mechanisms in terms of the probability of trade.

6.3 Quantitative Analysis of Bargaining Efficiency

The graphical analysis in the preceding section does not capture bounds on the difference in welfare between the different mechanisms; it only evaluates these mechanisms at the upper and lower bounds on seller valuations. This section presents a quantitative analysis of the bounds on welfare measures and the bounds on differences in welfare described in Section 5. Tables 4 and 5 contain numerical values for each of these bounds. Panel A displays the expected gains from trade, buyer and seller gains, and probability of trade for the ex-post efficient mechanism, the second-best mechanism, the buyer-optimal mechanism, the seller-optimal mechanism, and the real-world bargaining. Panel B displays the expected gains and probability of trade for the difference between ex-post efficiency and the second-best, the second-best and the real-world mechanism, and ex-post efficiency and the real-world mechanism. Gains are reported in units of \$1,000. I do not include the auction house revenue in the total expected gains from trade in these tables, but I do in my analysis of auction house fees in Appendix D. Tables 4 displays results using the dealers sample and Table 5 displays results using the fleet/lease sample. The estimated bounds are reported in square brackets and confidence sets are in parentheses. These confidence sets are constructed using a nonparametric bootstrap of the full estimation procedure (steps 1–6) and the counterfactual computations, thus accounting for uncertainty in the estimation procedure and counterfactuals. Appendix C.8 describes the confidence sets and proves that they are conservative.

Efficiency and Trade Volume in Dealer Sales. I begin by discussing the estimates in panel A of Table 4. The first column demonstrates that, in a full-information world, where ex-post

efficiency would be achievable, the gain from trade for the bargaining game would lie in a range \$2,442–5,045. Taking the lower bound (\$2,442) for illustrative purposes, this number suggests that (for example) if a car sells for \$7,000, the seller might have been willing to sell the car for as low as \$5,558 and the buyer willing to buy the car for as much as \$8,000 (because $8,000 - 5,558 = 2,442$). The probability of trade under ex-post efficiency mechanism ranges from 0.818–0.871. This latter quantity (or rather, one minus this quantity) serves as a direct measure of the amount (in mass) of overlap between the buyer and seller valuations.

The second column of panel A displays the second-best mechanism—the direct revelation mechanism maximizing the equally weighted expected gains subject to information constraints. I find that the range of surplus for this mechanism is only slightly below that of ex-post efficiency, suggesting that there is very little loss due solely to incomplete information in this setting. Moving to panel B of Table 4, the results in the first column confirm this finding, where I display bounds on the difference between the ex-post and second-best gains from trade and probability of trade. These bounds indicate that the second-best gains from trade fall below ex-post efficiency by \$9–59. Interestingly, however, the probability of trade in the second-best mechanism can be substantially lower than under ex-post efficiency (a lower bound of 0.699 as opposed to 0.818 in panel A, and an upper bound on the ex-post efficient probability of trade minus the second-best of 0.128 in panel B). Thus the second-best mechanism can miss out on trades that would be ex-post efficient (i.e. cases where the buyer values the car more than the seller), but these missed trades appear to be cases where the difference in valuations is small, and hence the surplus level is still close to ex-post efficient.

The final column of panel A indicates that the expected gains from trade in the real-world mechanism range from \$1,993 to \$3,933, with the buyer’s expected gains lying in \$822–845 and the seller’s lying in \$1,171–3,088. Relative to the real-world bargaining outcomes, column 3 of panel B indicates that ex-post efficiency would entail an increase in expected surplus of \$422–1,139 per bargaining transaction. This surplus lost in the real bargaining represents 17–23% of the ex-post efficient surplus. This lost surplus is a deadweight loss, uncaptured by either party (or by the auction house).

Efficient bargaining would also yield a higher conversion rate. Table 4 shows that the probability of trade in the real-world bargaining is 0.646, meaning trade fails 0.354 of the time. The final column of Panel B demonstrates that probability of trade would increase by 0.172–0.225 under ex-post efficiency. This implies that 17.2–22.5% of negotiating pairs consist of cases where the buyer values the car more than the seller but trade fails. Comparing this to the overall failure rate (35.4%) suggests that approximately half (48.5–63.6%) of failed negotiations are cases where gains from trade do exist but fail to be realized.

Panel B of Table 4 also demonstrates that the probability of trade would increase by up to 0.182 when moving from real-world bargaining to the second-best mechanism. Note that the lower bound

on the improvement in the probability of trade is much lower in the second-best mechanism (0.052), implying that I cannot reject the possibility that the real bargaining achieves only a slightly lower trade volume than the second-best mechanism. This again highlights the feature that the second-best mechanism would guarantee an increase in expected surplus relative to the real bargaining by capturing higher-value trades, not necessarily a higher volume of trade. However, industry participants suggest that it is high conversion—a high probability of trade—that is the primary goal of wholesale auto auction houses (see Treece 2013; Lacetera, Larsen, Pope, and Sydnor 2016), and thus the real-world mechanism may be achieving this goal relatively well.

Efficiency and Trade Volume in Fleet/lease Sales. The findings for the fleet/lease sample are shown in Table 5 and display similar patterns. As in the dealers sample, the second-best mechanism achieves a surplus level that is similar to ex-post efficiency in panel A, and panel B indicates that an upper bound on the gap between second-best and ex-post efficient gains from trade is \$77. The real-world bargaining falls short of the ex-post efficient gains by \$289–864, a deadweight loss of about 12–20% of the ex-post efficient surplus. Panel A indicates that the probability of trade in the second-best mechanism is strictly lower than under ex-post efficiency, although still higher than in the real-world bargaining. This is confirmed in panel B, where the gap in probability of trade between the second-best and real-world bargaining ranges from 0.045–0.174, but between ex-post efficiency and real-bargaining the gap ranges from 0.199–0.235. Trade in the bargaining stage of the game occurs with probability 0.658 in the fleet/lease sample; the 0.199–0.235 improvement under ex-post efficiency shown in panel B suggests that 58.2–68.7% of failed trades are cases where the buyer values the car more than the seller but trade fails to occur.

The estimated probabilities of trade for the real bargaining (0.646 in Table 4 and 0.658 in Table 5) come from integrating the estimated allocation function $x^{RW}(\rho(S), B; P^A)$ over all three of its arguments, as described in Section 5. The corresponding raw probability of trade in the data for the bargaining stage can be calculated by combining periods 2 and higher from Table 2 (and from Appendix Table A4). For the dealers sample this number is 0.684 and for the fleet/lease sample this number is 0.656, each of which is close to the estimated probability of trade from Tables 4–5 for the real bargaining. This comparison is one indication of the good fit of the overall estimation exercise; other measures are discussed in Appendix C, where I show that the integrated absolute error for each estimation step is small.

Overall, the results are quite similar in the dealers and fleet/lease samples. This suggests that, while the observable characteristics of these cars differ (as shown in Table 1), the residual variation after controlling for these observables—and the players’ behavior in the game—does not differ drastically between the two samples. As discussed at the end of Section 4, identification of the key objects of the paper relies on several key objects in the data: the joint distribution of (\tilde{R}, \tilde{P}^A) , $\Pr(N = n)$, and several conditional probability statements. Appendix Figures A3 and A5

display estimates of many of these objects for the two samples, illustrating that these drivers of identification do not differ much between the two samples.

Buyer-Optimal and Seller-Optimal Mechanisms. The third and fourth columns of panel A in Tables 4–5 display bounds on welfare outcomes under the buyer-optimal and seller-optimal mechanisms. These mechanisms also lie along the ex-ante efficient frontier, but place all of the welfare weight on one party or the other. One interesting feature of these mechanisms is that they are easy to implement; they simply require letting one party make a take-it-or-leave-it offer to the opposing party.

I find in Table 4 that the buyer-optimal mechanism would yield a much higher payoff for the buyer (\$1,416–3,962) and much lower payoff for the seller (\$360–370) than under the current mechanism. The probability of trade, however has the potential to drop as low as 0.266 under the buyer-optimal mechanism (with an upper bound of 0.609). Some of these changes are due to the fact that in this buyer-optimal bargaining the buyer is no longer forced to treat the auction price as a lower bound on the available bargaining prices.

The seller-optimal mechanism would potentially yield improvements for the seller, with the seller’s expected gains from trade lying in a range from \$1,870–4,349, and the buyer’s gains in this mechanism dropping to \$474–649. The probability of trade under the seller-optimal mechanism can be as low as 0.634, not nearly as low as in the buyer-optimal mechanism, and close to that of the real-world bargaining. The bounds on the total expected gains from trade in the seller-optimal mechanism are similar to those in the second-best mechanism (\$2,344–4,999). Table 5 displays similar results for the fleet/lease sample. Comparing the seller-optimal and real bargaining columns suggests that the probability of trade in the fleet/lease sample would potentially be lower under the seller-optimal mechanism ([0.592, .747]) than under the real-world mechanism (0.658).

This highlights an interesting distinction between a *secret* reserve price and a *public* reserve price. In an independent private values environment, a public reserve auction is equivalent to an auction followed by the seller-optimal bargaining mechanism (Menezes and Ryan 2005) and is optimal for the *seller* in most standard auction environments. A secret reserve auction, however, may be preferred by the *intermediating platform*—in this case, the auction house—because it may yield a higher probability of trade. In particular, if industry fees are such that the auction house is only paid when trade occurs, as is the case in the wholesale auto industry, then an auction house would prefer to maximize trade volume rather than seller revenue. The combination of secret reserve prices and bargaining in this setting (and potentially in other settings as well) may therefore be motivated by a goal to maximize the volume of trade rather than seller revenue (a related point is made in Elyakime, Laffont, Loisel, and Vuong 1994). Appendix B.6 contains further discussion of secret vs. public reserve prices.

6.4 Bargaining Between a Random Buyer and Seller

I now analyze the estimated expected gains from trade in a setting where bargaining takes place between the seller and a random buyer rather than the high-value bidder. Note that this alters the support of the types in the bargaining game for both the buyer *and* the seller, and this shifts not only the outcomes achieved in the real-world bargaining, but also the ex-ante and ex-post efficient frontiers. The allocation function for the real-world bargaining in this setting is possible to simulate by evaluating the estimated $g(r, p^A)$ function at a very small realized value of p^A (I choose the 0.001 quantile). At this small realization, the valuation of the buyer who bargains approximately represents a draw from the full support of buyer valuations, $[\underline{b}, \bar{b}]$ rather than from the truncated support used in the main analysis above, $[\underline{b}(p^A), \bar{b}]$. Similarly, the valuation of the seller who bargains is a draw from the full support of seller valuations, $[\underline{s}, \bar{s}]$ rather than from $[\underline{s}(p^A), \bar{s}]$.

Table 6 displays the results. In the dealers sample, the range for the gains from trade under ex-post efficiency is slightly tighter than in the main results in Table 4, and the estimated bounds on the gains in the real-world bargaining lie closer to zero than in the main results. The overall loss in efficiency between the real bargaining and ex-post efficiency is \$976–2,204 in the dealers sample, which corresponds to a loss of 39–45% of the ex-post efficient surplus. This loss is much larger than the percentage loss in the main results in Table 4, suggesting that bargaining between a random buyer and seller is more inefficient than bargaining between the high bidder and the seller. For the fleet/lease sample, the overall efficiency loss of the real bargaining is \$1,525–2,765, corresponding to a percentage loss of about 49–56% of the ex-post efficient surplus, a range in losses that is much larger than in the main results in Table 5. The results of this exercise suggest that the presence of the auction (through its roles of truncating the support of the types who arrive at bargaining and constraining from below the level of the final offer) does indeed improve the efficiency of the bargaining. The main findings of the paper still hold in this analysis: the real-world bargaining is inefficient, and nearly all of this inefficiency is due to factors other than the information constraints highlighted in Myerson and Satterthwaite (1983), as the second-best mechanism yields outcomes that are nearly ex-post efficient.

7 Discussion and Conclusion

The finding of this paper that the ex-ante and ex-post efficient frontiers are close to one another in this market stands in stark contrast to the result in the most popularly studied theory example of bilateral bargaining, that of symmetric uniform values (where both buyer and seller valuations are uniformly distributed on the interval $[0, 1]$; see, for example, Chatterjee and Samuelson 1983 and Myerson and Satterthwaite 1983). This case is known to yield a gap between the ex-post efficient

and second-best probability of trade. The large gap in this special case, however, may have little bearing on the gap to be expected in real-world settings, where the features of the distribution and the extent of asymmetries may diverge far from uniformity and symmetry. Also, as the results above highlight, even in situations where some efficient trades fail to occur, many of these failed trades may be cases where only very small gains from trade exist (i.e. where the buyer’s value is very close to the seller’s), and thus the loss in efficiency due to information constraints need not be large.

Overall, it is not obvious whether the results of this paper should be interpreted as implying that the real-world bargaining in this market is relatively efficient or relatively inefficient compared to other markets, particularly given that there are no existing empirical studies of bargaining with two-sided uncertainty to which these results may be compared. Estimating a structural model of one-sided uncertainty, Ambrus, Chaney, and Salitsky (2018) found an efficiency loss of 14% in studying ransom negotiations, and the losses I find are similar to these (17–23% for dealer cars and 12–20% for fleet/lease cars). Several papers in the experimental literature can also provide an interesting comparison. Bazerman, Gibbons, Thompson, and Valley (1998) argued that real-world bargaining can potentially yield *more* efficient outcomes than the theoretical ex-ante efficient frontier due to non-traditional utility functions (where one player’s utility nests the other’s), limits on players’ abilities to mimic other types, and other features of bounded rationality; and Valley, Thompson, Gibbons, and Bazerman (2002) found evidence in lab experiments that communication between players can allow them to outperform the ex-ante efficient frontier. In light of these arguments, the bargaining at wholesale auto auctions might be seen as relatively *inefficient* given that it falls short of that frontier at all.²²

As discussed in the introduction, a gap between the outcome of actual bargaining and the ex-ante efficient frontier can occur for a number of reasons. First, real-world bargaining mechanisms can have multiple equilibria, many of which may be inefficient (see Ausubel and Deneckere 1993), and the actions I observe in the data may correspond to one of these inefficient equilibria. Second, it may be that this particular bargaining protocol, even in its most efficient equilibria, falls short of the frontier.²³ Third, it may be the case that a gap exists because of a Wilson-doctrine-like argument: ex-ante efficient mechanisms can be unwieldy to implement in practice (in particular when $\eta \in (0, 1)$). These mechanisms require that players and the mechanism designer all have

²²Importantly, nothing in my estimation procedure forced the real-world outcome to lie within the ex-ante efficient frontier; this is a finding implied by my estimates, not a constraint placed on the model.

²³Ausubel and Deneckere (1993) explain this potential efficiency loss as follows: “In [the second-best mechanism], each player reveals his private information before hearing his opponent’s report. By way of contrast, in sequential bargaining, the player who reveals second may be less apt to report truthfully than if he were still ignorant of his opponent’s report. To the extent that information revelation is inhibited, this might further contribute to waste.” Ausubel and Deneckere (1993) demonstrated that this efficiency loss need not occur in sequential bargaining with one-sided incomplete information when the valuation distribution has a monotone hazard rate (unlike the distributions I estimate). Ausubel, Cramton, and Deneckere (2002) argued that these results may be extended to alternating-offer games, but no general exposition exists.

knowledge of buyer and seller distributions, and furthermore that the players comprehend that it is indeed incentive compatible for them to truthfully reveal their valuations. The implementation of alternating-offer bargaining, on the other hand, does not require such assumptions and the rules can be easily explained to both the players and the market designer, unlike the black box that theoretical mechanisms may appear to be from a player’s perspective. It may indeed be the case, as hypothesized by Wilson (1986) and Ausubel and Deneckere (1993), that “[real-world bargaining mechanisms] survive because they employ trading rules that are efficient for a wide class of environments.”²⁴

Actual quantitative estimates of real-world bargaining efficiency from other studies will be a welcome addition to the literature in the future for comparison to the estimates in this paper. A fruitful avenue for future empirical research would be to apply the bounding methodology developed in this paper to study the efficiency of bargaining in other settings, potentially exploiting more fully all of the offers observed in alternating-offer bargaining data—a form of data that is becoming increasingly available (e.g. Merlo and Ortalo-Magne 2004; Keniston 2011; Bagwell, Staiger, and Yurukoglu 2017; Backus, Blake, Larsen, and Tadelis 2018; Hernandez-Arenaz and Iriberry 2018).

My consistent finding is that the ex-ante and ex-post efficient frontiers lie close together in this market, while the real-world bargaining falls short of the ex-ante efficient frontier. This suggests that efficiency loss in this market may not be due to incomplete information alone, but to the other aspects of the real-world bargaining described above. It is important to note, however, that these other aspects all have roots in incomplete information; if players were to have complete information, many of these other barriers to efficiency might also disappear.

²⁴An interesting avenue for future theory research would be to apply recent robust mechanism design techniques (surveyed in Carroll 2018) to the analysis of alternating-offer bargaining to determine whether there is a designer’s objective function or information structure under which alternating-offer bargaining is preferable to other mechanisms for bilateral trade.

References

- AMBRUS, A., E. CHANEY, AND I. SALITSKY (2018): “Pirates of the Mediterranean: An empirical investigation of bargaining with transaction costs,” *Quantitative Economics*, 9, 217–246.
- AN, Y., AND X. TANG (2018): “Specification and negotiation in incomplete contracts: An econometric analysis,” *RAND Journal of Economics*, forthcoming.
- ARADILLAS-LÓPEZ, A., A. GANDHI, AND D. QUINT (2016): “A simple test for moment inequality models with an application to English auctions,” *Journal of Econometrics*, 194, 96–115.
- ASKER, J. (2010): “A study of the internal organization of a bidding cartel,” *American Economic Review*, 100(3), 724–762.
- ATHEY, S., AND P. HAILE (2007): “Nonparametric approaches to auctions,” *Handbook of Econometrics*, 6, 3847–3965.
- AUSUBEL, L., P. CRAMTON, AND R. DENECKERE (2002): “Bargaining with incomplete information,” *Handbook of Game Theory*, 3, 1897–1945.
- AUSUBEL, L., AND R. DENECKERE (1993): “Efficient sequential bargaining,” *Review of Economic Studies*, 60(2), 435.
- BACKUS, M., T. BLAKE, B. LARSEN, AND S. TADELIS (2018): “Sequential bargaining in the field: Evidence from millions of online bargaining interactions,” NBER Working Paper 24306.
- BAGWELL, K., R. W. STAIGER, AND A. YURUKOGLU (2017): “Multilateral trade bargaining: A first peek at the GATT bargaining records,” Working paper, Stanford University.
- BAZERMAN, M., R. GIBBONS, L. THOMPSON, AND K. VALLEY (1998): “Can negotiators outperform game theory?,” in *Debating Rationality: Nonrational Aspects of Organizational Decision Making*, ed. by J. Halpern, and R. Stern, pp. 79–99. Ithaca, NY: Cornell University Press.
- CARROLL, G. (2018): “Robustness in mechanism design and contracting,” *Annual Reviews*, forthcoming.
- CHATTERJEE, K., AND W. SAMUELSON (1983): “Bargaining under incomplete information,” *Operations Research*, 31(5), 835–851.
- CRAMTON, P. (1991): “Dynamic bargaining with transaction costs,” *Management Science*, 37(10), 1221–1233.
- CRAWFORD, G. S., AND A. YURUKOGLU (2012): “The welfare effects of bundling in multichannel television markets,” *American Economic Review*, 102(2), 643–685.

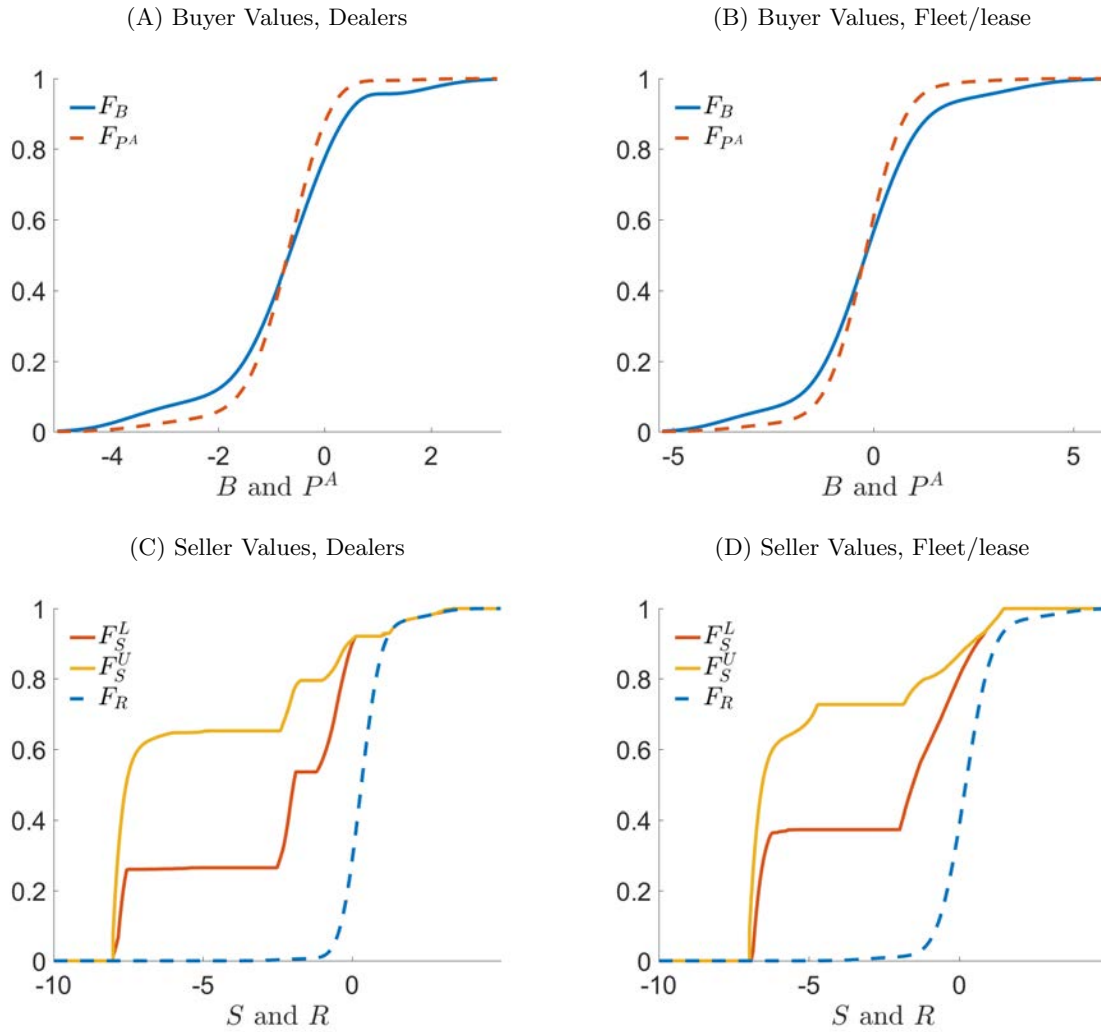
- DECAROLIS, F. (2018): “Comparing public procurement auctions,” *International Economic Review*, 59(2), 391–419.
- EDLIN, A., AND C. SHANNON (1998): “Strict monotonicity in comparative statics,” *Journal of Economic Theory*, 81(1), 201–219.
- ELYAKIME, B., J. J. LAFFONT, P. LOISEL, AND Q. VUONG (1994): “First-price sealed-bid auctions with secret reservation prices,” *Annales d’Economie et de Statistique*, 34, 115–141.
- (1997): “Auctioning and bargaining: An econometric study of timber auctions with secret reservation prices,” *Journal of Business & Economic Statistics*, 15(2), 209–220.
- EVDOKIMOV, K., AND H. WHITE (2012): “Some extensions of a lemma of Kotlarski,” *Econometric Theory*, 28(4), 925–932.
- FREYBERGER, J., AND B. LARSEN (2017): “Identification in ascending auctions, with an application to digital rights management,” NBER Working Paper 23569.
- FUDENBERG, D., D. LEVINE, AND J. TIROLE (1985): “Infinite-horizon models of bargaining with one-sided incomplete information,” in *Game Theoretic Models of Bargaining*, ed. by A. Roth, pp. 73–98. Cambridge: Cambridge University Press.
- FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. Cambridge: MIT Press.
- GALLANT, A., AND D. NYCHKA (1987): “Semi-nonparametric maximum likelihood estimation,” *Econometrica*, 55(2), 363–390.
- GENESOVE, D. (1991): “Coconuts, lemons and pies: Adverse selection, search and bargaining in the wholesale used car market,” Ph.D. thesis, Princeton University.
- GRENNAN, M., AND A. SWANSON (2019): “Transparency and negotiated prices: The value of information in hospital-supplier bargaining,” *Journal of Political Economy*, forthcoming.
- GUL, F., AND H. SONNENSCHNEIN (1988): “On delay in bargaining with one-sided uncertainty,” *Econometrica*, 56(3), 601–611.
- HAILE, P., H. HONG, AND M. SHUM (2003): “Nonparametric tests for common values at first-price sealed-bid auctions,” Working paper, Yale.
- HAILE, P., AND E. TAMER (2003): “Inference with an incomplete model of English auctions,” *Journal of Political Economy*, 111(1), 1–51.
- HERNANDEZ-ARENAZ, I., AND N. IRIBERRI (2018): “Women ask for less (only from men): Evidence from bargaining in the field,” *Journal of Economic Behavior & Organization*, 152, 192–214.

- HOLMSTRÖM, B., AND R. B. MYERSON (1983): “Efficient and durable decision rules with incomplete information,” *Econometrica*, 51(6), 1799–1819.
- HUH, W., AND K. PARK (2010): “A sequential auction-bargaining procurement model,” *Naval Research Logistics (NRL)*, 57(1), 13–32.
- JUDD, K. (1998): *Numerical Methods in Economics*. Cambridge: The MIT Press.
- KENISTON, D. (2011): “Bargaining and welfare: A dynamic structural analysis,” Working Paper, Yale.
- KIM, K. I., AND J. LEE (2014): “Nonparametric estimation and testing of the symmetric IPV framework with unknown number of bidders,” Working paper, University of Minnesota.
- KOTLARSKI, I. (1967): “On characterizing the gamma and normal distribution,” *Pacific Journal of Mathematics*, 20(1), 69–76.
- KRASNOKUTSKAYA, E. (2011): “Identification and estimation of auction models with unobserved heterogeneity,” *Review of Economic Studies*, 78(1), 293–327.
- LACETERA, N., B. J. LARSEN, D. G. POPE, AND J. R. SYDNOR (2016): “Bid takers or market makers? The effect of auctioneers on auction outcomes,” *American Economic Journal: Microeconomics*, 8(4), 195–229.
- LANG, S. (2011): “Hammer time: The chosen,” *The Truth About Cars*, Url: <http://www.thetruthaboutcars.com/2011/09/hammer-time-the-chosen>. Posted Sep. 24, 2011. Retrieved Dec. 21, 2011.
- LARSEN, B., AND A. ZHANG (2018): “A mechanism design approach to identification and estimation,” NBER Working Paper 24837.
- LI, H., AND N. LIU (2015): “Nonparametric identification and estimation of double auctions with bargaining,” Working paper, Penn State.
- LI, T., I. PERRIGNE, AND Q. VUONG (2000): “Conditionally independent private information in OCS wildcat auctions,” *Journal of Econometrics*, 98(1), 129–161.
- LI, T., AND Q. VUONG (1998): “Nonparametric estimation of the measurement error model using multiple indicators,” *Journal of Multivariate Analysis*, 65(2), 139–165.
- MENEZES, F., AND M. RYAN (2005): “Reserve price commitments in auctions,” *Economics Letters*, 87(1), 35–39.

- MERLO, A., AND F. ORTALO-MAGNE (2004): “Bargaining over residential real estate: Evidence from England,” *Journal of Urban Economics*, 56(2), 192–216.
- MERLO, A., F. ORTALO-MAGNÉ, AND J. RUST (2015): “The home selling problem: Theory and evidence,” *International Economic Review*, 56(2), 457–484.
- MERLO, A., AND X. TANG (2012): “Identification and estimation of stochastic bargaining models,” *Econometrica*, 80(4), 1563–1604.
- (2018): “Bargaining with optimism: Identification and estimation of a model of medical malpractice litigation,” *International Economic Review*, forthcoming.
- MYERSON, R. B. (1979): “Incentive compatibility and the bargaining problem,” *Econometrica*, 47, 61–73.
- (1981): “Optimal auction design,” *Mathematics of Operations Research*, 6(1), 58–73.
- MYERSON, R. B., AND M. A. SATTERTHWAITTE (1983): “Efficient mechanisms for bilateral trading,” *Journal of Economic Theory*, 29(2), 265–281.
- NAAA (2009): “2009 Annual Review,” *National Auto Auction Association*, www.naaa.com.
- PERRY, M. (1986): “An example of price formation in bilateral situations: A bargaining model with incomplete information,” *Econometrica*, 54(2), 313–321.
- ROBERTS, J. W. (2013): “Unobserved heterogeneity and reserve prices in auctions,” *RAND Journal of Economics*, 44(4), 712–732.
- RUBINSTEIN, A. (1982): “Perfect equilibrium in a bargaining model,” *Econometrica*, 50(1), 97–109.
- SATTERTHWAITTE, M. A., AND S. R. WILLIAMS (1989): “Bilateral trade with the sealed bid k-double auction: Existence and efficiency,” *Journal of Economic Theory*, 48(1), 107–133.
- SCOTT MORTON, F., J. SILVA-RISSO, AND F. ZETTELMEYER (2011): “What matters in a price negotiation: Evidence from the US auto retailing industry,” *Quantitative Marketing and Economics*, 9(4), 365–402.
- SIEG, H. (2000): “Estimating a bargaining model with asymmetric information: Evidence from medical malpractice disputes,” *Journal of Political Economy*, 108(5), 1006–1021.
- SILVEIRA, B. S. (2017): “Bargaining with asymmetric information: An empirical study of plea negotiations,” *Econometrica*, 85(2), 419–452.
- SONG, U. (2004): “Nonparametric estimation of an eBay auction model with an unknown number of bidders,” Working paper, University of British Columbia.

- STORMS, E. (2015): “Monotonicity of trade in PSBNE of bargaining with alternating offers,” Mimeo, Stanford University.
- TREECE, J. B. (2013): “Manheim exec: Why do so few sales at auctions close?,” *Automotive News*, Url: <http://www.autonews.com/article/20131202/RETAIL04/312029992/manheim-exec-why-do-so-few-sales-at-auctions-close>. Posted Dec. 2, 2013. Retrieved Dec. 2, 2013.
- VALLEY, K., L. THOMPSON, R. GIBBONS, AND M. H. BAZERMAN (2002): “How communication improves efficiency in bargaining games,” *Games and Economic Behavior*, 38(1), 127–155.
- WANG, R. (2000): “Bidding and renegotiation in procurement auctions,” *European Economic Review*, 44(8), 1577–1597.
- WATANABE, Y. (2009): “Learning and bargaining in dispute resolution: Theory and evidence from medical malpractice litigation,” Working Paper, Northwestern University.
- WILLIAMS, S. R. (1987): “Efficient performance in two agent bargaining,” *Journal of Economic Theory*, 41(1), 154–172.
- WILSON, R. (1986): “Game-theoretic analyses of trading processes,” in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, pp. 33–70. Cambridge: Cambridge University Press.

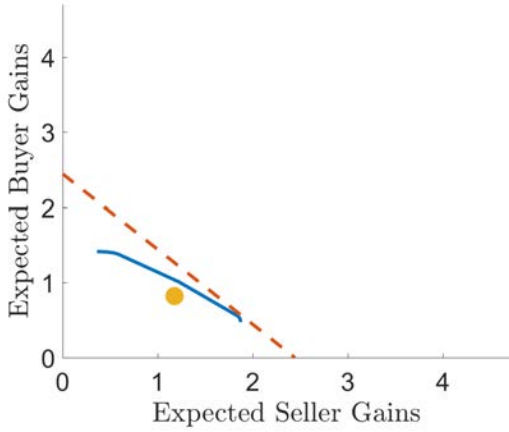
Figure 1: Distribution Estimates




Notes: Panels A and B display estimated distribution of auctions prices (dashed line) after removing observable and unobservable game-level heterogeneity, and estimated distribution of buyer valuations (solid line). Panels C and D display estimated distribution of reserve prices (dashed line) after removing observable and unobservable game-level heterogeneity, and estimated lower and upper bounds on distribution of seller valuations (solid lines). Panels on left use dealers sample and on right use fleet/lease sample. Units = \$1,000.

Figure 2: Graphical Evaluation of Bargaining Efficiency

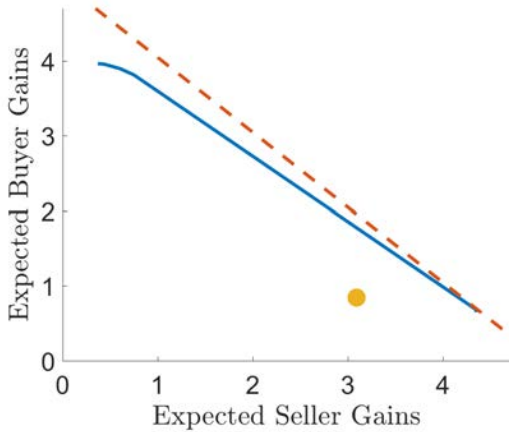
(A) Dealers, Gains Using Seller Lower Bound




(B) Fleet/lease, Gains Using Seller Lower Bound




(C) Dealers, Gains Using Seller Upper Bound



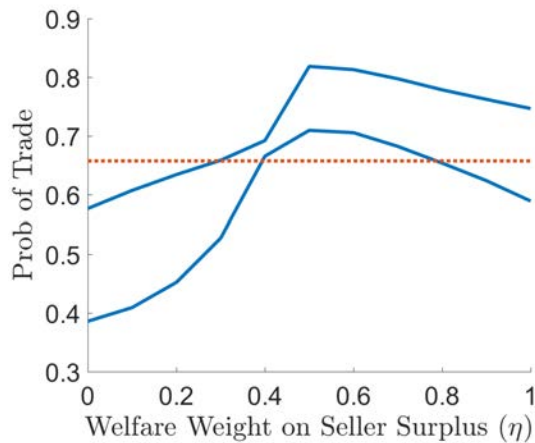
(D) Fleet/lease, Gains Using Seller Upper Bound



(E) Dealers, Prob of Trade



(F) Fleet/lease, Prob of Trade



Notes: Panels A and B display estimated expected seller and buyer gains on ex-post efficient frontier (dashed line), on ex-ante efficient frontier (solid line), and in real-world bargaining (solid dot) using the seller distribution lower bound. Panels C and D display the same estimates at the seller distribution upper bound. Panels E and F display, in solid lines, the estimated probability of trade at each η (seller's welfare weight) along the ex-ante efficient frontier evaluated at the seller valuation distribution upper and lower bounds. The dashed line in panels E and F displays the probability of trade in the real-world mechanism for comparison. Panels on left use dealers sample and on right use fleet/lease sample. Units = \$1,000.

Table 1: Descriptive Statistics

	Dealers		Fleet/lease	
A. Full Sample	Mean	S.D.	Mean	S.D.
Trade	0.705	0.456	0.768	0.422
Reserve price	\$7,405	\$5,196	\$10,307	\$5,789
Auction price	\$6,253	\$4,881	\$9,804	\$5,857
Auction price if \geq reserve	\$6,197	\$4,700	\$11,063	\$6,139
Auction price if $<$ reserve	\$6,258	\$4,899	\$9,160	\$5,600
Number of periods	2.096	0.681	1.774	0.690
Blue book	\$6,820	\$4,828	\$10,951	\$6,144
Age (years)	6.769	3.369	3.178	2.552
Odometer (miles)	97,938	46,445	57,481	40,389
Sample size	133,523		131,443	
B. Bid Log Sample	Mean	S.D.	Mean	S.D.
Number bidders lower bound	2.924	0.340	2.973	0.443
Sample size	13,150		102,172	

Notes: Mean and standard deviation of variables in dealers and fleet/lease samples. Trade is an indicator for whether trade occurred between the buyer and seller. Number of periods is 1 if game ends through auction price exceeding reserve price or through buyer opting out of bargaining, 2 if seller accepts at her first bargaining turn, etc. Blue book is an estimate of the market value of the car, provided by the auction house. Panel A displays full sample and panel B displays subsample containing bid log records.

Table 2: Outcomes of Game By Period: Dealers Sample

Ending period	Player's turn	# Obs	% of Sample	% Trade	<u>Full Sample</u>		<u>Conditional on Sale</u>		
					Reserve price	Auction price	Reserve price	Auction price	Final price
1	(Auction)	14,232	10.659%	88.58%	\$5,818 (\$4,688)	\$6,050 (\$4,696)	\$5,704 (\$4,626)	\$6,197 (\$4,700)	\$6,197 (\$4,700)
2	S	99,708	74.675%	77.28%	\$7,632 (\$5,245)	\$6,379 (\$4,936)	\$7,132 (\$4,922)	\$6,003 (\$4,651)	\$6,003 (\$4,651)
3	B	14,644	10.967%	11.81%	\$7,192 (\$5,041)	\$5,502 (\$4,598)	\$7,725 (\$5,222)	\$6,508 (\$4,968)	\$6,912 (\$5,054)
4	S	2,916	2.184%	65.47%	\$7,828 (\$5,168)	\$6,369 (\$4,833)	\$7,678 (\$5,022)	\$6,359 (\$4,733)	\$6,608 (\$4,798)
5	B	1,667	1.248%	38.15%	\$8,267 (\$5,283)	\$6,680 (\$4,933)	\$8,640 (\$5,469)	\$7,174 (\$5,173)	\$7,792 (\$5,317)
6	S	190	0.142%	75.79%	\$8,655 (\$5,286)	\$7,068 (\$4,988)	\$8,995 (\$5,460)	\$7,456 (\$5,171)	\$7,965 (\$5,298)
7	B	139	0.104%	54.68%	\$8,459 (\$5,260)	\$6,829 (\$4,963)	\$8,865 (\$5,394)	\$7,339 (\$5,153)	\$8,029 (\$5,315)
8	S	20	0.015%	70.00%	\$8,935 (\$4,933)	\$7,575 (\$4,888)	\$9,446 (\$5,669)	\$8,157 (\$5,614)	\$8,725 (\$5,692)
9	B	6	0.004%	66.67%	\$7,583 (\$4,362)	\$6,225 (\$4,517)	\$6,675 (\$2,871)	\$5,188 (\$2,964)	\$5,838 (\$2,800)
10	S	1	0.001%	100.00%	\$14,500	\$11,000	\$14,500	\$11,000	\$13,600

Notes: Dealers sample. For each period (period 1 = auction, period 2 = seller's first turn in bargaining, period 3 = buyer's turn, etc.), table reports the number of observations ending in that period, percent of total sample ending in that period, and percent of cases in which trade occurred. Table also reports reserve price and auction price for observations ending in a given period and, for those observations ending in trade, the reserve price, auction price, and final price conditional on trade. Corresponding statistics for the fleet/lease sample are found in Table A4.

Table 3: Monotonicity Results for Welfare Measures

A. Levels	Ex-post	Second-best	Buyer-optimal	Seller-optimal	Real bargaining
Expected gains from trade	↓	↓	–	–	↓
Buyer gains		–	↓	–	↓
Seller gains		–	–	↓	↓
Probability of trade	↓	*	*	↓	↓
B. Differences		Ex-post minus second-best	Second-best minus real		Ex-post minus real
Expected gains from trade		*	↓		↓
Probability of trade		*	*		↓

Notes: Table displays monotonicity results for welfare measures proved in Proposition 6. Each cell marked with a ↓ signifies that the specified measure will decrease given a first-order stochastically dominating shift in F_S . Each cell marked with an asterisk indicates that there is no mathematical proof of a monotonicity result and that the bounds must be determined numerically. Cases marked with “–” in the second-best column indicate that I will report bounds on these quantities corresponding to the F_S leading to the maximum and minimum bounds on the total expected gains from trade. Cases marked with “–” in the buyer-optimal column indicate that I will report bounds on these quantities corresponding to the F_S leading to the maximum and minimum bounds on the buyer gains from trade. Cases marked with “–” in the seller-optimal column indicate that I will report bounds on these quantities corresponding to the F_S leading to the maximum and minimum bounds on the seller gains from trade.

Table 4: Bounds on Welfare Measures, Dealers Sample

A. Levels	Ex-post	Second-best	Buyer-optimal	Seller-optimal	Real bargaining
Expected gains from trade	[2.442, 5.045] (2.184, 5.247)	[2.397, 5.029] (2.149, 5.227)	[1.776, 4.332] (1.473, 4.492)	[2.344, 4.999] (2.096, 5.207)	[1.993, 3.933] (1.812, 4.085)
Buyer gains		[0.553, 0.701] (0.521, 0.733)	[1.416, 3.962] (1.159, 4.102)	[0.474, 0.649] (0.457, 0.662)	[0.822, 0.845] (0.779, 0.913)
Seller gains		[1.844, 4.328] (1.578, 4.519)	[0.360, 0.370] (0.190, 0.441)	[1.870, 4.349] (1.612, 4.561)	[1.171, 3.088] (0.947, 3.202)
Probability of trade	[0.818, 0.871] (0.806, 0.894)	[0.699, 0.828] (0.672, 0.857)	[0.266, 0.609] (0.250, 0.637)	[0.634, 0.801] (0.627, 0.811)	[0.646, 0.646] (0.638, 0.672)
B. Differences		Ex-post minus second-best	Second-best minus real		Ex-post minus real
Expected gains from trade		[0.009, 0.059] (0.004, 0.066)	[0.377, 1.123] (0.318, 1.202)		[0.422, 1.139] (0.355, 1.219)
Probability of trade		[0.034, 0.128] (0.015, 0.142)	[0.052, 0.182] (0.026, 0.194)		[0.172, 0.225] (0.159, 0.236)

Notes: Dealers sample. Bounds on welfare measures under ex-post efficient, second-best, buyer-optimal, and seller-optimal mechanisms compared to real-world mechanism. Panel A displays levels and panel B displays differences. Estimated bounds are in square braces and 95% confidence set is in parentheses. Gains are in \$1,000 units.

Table 5: Bounds on Welfare Measures, Fleet/lease Sample

A. Levels	Ex-post	Second-best	Buyer-optimal	Seller-optimal	Real bargaining
Expected gains from trade	[2.408, 4.195] (2.168, 4.363)	[2.342, 4.165] (2.102, 4.337)	[1.762, 3.503] (1.682, 3.783)	[2.223, 4.088] (1.992, 4.253)	[2.080, 3.370] (1.843, 3.466)
Buyer gains		[0.867, 1.039] (0.797, 1.058)	[1.439, 2.953] (1.278, 3.080)	[0.664, 0.880] (0.628, 0.888)	[1.158, 1.192] (1.073, 1.219)
Seller gains		[1.475, 3.126] (1.281, 3.298)	[0.323, 0.550] (0.311, 0.720)	[1.558, 3.208] (1.350, 3.367)	[0.922, 2.178] (0.754, 2.302)
Probability of trade	[0.857, 0.893] (0.847, 0.898)	[0.703, 0.832] (0.694, 0.840)	[0.377, 0.660] (0.364, 0.678)	[0.589, 0.747] (0.551, 0.752)	[0.658, 0.658] (0.653, 0.677)
B. Differences		Ex-post minus second-best	Second-best minus real		Ex-post minus real
Expected gains from trade		[0.009, 0.077] (0.007, 0.082)	[0.223, 0.834] (0.195, 0.955)		[0.289, 0.864] (0.260, 0.985)
Probability of trade		[0.042, 0.155] (0.035, 0.157)	[0.045, 0.174] (0.029, 0.179)		[0.199, 0.235] (0.177, 0.238)

Notes: Fleet/lease sample. Bounds on welfare measures under ex-post efficient, second-best, buyer-optimal, and seller-optimal mechanisms compared to real-world mechanism. Panel A displays levels and panel B displays differences. Estimated bounds are in square braces and 95% confidence set is in parentheses. Gains are in \$1,000 units.

Table 6: Expected Gains From Trade in Bargaining Between a Random Buyer and Seller

	Ex-post	Second-best	Real bargaining	Ex-post minus real
Dealers Sample	[2.515, 4.888]	[2.423, 4.837]	[1.539, 2.684]	[0.976, 2.204]
	(2.113, 5.049)	(2.026, 4.991)	(1.189, 2.867)	(0.787, 2.675)
Fleet/lease Sample	[3.095, 4.971]	[2.972, 4.900]	[1.571, 2.207]	[1.525, 2.765]
	(2.926, 5.192)	(2.796, 5.144)	(1.464, 2.586)	(1.233, 2.995)

Notes: Bounds on expected gains from trade in ex-post efficient, second-best, real-world bargaining, as well as the gap between ex-post efficient and real-world bargaining, when a random buyer (rather than the high bidder) bargains with a seller. As explained in Section 6.4, this setting is computed by setting the auction price to a low quantile of the auction price distribution. Units are \$1,000. Estimated bounds are in square braces and 95% confidence set is in parentheses.

A Proofs

Before providing the proofs corresponding to results in the main text, I first introduce some additional notation and state some preliminary lemmas. Let $H_t \equiv \{P_\tau\}_{\tau=1}^{t-1}$ represent the set of offers made from period 1 up through period $t - 1$ of the bargaining game. The player whose turn it is at time t has not yet made an offer and so this offer does not enter into H_t . Let $D_t^S \in \{A, Q, C\}$ represent the seller's decision in period t , and let $D_{t+1}^B \in \{A, Q, C\}$ represent the buyer's decision in period $t + 1$.

The seller's payoff at period t of the bargaining game is given by the following. Conditional on a realization of the history $H_t = h_t$, which includes the buyer's most recent offer (p_{t-1}^B), a seller of type $S = s$, chooses to accept (A), quit (Q), or counter (C), yielding the following payoffs:

$$\begin{aligned}
\text{A} &: p_{t-1}^B \\
\text{Q} &: s \\
\text{C} &: V_t^S(s|h_t) \\
&= \max_p \left\{ p \Pr(D_{t+1}^B = A | \{h_{t+1}, p\}) + s \Pr(D_{t+1}^B = Q | \{h_{t+1}, p\}) - c_S \right. \\
&\quad \left. + \Pr(D_{t+1}^B = C | \{h_{t+1}, p\}) E_{P_{t+1}^B} \left[\max \{ P_{t+1}^B, s, V_{t+2}^S(s|\{h_{t+1}, p, P_{t+1}^B\}) \} \mid \{h_{t+1}, p\}, D_{t+1}^B = C \right] \right\}
\end{aligned}$$

where p is the counteroffer chosen by the seller. The seller's counteroffer payoff takes into account that the buyer may either accept, quit, or return a counteroffer. In the latter case, the seller receives her expected payoff from being faced with the decision in period $t + 2$ to accept, quit, or counter.

The buyer's payoff at period $t + 1$ of the bargaining game is defined similarly, with the buyer receiving $b - p$ if he accepts a price p , 0 if he quits, and an expected counteroffer payoff if he counters. Conditional on a realization of the history $H_{t+1} = h_{t+1}$, which includes the seller's most recent offer (p_t^S), a buyer of type $B = b$ chooses to accept (A), quit (Q), or counter (C), yielding the following payoffs:

$$\begin{aligned}
\text{A} &: b - p_t^S \\
\text{Q} &: 0 \\
\text{C} &: V_{t+1}^B(b|h_{t+1}) \\
&= \max_p \left\{ (b - p) \Pr(D_{t+2}^S = A | \{h_{t+1}, p\}) - c_B \right. \\
&\quad \left. + \Pr(D_{t+2}^S = C | \{h_{t+1}, p\}) E_{P_{t+2}^S} \left[\max \{ b - P_{t+2}^S, 0, V_{t+3}^B(b|\{h_{t+1}, p, P_{t+2}^S\}) \} \mid \{h_{t+1}, p\}, D_{t+2}^S = C \right] \right\}
\end{aligned}$$

where p is the counteroffer chosen by the buyer. The buyer's outside option is normalized to zero.

The expected payoff of a buyer of type $B = b$ in the bargaining subgame, conditional on winning the auction and conditional on entering bargaining when the auction price is p^A , is given by

$$\begin{aligned} \pi^B(p^A, b) &= (b - p^A) \Pr(D_2^S = A | p^A) \\ &\quad + \Pr(D_2^S = C | p^A) E_{P_2^S} \left[\max \{b - P_2^S, 0, V_3^B(b | \{p^A, P_2^S\})\} \middle| p^A, D_2^S = C \right] - c_B \end{aligned}$$

This expression is the payoff to the buyer from stating the auction price as a counteroffer, which is how the bargaining game begins.

Lemma 1. *If Assumptions A1–A3 are satisfied, then for any finite T and any realized histories h_t and h_{t+1} , $V_t^S(s | h_t)$ is weakly increasing in s and $V_{t+1}^B(b | h_{t+1})$ is weakly increasing in b for all $t \leq T$.*

Proof. The proof proceeds by induction on the number of periods remaining. I prove the result in the case where the buyer moves last; analogous reasoning proves that the result also holds if the seller moves last. Suppose there are T total periods in the game and there is currently one period remaining: it is the seller's turn and after her turn the buyer will only be allowed to accept or quit. At a given realization of $H_{T-1} = h_{T-1}$, the seller's payoff from countering at a price of p is then

$$U_T^S(s, p | h_{T-1}) \equiv p \Pr(D_T^B = A | \{h_{T-1}, p\}) + s(1 - \Pr(D_T^B = A | \{h_{T-1}, p\})) - c_S$$

Let $p^*(s | h_{T-1}) \in \arg \max_p U_T^S(s, p | h_{T-1})$. That is, $V_{T-1}^S(s | h_{T-1}) = U_{T-1}^S(s, p^*(s | h_{T-1}) | h_{T-1})$. Now let $V_{T-1}(s, s' | h_{T-1})$ represent the payoff to the seller of type s who mimics type $s' < s$ (note that the ability of a seller—or buyer—to mimic another type relies on the i.i.d. properties in Assumptions A1–A2). Clearly $V_{T-1}(s, s | h_{T-1}) \geq V_{T-1}(s, s' | h_{T-1})$ because $V_{T-1}(s, s | h_{T-1})$ is the maximized counteroffer payoff given the seller's true value, s . It remains to be shown that $V_{T-1}(s, s' | h_{T-1}) \geq V_{T-1}(s', s' | h_{T-1})$.

Below, $p^*(s' | h_{T-1})$ represents the offer that would be optimal for a seller of type s' given the realized history h_{T-1} . Observe that

$$\begin{aligned} V_{T-1}(s, s' | h_{T-1}) &= p^*(s' | h_{T-1}) \Pr(D_T^B = A | \{h_{T-1}, p^*(s' | h_{T-1})\}) \\ &\quad + s(1 - \Pr(D_T^B = A | \{h_{T-1}, p^*(s' | h_{T-1})\})) - c_S, \end{aligned}$$

and

$$\begin{aligned} V_{T-1}(s', s' | h_{T-1}) &= p^*(s' | h_{T-1}) \Pr(D_T^B = A | \{h_{T-1}, p^*(s' | h_{T-1})\}) \\ &\quad + s'(1 - \Pr(D_T^B = A | \{h_{T-1}, p^*(s' | h_{T-1})\})) - c_S \end{aligned}$$

Thus,

$$V_{T-1}(s, s'|h_{T-1}) - V_{T-1}(s', s'|h_{T-1}) = (s - s')(1 - \Pr(D_T^B = A|\{h_{T-1}, p^*(s'|h_{T-1})\})) \geq 0$$

Therefore, $V_{T-1}(s, s|h_{T-1}) \geq V_{T-1}(s', s'|h_{T-1})$, and the seller's counteroffer payoff is weakly increasing in her type when there is one period remaining.

To complete the proof by induction, let $V_{T-(t-1)}^S(s|h_{T-(t-1)})$ denote the seller's counteroffer payoff with $t - 1$ periods remaining, and suppose $V_{T-(t-1)}^S(s|h_{T-(t-1)})$ is weakly increasing in s . Note that, for $s' < s$, when there are t periods remaining, $V_{T-t}(s, s|h_{T-t}) \geq V_{T-t}(s', s'|h_{T-t})$ by the same argument as above for. It remains to be shown that $V_{T-t}(s, s'|h_{T-t}) \geq V_{T-t}(s', s'|h_{T-t})$.

Note that

$$\begin{aligned} & V_{T-t}(s, s'|h_{T-t}) - V_{T-t}(s', s'|h_{T-t}) \\ &= (s - s') \Pr \left(D_{T-(t-1)}^B = Q|\{h_{T-t}, p^*(s'|h_{T-t})\} \right) + \Pr \left(D_{T-(t-1)}^B = C|\{h_{T-t}, p^*(s'|h_{T-t})\} \right) \\ & \times E_{P_{T-(t-1)}^B} \left[\max \left\{ P_{T-(t-1)}^B, s, V_{T-(t-2)}^S \left(s, s'|\{h_{T-t}, p^*(s'|h_{T-t}), P_{T-(t-1)}^B \} \right) \right\} \right. \\ & \left. - \max \left\{ P_{T-(t-1)}^B, s', V_{T-(t-2)}^S \left(s', s'|\{h_{T-t}, p^*(s'|h_{T-t}), P_{T-(t-1)}^B \} \right) \right\} \right] \Big| \{h_{T-t}, p^*(s'|h_{T-t})\}, D_{T-(t-1)}^B = C \\ & \geq 0 \end{aligned}$$

Therefore, $V_{T-t}(s, s|h_{T-t}) \geq V_{T-t}(s', s'|h_{T-t})$, completing the proof. The proof that the buyer counteroffer payoff, $V_{t+1}^B(b|h_{t+1})$, is increasing in b follows by the same steps. \square

Lemma 2. *In any equilibrium with truth-telling in the auction, f_B being positive everywhere on the support of B implies f_{PA} is positive everywhere on the support of B .*

Proof. If f_B is positive everywhere on the support of B , then F_B is strictly increasing. Equation (3) demonstrates that, under truth-telling in the auction, F_{PA} is then also strictly increasing, and so f_{PA} is positive everywhere on the support of B . \square

Proof of Proposition 1

Proof. Consider an arbitrary bidder of type $B = b$. A bidder's strategy is the price at which he stops bidding as a function of his type. Suppose the current price of the ascending button auction is some value \bar{p} and suppose the bidder is one of at least two bidders still remaining in the auction up until the price reaches its current level \bar{p} . The auction will eventually end at some price $p^A \geq \bar{p}$.

If $b > \bar{p}$, it is optimal for the bidder to remain in the auction, as dropping out would yield a payoff of 0 and staying in would yield a non-negative expected payoff because there is some chance that the bidder will win at a price $p^A < b$. To see that the expected payoff of remaining in the

bidding is non-negative, consider first the case where the auction price and reserve price satisfy $p^A \geq R$; in this case, the car will sell through the auction and the bidder will receive a positive payoff. Now consider the case where instead $p^A < R$; in this case, the bidder will be given the option to enter into bargaining, and will only choose to enter if doing so yields a non-negative expected payoff. Recall that the bidder only pays the cost of bargaining, c_B , if he chooses not to opt out of bargaining at this point.

If $b < \bar{p}$, the buyer cannot receive a positive expected payoff from remaining in the auction. To see this, note that if the bidder remains in the auction there is some chance that he will win at some $p^A > b$. If this occurs and the auction price and reserve price satisfy $p^A \geq R$, the car will sell through the auction and the bidder will receive $b - p^A < 0$. If, on the other hand, the bidder wins and $p^A < R$, the bidder's payoff conditional on entering bargaining will necessarily be negative because the final bargained price must be greater than p^A and hence, in this case, the bidder will opt out of bargaining, receiving a payoff of 0. \square

Proof of Proposition 3

Proof. Note that for $b' > b$, $\pi^B(\chi(b), b') > 0$. This follows because, by Lemma 1, $V_3^B(\cdot)$ is weakly increasing in b , and this fact, combined with the term $(b - p^A)$ appearing in $\pi^B(p^A, b)$, which is strictly increasing in b , implies that $\pi^B(p^A, b)$ is strictly increasing in b . Thus, $\chi(b') > \chi(b)$, and hence χ is strictly increasing, and χ^{-1} exists and is also strictly increasing.

The property that $\chi^{-1}(p^A) > p^A$ follows from the following argument. A buyer must pay $c_B > 0$ if he opts to bargain, and the best possible outcome a buyer can expect from bargaining would be to only have to pay p^A . Therefore, for any auction price p^A , there exists some buyer with type close to p^A , say $p^A + \varepsilon$, where $\varepsilon < c_B$, who would prefer to opt out of bargaining rather than receive a payoff of (at most) $\varepsilon - c_B$, which is negative.

Strict monotonicity of $\rho(\cdot)$, along with the fact that $\rho(s) \geq s$, is proven separately in Lemma 3 below.

When the auction price is p^A and bargaining occurs, it will be common knowledge among the two bargaining parties that seller's type s satisfies $\rho(s) \geq p^A$, and thus $s \in [\rho^{-1}(p^A), \bar{s}]$. Similarly, bargaining occurring means the buyer did not opt out, so $\chi(b) \geq p^A$, implying $b \in [\chi^{-1}(p^A), \bar{b}]$. \square

Lemma 3. *If Assumptions A1–A5 are satisfied, then in any BNE satisfying Assumption A4, the seller's optimal secret reserve price, $\rho^*(s)$, is strictly increasing in s and satisfies $\rho^*(s) \geq s$.*

Proof. Suppose for simplicity that there is some positive probability that the buyer does not opt out of bargaining when $R > P^A$; as I discuss below, the proposition still holds without this assumption. In choosing her secret reserve price, $\rho(s)$, a seller of type $S = s$ wishes to maximize her ex-ante

payoff, given by

$$E_{P^A} \left[E_B \left[P^A * 1 \{P^A \geq \rho(s)\} + s * 1 \{P^A < \rho(s), \pi^B(P^A, B) < 0\} \right. \right. \\ \left. \left. + \pi^S(P^A, s) * 1 \{P^A < \rho(s), \pi^B(P^A, B) \geq 0\} \middle| P^A \right] \right] \quad (9)$$

This term consists of three pieces: 1) the auction price, which the seller receives if it exceeds the reserve price; 2) the seller's type s , which the seller receives if the auction price is below the reserve price and the buyer opts out of bargaining; and 3) the seller's bargaining payoff, $\pi^S(P^A, s) = \max \{P^A, s, V_2^S(s|P^A)\}$, which the seller receives when the price is below the reserve price and bargaining occurs.

The seller's payoff can be re-written as

$$\int_{\rho}^{\bar{b}} p^A f_{P^A}(p^A) dp^A + \int_{\underline{b}}^{\rho} \left[\int_{p^A}^{\chi^{-1}(p^A)} s f_B(b) db + \int_{\chi^{-1}(p^A)}^{\bar{b}} \pi^S(p^A, s) f_B(b) db \right] \frac{f_{P^A}(p^A)}{1 - F_B(p^A)} dp^A \\ = \int_{\rho}^{\bar{b}} p^A f_{P^A}(p^A) dp^A + \int_{\underline{b}}^{\rho} \left[s (F_B(\chi^{-1}(p^A)) - F_B(p^A)) + \pi^S(p^A, s) (1 - F_B(\chi^{-1}(p^A))) \right] \frac{f_{P^A}(p^A)}{1 - F_B(p^A)} dp^A$$

Assumption A4 implies $\pi^S(\cdot, s)$ is continuous and thus the payoff is differentiable. Differentiating the above expression using Leibniz Rule yields the following derivative with respect to ρ :

$$\frac{\partial}{\partial \rho} = f_{P^A}(\rho) \left[-\rho + s \frac{F_B(\chi^{-1}(\rho)) - F_B(\rho)}{1 - F_B(\rho)} + \pi^S(\rho, s) \frac{1 - F_B(\chi^{-1}(\rho))}{1 - F_B(\rho)} \right] \quad (10)$$

I next show that $f_{P^A}(\rho) > 0$ for any ρ in the support of R . To see this, first note that $f_{P^A}(v) > 0$ for all $v \in [\underline{b}, \bar{b}]$ by Lemma 2. Second, choosing any $\rho < \underline{b}$ would be dominated by instead choosing a reserve price of \underline{b} because every buyer has a valuation of at least \underline{b} . Third, a seller would be indifferent between any reserve price above \bar{b} (because no buyer would be willing to pay more than \bar{b}). Therefore, any $\rho \notin [\underline{b}, \bar{b}]$ would be weakly dominated, and so I can remove $f_{P^A}(\rho)$ from the above expression without dividing by zero. Also note that $1 - F_B(\rho) > 0$ because f_B is positive by Assumption A5.

Lemma 1 shows that $\pi^S(p^A, s)$ is weakly increasing in s and so $\frac{\partial}{\partial \rho}$ is weakly increasing in s . Topkis Theorem then implies that $\rho^*(s)$ is weakly increasing in s .

A stronger, strictly increasing result for $\rho^*(s)$ is then obtained as follows. The proof of Proposition 3 demonstrates (due to costly bargaining) that $\chi^{-1}(p^A) > p^A$, and thus $F_B(\chi^{-1}(\rho)) > F_B(\rho)$. Combining these arguments implies that $\frac{\partial}{\partial \rho}$ is strictly increasing in s , which in turn implies, by the Edlin and Shannon (1998) Theorem, that $\rho^*(s)$ is strictly increasing on the interior of the support of R . Let \underline{r} and \bar{r} be the infimum and supremum of $\rho(s)$ that are optimal for any s . Suppose the

support of R is an interval. For the purposes of applying the Edlin and Shannon (1998) result, the support of r can be considered to be $(\underline{r} - \varepsilon, \bar{r} + \varepsilon)$ for some $\varepsilon > 0$, and $\rho^*(s)$ will be strictly increasing on the interior of this interval. Now suppose the support of R is not an interval, i.e. suppose there exist a point or points on the interior of $[\underline{r}, \bar{r}]$ that are not optimal for any $s \in [\underline{s}, \bar{s}]$. Such points would constitute discontinuities of the function $\rho^*(s)$. By the weakly increasing result from above (due to Topkis Theorem), any such discontinuities are positive jumps in the function $\rho^*(s)$, and therefore $\rho^*(s)$ will be strictly increasing even if the support of R is not an interval.

The fact that $\rho^*(s) \geq s$ follows from a simple rationality argument (no seller would offer a reserve price less than s given that a reserve price of s yields a weakly higher payoff), but it can also be seen by noting that the first-order condition above implies that the reserve price is given by a convex combination of s and a quantity weakly greater than s (i.e. $\pi^S(\rho, s)$).

Note that if the equilibrium of the game is such that the buyer opts out of bargaining with probability one when $R > P^A$ then the expression in (10) becomes $\frac{\partial}{\partial \rho} = f_{PA}(\rho) [-\rho + s]$, and thus the optimal reserve price is $\rho^*(s) = s$, again satisfying the proposition. \square

Proof of Proposition 4

Proof. Theorem 1 of Storms (2015) (included below as Lemma 4, modified to fit this setting) implies that, in any BNE of this game, conditional on a realization of the auction price, p^A , for each seller type s , there is a cutoff value $g_0(s, p^A)$ such that trade occurs if and only if the buyer's type b satisfies $b \geq g_0(s, p^A)$. Given the strict monotonicity of $\rho(\cdot)$ (Proposition 3), such a cutoff function also exists with realizations s replaced with realizations of the reserve price r . Call this cutoff function $g(r, p^A)$. \square

Lemma 4. *(Due to Storms 2015) If Assumptions A1–A5 are satisfied, then, conditional on any realization of the auction price $P^A = p^A$, in any BNE of the bargaining subgame satisfying A4, for each seller type s there is a cutoff value $g_0(s, p^A)$ such that s trades with a buyer b if and only if $b \geq g_0(s, p^A)$.*

Proof. Fix $P^A = p^A$ throughout this proof. I first prove a preliminary property. Fix any arbitrary BNE. Let $\Pr(\mathcal{A} = 1|b, h^t)$ represent the probability of trade for a buyer who mimics the strategy of a buyer of type b when the history so far in the game is h_t . Here, $\mathcal{A} \in \{0, 1\}$ is a random variable indicating whether or not trade occurs, where, from the buyer's perspective, the seller's valuation is unknown. Let $y(b, h_t)$ represent the expected transfer from playing such an action. Also, let $h_t(s, b)$ denote the history of the game in time t when the players' types are s, b and when they play their equilibrium strategies.

I will discuss properties that must hold on histories that have a positive probability of being played in equilibrium (i.e. histories that at least some buyer and seller pair would play). In such

histories, in any BNE, each buyer type must weakly prefer to play his own strategy from any history onward to playing that of another type. Thus, for $b' > b$, we have

$$\begin{aligned} b \Pr(\mathcal{A} = 1|b, h_t) - y(b, h_t) &\geq b \Pr(\mathcal{A} = 1|b', h_t) - y(b', h_t) \\ b' \Pr(\mathcal{A} = 1|b', h_t) - y(b', h_t) &\geq b' \Pr(\mathcal{A} = 1|b, h_t) - y(b, h_t) \end{aligned}$$

Combining inequalities demonstrates that, for $b' > b$,

$$\Pr(\mathcal{A} = 1|b', h_t) \geq \Pr(\mathcal{A} = 1|b, h_t) \quad (11)$$

A similar result holds for $s' < s$,

$$\Pr(\mathcal{A} = 1|s', h_t) \geq \Pr(\mathcal{A} = 1|s, h_t) \quad (12)$$

Using this property, Lemma 4 can be proved by contradiction. Such a contradiction would be a triple s, b , and b' with $b' > b$ such that s eventually (at some unspecified time period of the game) trades with b , but does not at any period of the game reach agreement with a type b' . For the sake of clarity, I will give such triples a name, referring to them as *Type A* triples. Let h_t^* be the longest history of play *among all Type A triples* such that the strategy for b is the same as that for b' up to time t when the seller's type is s (that is, $h_t^* = h_t(s, b) = h_t(s, b')$). Throughout the remainder of the proof, let s, b , and b' be a Type A triple at which h_t^* is achieved. The result in (11) implies that b' must trade with weakly greater probability than b from h_t onward. This weak inequality, combined with (s, b, b') being a Type A triple, implies that there must be some seller type s' who reaches history h_t against both b and b' and who trades with b' but not b .

Now consider two cases.

1. Case where $s' > s$. Since s does not trade with b' from the history $h_{t+1}(s, b')$, s cannot trade with any types $\tilde{b} < b'$ from $h_{t+1}(s, b')$, or else (s, \tilde{b}, b') would form a counterexample to h_t^* because it would constitute a Type A triple with buyers having $t + 1$ periods of identical strategies. But by (12), s must trade more often than s' conditional on the history $h_{t+1}(s, b')$, and hence there must be some type $b'' > b'$ such that b'' eventually trades with s but not with s' when the history is $h_{t+1}(s, b')$. The triple (s', b', b'') then gives a contradiction because it constitutes a Type A triple with buyers having $t + 1$ stages of their strategies being identical.
2. Case where $s' < s$. Since s trades with b from the history $h_{t+1}(s, b)$, s must trade with all types $\tilde{b} > b$ from $h_{t+1}(s, b)$, or else (s, b, \tilde{b}) would form a counterexample to h_t^* because it would constitute a Type A triple with buyers having $t + 1$ periods of identical strategies. By (12), s' must trade more often than s conditional on the history $h_{t+1}(s, b)$. It follows that there must be some type $b''' < b$ that trades with s' but not s . The triple (s', b, b''') then gives

a contradiction because it constitutes a Type A triple with buyers having $t + 1$ stages of their strategies being identical.

□

Proof of Proposition 5

Proof. Given the structure of additive separability in the willingness to pay/sell, the goal is to show that the auction price, players' bargaining counteroffers, and the seller's secret reserve price will also be additively separable in the game-level heterogeneity. The buyer's type is given by $\tilde{B} \equiv B + W \sim F_{\tilde{B}}$, with density $f_{\tilde{B}}$. The seller's type is given by $\tilde{S} \equiv S + W$. For this proof, let the realization of W be w .

That the auction price will be additively separable in w is obvious, given that there is no incentive for bidders to deviate from truthful bidding by Proposition 1.

To demonstrate that bargaining offers are also additively separable, the proof proceeds by induction on the number of periods remaining. Before proving this result, I highlight here that Proposition 5 only states that bargaining offers will be additively separable if they are accepted with positive probability; the equilibrium framework does not rule out equilibria in which a player makes an offer that would not be accepted by any type, and additive separability will not necessarily hold for such offers.

Suppose there is currently one period remaining in the bargaining game: it is the seller's turn and after her turn the buyer will only be allowed to accept or quit (I prove the result in the case where the buyer moves last; analogous reasoning proves that the result also holds if the seller moves last).

Suppose for simplicity that the equilibrium does not entail the buyer rejecting all offers with probability one in the final period (if not, the seller would not choose to counter in period $T - 1$). In the final period, a buyer with type $\tilde{B} = \tilde{b}$ will accept a price, \tilde{p} , if and only if $\tilde{p} \leq \tilde{b}$. In period $T - 1$, the seller of type $\tilde{S} = \tilde{s}$ chooses \tilde{p}^* to solve

$$\begin{aligned} \tilde{p}^* &\in \arg \max_{\tilde{p}} \tilde{p}(1 - F_{\tilde{B}}(\tilde{p})) + \tilde{s}F_{\tilde{B}}(\tilde{p}) - c_S \\ &= \arg \max_p p(1 - F_B(p)) + sF_B(p) - c_S + w(1 - F_B(p)) + wF_B(p) \\ &= w + \arg \max_p \{p(1 - F_B(p)) + sF_B(p) - c_S\} \end{aligned}$$

Therefore, the penultimate bargaining offer in the heterogeneous setting will be w above the bargaining offer from the homogeneous good setting, and similarly for the seller's maximized payoff.

To complete the proof by induction, suppose that offers and payoffs in periods $T - (t - 1)$ and $T - (t - 2)$ are w higher than their homogeneous good counterparts and the probability of the *buyer* accepting, quitting, or countering in period $T - (t - 1)$ will be the same in the heterogeneous good

model as in the homogeneous good model. It remains to be shown that the same holds true for the offers and payoffs in period $T - t$. Let all $(\tilde{\cdot})$ expressions represent the heterogeneous model expressions. The seller's payoffs from accepting, quitting, or countering in period $T - t$ can be written as follows:

$$\begin{aligned}
\text{A} : \tilde{p}_{T-(t+1)}^B &= w + p_{T-(t+1)}^B \\
\text{Q} : \tilde{s} &= w + s \\
\text{C} : \tilde{V}_{T-t}^S(\tilde{s}|\tilde{h}_{T-t}) & \\
&= \max_{\tilde{p}} \tilde{p} \Pr\left(D_{T-(t-1)}^B = A|\{\tilde{h}_{T-t}, \tilde{p}\}\right) + \tilde{s} \Pr\left(D_{T-(t-1)}^B = Q|\{\tilde{h}_{T-t}, \tilde{p}\}\right) \\
&+ \Pr\left(D_{T-(t-1)}^B = C|\{\tilde{h}_{T-t}, \tilde{p}\}\right) \\
&\times E_{\tilde{P}_{T-(t-1)}^B} \left[\max \left\{ \tilde{P}_{T-(t-1)}^B, \tilde{s}, \tilde{V}_{T-(t-2)}^S(\tilde{s}|\{\tilde{h}_{T-t}, \tilde{p}, \tilde{P}_{T-(t-1)}^B\}) \right\} \middle| \{\tilde{h}_{T-t}, \tilde{p}\}, D_{T-(t-1)}^B = C \right] - c_S \\
&= w + \max_p p \Pr\left(D_{T-(t-1)}^B = A|\{h_{T-t}, p\}\right) + s \Pr\left(D_{T-(t-1)}^B = Q|\{h_{T-t}, p\}\right) \\
&+ \Pr\left(D_{T-(t-1)}^B = C|\{h_{T-t}, p\}\right) \\
&\times E_{P_{T-(t-1)}^B} \left[\max \left\{ P_{T-(t-1)}^B, s, V_{T-(t-2)}^S(s|\{h_{T-t}, p, P_{T-(t-1)}^B\}) \right\} \middle| \{h_{T-t}, p\}, D_{T-(t-1)}^B = C \right] - c_S
\end{aligned}$$

The last line follows by removing w from each expression. Note that if the equilibrium entails the buyer quitting with probability 1 in period $T - (t - 1)$ then the seller would have no incentive to counter in period $T - t$ and so the additive separability result still holds for offers observed in equilibrium.

To see that the probability of the buyer accepting, quitting, or countering in period $T - (t - 1)$ will be the same in the heterogeneous good model as in the homogeneous good model, note that the buyer's payoffs for each action are given by:

$$\begin{aligned}
\text{A} : \tilde{b} - \tilde{p}_{T-t}^S &= b - p_{T-t}^S \\
\text{Q} : 0 & \\
\text{C} : \tilde{V}_{T-(t-1)}^B(\tilde{b}|\tilde{h}_{T-(t-1)}) & \\
&= \max_{\tilde{p}} (\tilde{b} - \tilde{p}) \Pr\left(D_{T-(t-2)}^S = A|\{\tilde{h}_{T-(t-1)}, \tilde{p}\}\right) - c_B \\
&+ \Pr\left(D_{T-(t-2)}^S = C|\{\tilde{h}_{T-(t-1)}, \tilde{p}\}\right) \\
&\times E_{\tilde{P}_{T-(t-2)}^S} \left[\max \left\{ \tilde{b} - \tilde{P}_{T-(t-2)}^S, 0, \tilde{V}_{T-(t-3)}^B(\tilde{b}|\{\tilde{h}_{T-(t-1)}, \tilde{p}, \tilde{P}_{T-(t-2)}^S\}) \right\} \middle| \{\tilde{h}_{T-(t-1)}, \tilde{p}\}, D_{T-(t-2)}^S = C \right] \\
&= V_{T-(t-1)}^B(b|h_{T-(t-1)})
\end{aligned}$$

And thus the claim is true because the buyer's bargaining payoffs are the same in the homogeneous good and heterogeneous good models. It also immediately follows that $\pi^B(\tilde{\chi}, \tilde{b}) = \pi^B(\chi, b)$ by the above arguments for the buyer's bargaining payoff, where $\tilde{\chi}$ satisfies $0 = \pi^B(\tilde{\chi}, \tilde{b})$.

Now consider the seller's secret reserve price in the setting with game-level heterogeneity w . Suppose for now that the equilibrium does not entail the buyer opting out of bargaining with probability 1 when $R > P^A$. From the proof of Lemma 3, the derivative of the seller's payoff with respect to the seller's choice of secret reserve price, $\tilde{r} = \rho(\tilde{s})$, will be given by

$$\frac{\partial}{\partial \tilde{r}} = f_{\tilde{P}^A}(\tilde{r}) \left[-\tilde{r} + \tilde{s} \frac{F_{\tilde{B}}(\tilde{\chi}^{-1}(\tilde{r})) - F_{\tilde{B}}(\tilde{r})}{1 - F_{\tilde{B}}(\tilde{r})} + \pi^S(\tilde{r}, \tilde{s}) \frac{1 - F_{\tilde{B}}(\tilde{\chi}^{-1}(\tilde{r}))}{1 - F_{\tilde{B}}(\tilde{r})} \right] \quad (13)$$

$$= f_{P^A}(\tilde{r} - w) \left[-\tilde{r} + w + s \frac{F_B(\chi^{-1}(\tilde{r} - w)) - F_B(\tilde{r} - w)}{1 - F_B(\tilde{r} - w)} + \pi^S(\tilde{r} - w, s) \frac{1 - F_B(\chi^{-1}(\tilde{r} - w))}{1 - F_B(\tilde{r} - w)} \right] \quad (14)$$

Therefore, the optimal secret reserve price in the heterogeneous setting will be w above the optimal reserve in the homogeneous setting. Now suppose the equilibrium does involve the buyer opting out of bargaining with probability 1 when $R > P^A$. In this case, the expression in (13) becomes $\frac{\partial}{\partial \tilde{r}} = -\tilde{r} + \tilde{s}$, and thus the optimal secret reserve price is $s + w$, again satisfying the proposition.

An immediate implication of these results is a generalization of Proposition 4: At a general realization $W = w$, trade occurs if and only if $\tilde{b} \geq g(\tilde{r}, \tilde{p}^A) \Rightarrow b \geq g(\tilde{r} - w, \tilde{p}^A - w) = g(r, p^A)$. \square

Proof of Proposition 6

Proof. I first introduce notation. Let $\mathcal{W}(x, F_S)$ refer to any welfare measure under a given allocation function x and at a given seller valuation distribution F_S (suppressing dependence on F_B given that I have point estimates of F_B). I will add specific subscripts to \mathcal{W} to denote a particular type of welfare measure: the total expected gains from trade is $\mathcal{W}_{EG}(x, F_S)$, given by (8); the expected gains from trade for the buyer alone is $\mathcal{W}_B(x, F_S)$ and for the seller alone is $\mathcal{W}_S(x, F_S)$; and the expected probability of trade is $\mathcal{W}_{Pr}(x, F_S)$. The gap between the expected gains from trade for the ex-post efficient and second-best mechanisms, between the second-best and real-world mechanisms, and between the ex-post efficient and real-world mechanisms are given by the following three expressions:

$$\begin{aligned} \mathcal{W}_{EG}(x^* - x^{\frac{1}{2}}, F_S) &\equiv \mathcal{W}_{EG}(x^*, F_S) - \mathcal{W}_{EG}(x^{\frac{1}{2}}, F_S) \\ \mathcal{W}_{EG}(x^{\frac{1}{2}} - x^{RW}, F_S) &\equiv \mathcal{W}_{EG}(x^{\frac{1}{2}}, F_S) - \mathcal{W}_{EG}(x^{RW}, F_S) \\ \mathcal{W}_{EG}(x^* - x^{RW}, F_S) &\equiv \mathcal{W}_{EG}(x^*, F_S) - \mathcal{W}_{EG}(x^{RW}, F_S) \end{aligned}$$

Similar notation can be used to denote the gap between the probability of trade in different mechanisms.

I begin by proving that $\mathcal{W}_{EG}(x^*, F_S)$ will decrease given a stochastically dominating change in F_S . Let $\mathcal{W}_{EG}(x^*, F_S; p^A)$ be the expected gains from trade under ex-post efficiency conditional on a realization of the auction price p^A . This object can be written as

$$\mathcal{W}_{EG}(x^*, F_S; p^A) = \int_{\underline{s}(p^A)}^{\bar{s}} \Lambda(s) f_S(s|p^A) ds \quad (15)$$

where $\Lambda(s, p^A)$ is given by

$$\Lambda(s, p^A) = \int_{\underline{b}(p^A)}^{\bar{b}} (b - s) 1\{b \geq s\} f_B(b|p^A) db$$

Note that this function is weakly decreasing in s . Also recall that $f_S(s|p^A) = \frac{f_S(s)}{1 - F_S(\underline{s}(p^A))}$, and note that the denominator of this term does not vary with F_S , because $F_S(\underline{s}(p^A)) = F_S(\rho^{-1}(p^A)) = F_S(F_S^{-1}(F_R(p^A))) = F_R(p^A)$. It is well known that, for any weakly decreasing $\Lambda(\cdot)$ in an expectation such as (15), replacing f_S with a density corresponding to a distribution that first-order stochastically dominates f_S will lead to a lower value for the evaluated integral $\mathcal{W}_{EG}(x^*, F_S; p^A)$. Integrating $\mathcal{W}_{EG}(x^*, F_S; p^A)$ over p^A yields the desired result. The monotonicity of $\mathcal{W}_{Pr}(x^*, F_S)$ follows by the same result.

For the second-best mechanism, the proof of monotonicity of the expected gains from trade is much more involved and is found in recent work by Zhang (2017). It does not yield an mathematical proof of the monotonicity of the expected buyer or seller gains or the probability of trade in the second-best mechanism. For the buyer-optimal mechanism, monotonicity of the buyer gains follows from the same line of reasoning as in Zhang (2017), but there is no proof of monotonicity of the total expected gains from trade, seller gains, or probability of trade.

For the seller-optimal mechanism, the allocation function is $x^1 = 1\{b - \frac{1 - F_b(b)}{F_B(b)} \geq s\}$. As in the ex-post efficient case, this allocation function does not depend on F_S , and thus monotonicity of all the welfare measures in the seller-optimal mechanism are guaranteed (even for those marked with “ $_$ ” in the seller-optimal column of Table 3).

For the real-world mechanism, the proof of monotonicity of the expected gains from trade exploits that I have defined the real-world mechanism as a function of the reserve price, R , given that R is a strictly increasing function of S , and hence, holding fixed F_R , the allocation function, $x^{RW}(\rho(s), b; p^A) \equiv 1\{b \geq g(\rho(s), p^A)\}$, does not depend on the unknown distribution F_S . The expected gains from trade in the real-world mechanism can be written

$$\mathcal{W}_{EG}(x^{RW}, F_S) = \int_{\underline{b}}^{\bar{b}} \left[\int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} (b - s) 1\{b \geq g(\rho(s), p^A)\} f_S(s|p^A) f_B(b|p^A) ds db \right] f_{p^A}(p^A) dp^A$$

where $\rho(s) \equiv F_R^{-1}(F_S(s))$. Holding F_R fixed, the object $\rho(s)$ would be unchanged by such

a first-order stochastically dominating shift in F_S . This is because $\rho(s)$ always returns the $F_S(s)$ quantile of F_R . To see this, let \check{F}_S first order stochastically dominate F_S and let \check{s} and s be realizations such that $\check{s} > s$ and $\check{F}_S(\check{s}) = F_S(s)$. Thus, $\check{s} = \check{F}_S^{-1}(F_S(s))$. Then $\check{\rho}(\check{s}) \equiv F_R^{-1}(\check{F}_S(\check{s})) = F_R^{-1}(\check{F}_S(\check{F}_S^{-1}(F_S(s)))) = \rho(s)$, proving the result. The monotonicity of $\mathcal{W}_B(x^{RW}, F_S)$, $\mathcal{W}_S(x^{RW}, F_S)$, and $\mathcal{W}_{Pr}(x^{RW}, F_S)$ follow by the same argument. Note that these results should not be interpreted as suggesting that the real-world mechanism does not depend on the distribution of seller valuations; these results simply imply that, because reserve prices are one-to-one with seller valuations, after conditioning on F_R , the estimation of the object $\rho(\cdot)$ will not depend on the estimate of F_S .

For the first column in panel B of Table 3, there is no mathematical proof of monotonicity of the gap between the ex-post and second-best gains from trade, $\mathcal{W}_{EG}(x^* - x^{\frac{1}{2}}, F_S)$, and similarly for the gap for the probability of trade. For the gap between the expected gains from trade in the second-best mechanism and the real-world mechanism, $\mathcal{W}_{EG}(x^{\frac{1}{2}} - x^{RW}, F_S)$, monotonicity follows from the arguments above (the allocation function x^{RW} does not depend on F_S , and the expected gains from trade for the second-best mechanism are monotonic). For the gap in the probability of trade there is no proof of monotonicity. For the final column in panel B of Table 3, monotonicity of the gap between ex-post efficient and real-world expected gains from trade, $\mathcal{W}_{EG}(x^* - x^{RW}, F_S)$, follows by the same arguments as above for the real-world and ex-post efficient mechanisms, and similarly for the probability of trade gap. \square

B Additional Model Discussion

B.1 Discussion of BNE vs. Other Equilibrium Concepts

Throughout the paper, I assume play corresponds to a BNE. The use of BNE, rather than a Perfect Bayesian Nash (PBE) framework, for example, does not mean I am modeling the game as a static game; I am modeling the game as dynamic—I am just not specifying how players update their beliefs at histories of the game that are never reached. In any BNE, players' beliefs about the type of opponent they are facing are updated using Bayes Rule at any history of the game reached on the equilibrium path, but BNE places no restrictions on how these beliefs are updated when off-equilibrium (i.e. zero-probability) events occur.²⁵ PBE is a refinement of BNE (and thus, every PBE is also a BNE) requiring that the researcher also specify how beliefs are updated when these zero-probability events occur. Refinements such as PBE are useful in some applied theory settings to narrow down the set of possible equilibria, but they have not proven very useful

²⁵Because Bayes Rule is the only updating rule allowed on-path in BNE, and because BNE includes no statement of off-equilibrium beliefs, it is not technically necessary to discuss beliefs at all a BNE framework; it is only for clarity of exposition that I explicitly state the feature that beliefs are updated using Bayes Rule at all histories of the game that are reached in equilibrium.

for sequential incomplete-information bargaining games like mine.²⁶ In spite of the independence assumptions in *A1–A2*, and the simple bargaining structure imposed in Assumption *A3*, this game is difficult to analyze, both theoretically and empirically, as two-sided bargaining problems (i.e. games where both parties can make offers) with two-sided uncertainty (where both parties have private information) are known to have a multiplicity of equilibria that can be qualitatively quite dissimilar, and no known characterization of these equilibria exists (see discussion in Ausubel, Cramton, and Deneckere 2002). The approach I take in this paper will circumvent this problem by only assuming that participants play *some* BNE of the game, without attempting to fully characterize or solve for that equilibrium.

B.2 Additional Discussion of Proposition 1 Truthful Bidding Result

The proof of Proposition 1 is greatly simplified by the rule, discussed in Section 3, that the bargained price cannot be below the auction price; if this were not the case, proving the proposition would require ruling out the possibility of buyers bidding above their valuations and then attempting to bargain down to a lower final price afterward. It is possible to show that such behavior cannot occur in equilibrium, but the proof is more involved. I state this as the following lemma:

Lemma 5. *If the rules of the game are relaxed such that the bargained price is allowed to be lower than the auction price, then it is still the case that the following cannot occur in equilibrium: Some bidder i remains in the bidding even after the current bid exceeds b_i and, if bargaining occurs and if given the chance to counteroffer, the bidder makes a counteroffer less than b_i and this offer is accepted by the seller.*

Proof. Note that a buyer who bids above his value bears a risk of winning the auction at a price that exceeds the reserve price, giving the bidder a negative payoff. Therefore, if some bidder i does want to bid above his value in the auction, it must be because that bidder hopes to eventually end up in bilateral bargaining with the seller and hopes to have the bargaining game end at a price weakly below his value, and the bidder believes that the chance of this happening is sufficiently high to warrant the risk of bidding above his value. That is, for this behavior to occur in equilibrium,

²⁶Previous work has found that such refinements can lead to predictions of immediate trade or immediate disagreement (Perry 1986) or that equilibria satisfying proposed refinements often fail to exist (see the discussion in Sections 3.1.2 and 5 of Ausubel, Cramton, and Deneckere (2002) of equilibrium refinements of incomplete-information bargaining games proposed in Cramton (1984), Rubinstein (1985), Grossman and Perry (1986), Cho (1990), and Bikhchandani (1992). Gul and Sonnenschein (1988) pointed out that traditional refinements (such as Sequential Equilibrium or Perfect Bayesian Equilibrium) often do not meaningfully narrow down equilibria in sequential bargaining games because BNE can often be made sequential/perfect by specifying optimistic beliefs: whenever the seller makes an off-equilibrium offer the buyer believes the seller is the weakest type, and vice versa. The bargaining game described in Perry (1986) is nearly identical to the setup of the bargaining subgame described herein, but the author focuses on Sequential Equilibria rather than BNE and finds a unique equilibrium that involves only immediate trade or immediate disagreement. Cramton (1991) discussed how the Perry (1986) result can be overturned by allowing for a small amount of time discounting. Cramton (1992) derived a particular equilibrium in a two-sided bargaining game involving a war of attrition.

it must be the case that, in the bargaining game, i makes a counteroffer lower than p^A and this offer is accepted by some seller (the seller should never make such an offer herself because accepting p^A would be preferable). Let the term *high-low strategy* refer to this strategy of remaining in the bidding even after the current bid exceeds b_i and then, if given the chance to do so, making an offer less than b_i (and hence also less than p^A).

In order for the high-low strategy to be optimal for b_i , the following must be true (I will refer to this as the *high-low supposition*): the seller cannot distinguish between type b_i and some type $\check{b} > b_i$ whom the seller believes may not be playing the high-low strategy (and hence may not be attempting to later counter below p^A). If the seller could distinguish between the two, the seller would simply accept p^A when faced with the bidder known to be playing the high-low strategy. Note that it cannot be the case that all bidders play the high-low strategy, or else the seller would always immediately accept the auction price and the bidders would all obtain a negative surplus and hence all bidders would have been better off not bidding above their values. Thus, in order for the high-low strategy to be optimal for b_i , it must be the case that the seller cannot distinguish between type b_i and some type $\check{b} > b_i$ whom the seller believes may not be playing the high-low strategy.

Suppose the high-low supposition is true. Note that, if type $b_i > p^A$ finds it optimal to play the high-low strategy, so will all types higher than b_i , because $\pi^B(p^A, b_i) \geq 0 \Rightarrow \pi^B(p^A, \check{b}) > 0$ for all $\check{b} > b_i$ (because $V_3^B(\cdot)$ is weakly increasing in b_i by Lemma 1 and $(b_i - p^A)$ is strictly increasing in b_i , so $\pi^B(p^A, b_i)$ is strictly increasing in b_i). Therefore, if any buyer finds it optimal to play the high-low strategy, so will buyer type $\check{b} > b_i$, which contradicts our supposition. Therefore, if any buyer finds it optimal to play the high-low strategy, and if p^A is higher than the lowest buyer type who would find it optimal to play the high-low strategy, then the seller's best response when facing an auction price of p^A at the beginning of a bargaining game would be to immediately accept p^A if $p^A \geq s$ and quit otherwise. \square

B.3 Sharpness of Proposition 2 Bounds for Seller Valuations

Proposition 7. *If the only maintained assumptions are Assumption A7 and Assumption A2, the bounds in Proposition 2 are sharp.*

Proof. Suppose Assumptions A2 and A7 are satisfied. Let the bounds on F_S be given by $\Pr(D^S = A|P^A = p^A)$ and $\Pr(D^S \neq Q|P^A = p^A)$. I must demonstrate that any F_S lying in these bounds could have generated these exact bounds. The following behavior will generate the bounds exactly. Given an auction price p^A , if $s \leq p^A$ the seller accepts with probability $\frac{\Pr(D^S=A|P^A=p^A)}{F_S(p^A)}$ and otherwise the seller counters or the buyer opts out of bargaining; if $s \geq p^A$ the seller quits with probability $\frac{1-\Pr(D^S \neq Q|P^A=p^A)}{1-F_S(p^A)}$ and otherwise the seller counters or the buyer opts out of bargaining. This behavior generates the bounds exactly. \square

An interesting feature of these bounds is that $\Pr(D^S = A|P^A = v)$ and $\Pr(D^S \neq Q|P^A = v)$ are not necessarily monotonic, but Proposition 7 still holds. This is because these are bounds on a distribution function, and hence they are bounds on a weakly increasing function (i.e. lying in \mathcal{F}). An alternative, monotonized version of these bounds that emphasizes this point is

$$\sup_{v < p^A} \Pr(D^S = A|P^A = v) \leq F_S(p^A) \leq \inf_{v > p^A} \Pr(D^S \neq Q|P^A = v) \quad (16)$$

This ironing-like procedure results in bounds that are monotonic, as illustrated in Figure A1, but this does not result in bounds on F_S that are any tighter than those originally stated in Proposition 2 because F_S is already restricted to be a proper CDF.

Importantly, if these monotonized bounds from (16) *cross* at any point, this would be an indication of a violation of Assumption A7. Such a violation would mean that there is no weakly increasing function lying between the monotonized bounds or between the Proposition 2 bounds.

Proposition 7 demonstrates sharpness if the only maintained assumptions are A2 and A7. However, under the assumption of BNE (maintained throughout the paper, but not actually required for Proposition 2), the bounds are not necessarily sharp because there is no guarantee that the behavior described in the proof of Proposition 7 describes a BNE. Proving sharpness of these bounds is difficult within an equilibrium framework because, as highlighted above, no full characterization of equilibria exists for two-sided incomplete-information bargaining settings in the current theory literature, and hence it is difficult to determine an equilibrium in which the bounds hold with equality. However, either relaxing the assumptions or restricting the available data can lead to sharpness. For example, the behavior could indeed constitute a BNE if the buyer were to commit to never bargain, in which case the seller’s best response would always be to only accept or reject the auction price. In this case the bounds would hold with equality, but such an equilibrium—where buyers never bargain—would be unlikely to describe the data well. In light of these arguments, the bounds in Proposition 2 are conservative.

B.4 Continuation Values from Re-entering the Market

One limitation with the analysis in this paper is that it does not provide a sophisticated model of player’s continuation payoffs when trade fails: the buyer is modeled as receiving a payoff of 0 and the seller receives her valuation S .²⁷ In practice, buyers who fail to acquire a car may choose to later bid on a similar car. This is precisely what the market thickness controls described in

²⁷As highlighted in Section 3.1, this abstraction is only a limitation in that it ignores the possibility that a *permanent* change to a new bargaining mechanism would potentially shift the distribution of types who are willing to participate and also shift the payoff they receive if they choose to re-enter the game when trade fails. My counterfactual analysis is consistent with a one-time change in the bargaining mechanism: holding fixed the distribution of these inter-game continuation payoffs, the counterfactuals analyze the effects of a shift from the current bargaining mechanism to a counterfactual mechanism for the current instance of the game.

Appendix C.1 capture in the observed heterogeneity regression, but only in a reduced-form way through an additively separable effect on prices. The approach I adopt—treating buyers’ outside option as a 0 payoff—means that what I model as the buyer’s valuation is actually the buyer’s value of owning the car minus the buyer’s discounted continuation payoff of re-entering the market. In the third row of Table A6, I restrict the sample to observations in which the number of cars of the same make and model remaining to be sold on a given day is above the median value of this measure. In this subsample of the data, these continuation dynamics may look different than in the full sample, as players have more upcoming opportunities to obtain similar cars. I find that the estimates do differ from the full-sample estimates; in particular in the fleet/lease sample the upper bound on the efficiency loss between the real-world mechanism and ex-post efficiency is lower (\$654 as opposed to \$864 in the main sample). However, overall the qualitative findings in this subsample are similar to those in the main sample.

When the car fails to sell, the seller may take the car back to her own store or lot, or may leave the car at the auction house where another sale attempt will take place, typically one month later. What I model as the seller’s value S is actually the seller’s discounted continuation payoff from re-entering the market. Similarly, a seller selling multiple cars on a given day may treat these cars as a portfolio, and the seller’s willingness to sell one car may be related to her set of remaining cars to be sold. In the fourth row of each panel of Table A6, I restrict the sample to those observations in which the number of cars remaining to be sold by this same seller on a given day is above the median value of this measure. In this subsample I find similar results to the full sample for dealers, and a smaller upper bound on efficiency loss for fleet/lease sellers (\$634, close to the estimate from the third row of Table A6).

In the final row I limit the sample to observations of the seller’s first attempt to sell a given car (i.e. a given VIN). The main estimates in the paper treat each attempt to sell the car as an independent observation (Assumption A10), but if sellers’ outside options are not an independent draw from F_G with each attempt to sell the car, the model estimates might differ for cars with previous failed attempts to sell. In this final row, in both dealers and fleet/lease cars, I find a higher estimated gains from trade and lower efficiency shortfall for the real-world mechanism than in the main sample. Importantly, however, the qualitative implications are the same in every subsample: the real-world bargaining appears to be inefficient, and not solely due to the information constraints highlighted in Myerson and Satterthwaite (1983), as the second-best mechanism achieves a similar range of surplus as ex-post efficiency.

B.5 Sellers Accepting Offers Below the Secret Reserve Price

This section provides a discussion of several explanations rationalizing why some sellers would set a secret reserve price and then, in the bargaining stage, accept an offer below the reserve price. These

explanations include a specific equilibrium example where this behavior can occur; an extension to the main model in which seller's have market-level uncertainty at the time they set their reserve prices; and seller biases due to over-optimism about buyer demand or due to attempts to influence auctioneer effort.

B.5.1 An Equilibrium Example

This section demonstrates that the model in the body of the paper can contain BNE in which sellers accept offers that lie below their previously set reserve prices.²⁸ For simplicity of exposition, assume the seller has a value of $S = 0$ and buyer value B is uniformly distributed on $[0, 1]$.

Suppose the buyer commits to reject all counteroffers when $p^A \geq 1/2$ and the seller refuses to consider any counteroffer after her first counteroffer when $p^A < 1/2$. The seller's optimal secret reserve price must then be at least $1/2$. When $p^A < 1/2$, let $y(p^A)$ denote the optimal take-it-or-leave-it offer for the seller at her first chance to counteroffer, given the realized auction price. Suppose also that for each $p^A < 1/2$, there is a unique counteroffer as a function of p^A , call it $z(p^A)$, that the buyer will accept if $z(p^A) < B$. Define this $z(p^A)$ to be equal to p^A at $p^A = 1/4$, to be equal to $y(p^A)$ outside of $[1/4 - \varepsilon, 1/4 + \varepsilon]$ for some small $\varepsilon > 0$, and to be the linear interpolation of p^A and $y(p^A)$ along this interval:

$$z(p^A) = \begin{cases} \frac{|p^A - 1/4|}{\varepsilon} y(p^A) + \left(1 - \frac{|p^A - 1/4|}{\varepsilon}\right) p^A & \text{if } p^A \in [1/4 - \varepsilon, 1/4 + \varepsilon] \\ y(p^A) & \text{otherwise.} \end{cases}$$

For small enough ε , the seller's optimal reserve price will be $1/2$. Moreover, the seller's best response counteroffer for $p^A \in [0, 1/2]$ and $p^A \notin [1/4 - \varepsilon, 1/4 + \varepsilon]$ is $y(p^A)$. Now consider the seller's response when $p^A \in [1/4 - \varepsilon, 1/4 + \varepsilon]$. Since the buyer will not consider any counteroffer other than $z(p^A)$, the seller's only options are to accept the auction price or to counter at $z(p^A)$. Which option she chooses will depend on whether the price increase from countering is higher than the cost of making a counteroffer. The price increase is $z(p^A) - p^A$, which is non-negative since $y(p^A) > p^A$, while the cost of an making an offer is c_S . Since $y(p^A) - p^A = 0$ for $p^A = 1/4$, for p^A sufficiently close to $1/4$, the seller will accept the auction price, while for p^A further away from $1/4$ she will counter at $z(p^A)$. This is the seller's optimal response given that the buyer rejects any counteroffer other than $z(p^A)$, while the buyer's strategy of accepting only $z(p^A)$ and accepting it if and only if $z(p^A)$ is below his valuation is optimal since the seller never offers a lower price. Thus, in this equilibrium, the seller accepts some auction prices below her reserve price. Note also that this equilibrium example satisfies Assumption A4.

²⁸I thank Evan Storms for this example.

B.5.2 Model Extension with Market-Level Uncertainty

This section extends the model of Section 3 to allow sellers to have uncertainty, at the time they set their reserve prices (which is several days in advance of the actual auction day), about what their own valuation and what the distribution of buyer valuations will be on the day the auction will take place. This uncertainty is then resolved once the auction takes place and the seller sees (or learns over the phone through an auction house employee) additional information, such as the level of buyer turnout/interest, which can be affected by weather, financial news, or other shocks. Such uncertainty can rationalize why some sellers would set a secret reserve price and then accept an auction price below the reserve price.

The existence of such uncertainty would not affect certain key results of the model, such as bidding behavior in the auction. It would also not affect the result that the seller's secret reserve price strategy is strictly increasing, as I demonstrate below.

Let ζ be a finite vector parameterizing a seller's uncertainty about buyer valuations, where ζ is independent of buyer and seller valuations. Let $F_b(\cdot; \zeta)$, $f_b(\cdot; \zeta)$, $f_{PA}(\cdot; \zeta)$, $\chi^{-1}(b; \zeta)$, and $\pi^S(p^A, s; \zeta)$ be equivalent to the analogous objects in the main model but conditional on ζ , and let $u(s; \zeta)$ represent the seller's valuation conditional on the realization of ζ . At $\zeta = 0$, let each of these functions be equal to its counterpart in the main model (so $u(s; 0) = s$, $F_b(\cdot; 0) = F_b(\cdot)$, $f_b(\cdot; 0) = f_b(\cdot)$, $f_{PA}(\cdot; 0) = f_{PA}(\cdot)$, $\chi^{-1}(b; 0) = \chi^{-1}(b)$, and $\pi^S(p^A, s; 0) = \pi^S(p^A, s)$). At the time the seller chooses the reserve price, she knows each of these functions but does not know the realization of ζ . I also assume that $u(s; \zeta)$ is weakly increasing in ζ .

The following argument follows the steps of the proof of Lemma 3 and demonstrates that in this model reserve prices would still be increasing in s . To simplify this proof I assume that the optimal $\rho(s)$ satisfies $f_{PA}(\rho(s); \zeta) > 0$ for any ζ . The seller's expected payoff, prior to knowing the realization of ζ , can be written as

$$\begin{aligned}
& E_\zeta \left\{ \int_\rho^{\bar{b}} p^A f_{PA}(p^A; \zeta) dp^A \right. \\
& \quad \left. + \int_{\underline{b}}^\rho \left[\int_{p^A}^{\chi^{-1}(p^A; \zeta)} u(s; \zeta) f_B(b) db + \int_{\chi^{-1}(p^A; \zeta)}^{\bar{b}} \pi^S(p^A, u(s; \zeta); \zeta) f_B(b; \zeta) db \right] \frac{f_{PA}(p^A; \zeta)}{1 - F_B(p^A; \zeta)} dp^A \right\} \\
& = E_\zeta \left\{ \int_\rho^{\bar{b}} p^A f_{PA}(p^A; \zeta) dp^A \right. \\
& \quad \left. + \int_{\underline{b}}^\rho \left[u(s; \zeta) (F_B(\chi^{-1}(p^A; \zeta); \zeta) - F_B(p^A; \zeta)) + \pi^S(p^A, u(s; \zeta); \zeta) (1 - F_B(\chi^{-1}(p^A; \zeta); \zeta)) \right] \right. \\
& \quad \left. \times \frac{f_{PA}(p^A; \zeta)}{1 - F_B(p^A; \zeta)} dp^A \right\}
\end{aligned}$$

Differentiating the above expression using Leibniz Rule yields the following first-order condition

for the reserve price:

$$\begin{aligned} \frac{\partial}{\partial \rho} = & -\rho E_{\zeta} [f_{PA}(\rho; \zeta)] + u(s; \zeta) E_{\zeta} \left[\frac{F_B(\chi^{-1}(\rho; \zeta); \zeta) - F_B(\rho; \zeta)}{1 - F_B(\rho; \zeta)} \times f_{PA}(\rho; \zeta) \right] \\ & + E_{\zeta} \left[\pi^S(\rho, u(s; \zeta); \zeta) \frac{1 - F_B(\chi^{-1}(\rho; \zeta); \zeta)}{1 - F_B(\rho; \zeta)} \times f_{PA}(\rho; \zeta) \right] \end{aligned}$$

Conditional on a realization of ζ , Lemma 1 applies, and hence $\pi^S(p^A, u(s; \zeta); \zeta)$ is weakly increasing in s for any realization of ζ . Also, conditional on a realization of ζ , $\chi^{-1}(p^A; \zeta) > p^A$ by the same arguments as in Proposition 3, and thus $F_B(\chi^{-1}(\rho; \zeta); \zeta) > F_b(\rho; \zeta)$. Combining these results demonstrates that $\frac{\partial}{\partial \rho}$ will be strictly increasing in s , and thus the Edlin and Shannon (1998) Theorem implies that the reserve price will be strictly increasing in s for a given realization of ζ , and taking expectations over ζ yields the desired result.

However, even simple versions of this market uncertainty model would likely not be empirically tractable. Suppose buyer and seller valuations are additively separable in ζ , such that $u(s; \zeta) = s + \zeta$, $F_b(b; \zeta) = F_{b+\zeta}(b)$, etc., and suppose ζ can be estimated in the data (such as through a location-by-date fixed effect). In this case, the arguments in Proposition 5 imply that players' continuous decisions taking place after ζ is realized (the auction price and bargaining offers) will be shifted by ζ , but the secret reserve price, set before ζ is realized, will be shifted instead by $E[\zeta]$. The advantage of Proposition 5 in the main model is that it allows me to account for game-level heterogeneity in steps 1–2 and then ignore it in subsequent steps. This would not be feasible in this alternative model because some later estimation steps would require carrying around the realization of ζ . For example, the function determining the probability of trade, $g(\cdot)$, is a function of r , b , and p^A , all of which are additively separable in game-level heterogeneity in the main model; in this additively separable market uncertainty model, however, b and p^A are additively separable in ζ but r is not.

B.5.3 Optimistic Beliefs and Influencing Auctioneers

Sellers accepting prices below their secret reserve price can also be rationalized by sellers having overly optimistic beliefs about auction prices prior to the auction taking place (Treece 2013) or by an attempt to influence auctioneers to exert greater effort to achieve higher prices (Lacetera, Larsen, Pope, and Sydnor 2016; Treece 2013). Such situations can be modeled as the sellers choosing a reserve price given by the optimal, un-biased reserve price, $\rho^*(s)$, plus a bias term, $h(s)$, that is weakly increasing in s . Under such a weakly increasing bias assumption, the observed reserve prices would be strictly increasing in the underlying seller valuation.²⁹

²⁹Lacetera, Larsen, Pope, and Sydnor (2016) demonstrated that who gets assigned as an auctioneer for a given sale can affect the probability of sale, the speed of the sale, and, to a lesser extent, the auction price, and Coey, Larsen, and Sweeney (2014) demonstrated that this assignment can have meaningful effects on seller revenue. I do not model any role of human auctioneers in this paper. Lacetera, Larsen, Pope, and Sydnor (2016) also provided evidence in

B.6 Secret vs. Public Reserve Prices

I present here a discussion of public vs. secret reserve prices. Consider a modified version of the game consisting of no bargaining and only a public reserve price set optimally by the seller prior to the auction. Such a modified game would be equivalent to an auction with no reserve price followed by a bargaining game in which the seller makes a take-it-or-leave-it offer to the high bidder (see Menezes and Ryan 2005). In this modified setup, the seller’s valuation would be related to the optimal public reserve price (R_P) and to the distribution of buyer valuations according to

$$S = R_P - \frac{1 - F_B(R_P)}{f_B(R_P)} \quad (17)$$

The existing literature has provided a number of possible explanations for why some real-world auctions include secret instead of public reserve prices. The findings of my paper are most closely related to two previous papers, Elyakime, Laffont, Loisel, and Vuong (1994, 1997), both studying timber auctions in France, which follow a format of a secret reserve auction followed by post-auction bargaining (but the auction is a first price auction instead of an ascending auction). An important point raised in the first paper is that the use of secret reserve prices may be motivated by a goal to increase the sellers’ *revenue* rather than seller’s *profits* (where the latter includes the seller’s private value). Using their structural model estimates, the authors found that the expected revenue is 22% higher in the secret reserve auction than it would be in a public reserve auction. The authors state, “If this superiority in expected sales of the secret reservation price was general, then this might give a plausible explanation for the frequent use of auctions with secret reservation price since the organizers of the auctions are paid a percentage of sales.” Wholesale auto auctions have a similar feature where the auction house is only paid when trade occurs; as explained in Appendix D, the fee in my setting is primarily a fixed fee. My results in Tables 4–5 suggest that the seller-optimal mechanism can potentially have a lower probability of trade than the real-world secret-reserve-auction-plus-bargaining mechanism. This is consistent with the findings and conjecture of Elyakime, Laffont, Loisel, and Vuong (1994): the use of secret reserve prices may be motivated by auction house incentives to increase trades.

An important point raised in the second paper, Elyakime, Laffont, Loisel, and Vuong (1997), is that post-auction bargaining “is a consequence of the inability of the seller to commit to a secret

favor of random assignment of auctioneers in their data. Under such random assignment, ignoring these auctioneer effects in my model would have no effect on the model’s interpretation if the auctioneer primarily affects the number of bidders N or the bidders’ valuations. This is because N , B , and S are already considered to be mutually independent in the model, and random auctioneer assignment would serve as one source of the random variation in B and N . If, however, some auctioneers are better than others at actually influencing sellers to *change* their valuations then this would not fit easily into the current model, because the reserve price is chosen before any auctioneer influence. Lacetera, Larsen, Pope, and Sydnor (2016) provided survey evidence that sellers are not influenced by auctioneers in this fashion; instead, the authors find survey and quantitative evidence that better auctioneers create more excitement among bidders, which could be consistent with influencing N or B .

or even public reservation price.” A secret reserve price itself may also reflect a seller’s inability to commit to the optimal bargaining outcome (i.e. a public reserve price). At wholesale auto auctions, the auction house does not provide an avenue for sellers to commit to or announce a reserve price before the auction, and the auction house calls up the seller when the secret reserve price is not met, also reducing a seller’s commitment power. Here I present evidence that the current secret reserve price mechanism leads to higher reserve prices than would a public reserve auction, potentially due to this lack of commitment power in the current mechanism.

A simple way to compare secret vs. public reserve prices is to compare the seller valuations estimated in the body of the paper to those that would be implied from naively treating the secret reserve prices in the data as optimally set *public* reserve prices.³⁰ The distribution of these implied values can be computed by simply plugging in draws from \hat{F}_R into (17). Figure A2 displays the results of this comparison. The results in each panel indicate that the seller valuations estimated in the body of the paper are, for the most part, lower than those that would be inferred from treating reserve prices as optimal public reserve prices. Equivalently, the secret reserve prices observed in the data appear to be for the most part *higher* than optimal public reserve prices would be. This may arise because, in the current mechanism, when sellers set a high reserve price and the auction price falls short of the reserve price, sellers can still have the option to accept the auction price or take other bargaining actions, whereas in the public reserve setting the seller must commit to a non-negotiable reserve price. Thus, the downside to setting a high secret reserve price may be smaller than the downside to setting a high public reserve price. Along these lines, Kim (2013), Coey, Larsen, Sweeney, and Waisman (2018), and others have noted the asymmetric losses from choosing a non-optimal public reserve price: the losses in expected revenue from setting *too high* of a public reserve price can be large, and are much larger than the losses from setting a public reserve price *too low*. Industry commentators suggest that sellers’ reserve prices at wholesale auto auctions do indeed appear to be surprisingly high (Treece 2013). The exercise presented here suggests that this feature may be driven at least in part by the fact that the auction mechanism is a secret reserve auction followed by bargaining, where sellers do not commit ex-ante to a publicly announced reserve price.

³⁰Note that the seller CDF bounds estimated in the body of the paper do not rely on the distribution of secret reserve prices other than through the inequality $R \geq S$.

C Technical Details on Estimation Steps

C.1 Additional Details on Sample Restrictions and Controlling for Game-level Observables

C.1.1 Sample Restrictions

Table A2 describes the order in which I impose sample restrictions and also the number of observations dropped due to each restriction. Panel A describes the first restrictions I impose using the full dataset, before splitting into dealer and fleet/lease sales, and panel B describes restrictions I impose separately for the dealers and fleet/lease samples. The full sample contains 1,008,847 runs of vehicles. I first drop observations having missing values for the auction house blue book estimate or odometer reading (158,520) or the timestamp of the attempted sale (21,246). These missing values, and other misrecorded observations described below, are in many cases an indication that a planned run of the vehicle did not actually take place. I drop 126 observations with misrecorded Vehicle Identification Numbers (VINs), such as those having all zeros or less than the correct number of characters for a VIN, and 4 duplicate observations at the VIN-seller-time level. I also drop 1,578 observations in which the same car is recorded as having been sold by one seller, then another seller, and then the original seller; I drop these because I do not trust that they are recorded correctly. I drop 27,027 observations with *both* the auction price and the reserve price missing because these cannot be used in my analysis. I compute each car’s age as the difference between the model year and the year corresponding to the date the vehicle is run. I drop 5,855 observations that are very old (age > 16 years) or very new (i.e. next year’s model, in which case age < 0) or that have less than 100 miles or greater than 300,000 miles on the odometer.

I drop a number of atypical observations (18,309) in which the data contains conflicting indicators of whether the vehicle sold (e.g. the vehicle was recorded as having sold but no sales date is recorded or the bargaining sequence suggests the agents ended in disagreement; or the vehicle was recorded as having not sold but the auction price exceeded the reserve price). I also take several steps to clean the bargaining sequences. I drop observations (10,269) in which bargaining began through a bidder other than the high bidder contacting the auction house with an offer immediately after a failed auction. In 2,595 observations, this same process occurs but with multiple bidders reporting what are referred to as “back-up offers” to the auction house in case trade fails with the current buyer in bargaining. These offers and back-up offers occur infrequently and are not the primary methods of sale. I also drop 179 observations with inexplicable bargaining sequences (e.g. the buyer accepts and then the seller counters); 150 observations in which the auction price appears to have missing or extra zeros; 1,423 observations in which the sale price in the data is not equal to the auction price or to the final negotiated price; and 9 observations in which the final price is below the auction price. After all of these drops described in panel A, 761,557 observations remain,

427,607 of which are dealer sales and 333,950 of which are fleet/lease sales.

I then impose a number of restrictions separately for the dealers and fleet/lease samples, shown in panel B. I first drop observations for which the following variables lie outside their respective 0.01 and 0.99 percentiles: blue book price, final price, auction price, and reserve price. Panel B lists the number of observations dropped in each case along with the particular cutoff values of the cutoff defining the sample restriction (i.e. the 0.01 or 0.99 quantiles of the corresponding variable). I also drop a small number of observations in which the later bargaining offers (i.e. beyond the auction price) take on extreme values (lying outside the 0.01 or 0.99 quantiles of the auction price). Finally, I drop observations from days on which fewer than 100 cars were offered for sale at a given auction house on that day and observations in which fewer than ten vehicles were observed at a given make-model-year-trim-age combination. In the end I am left with 300,740 observations in the dealers sample and 211,656 observations in the fleet/lease sample. I will refer to these samples as the regression samples.

Some of these remaining observations record only a secret reserve price or only an auction price but not both. These observations are not suitable for my final analysis (which requires observing both the auction price and the reserve price for a given attempted sale) but are still useful in controlling for observable heterogeneity, and therefore I include these observations in the *regression sample* used in the step 1 regression but not in the subsequent estimation steps. Panel C of Table A2 provides counts for these observations. In the end, I am left with 133,523 dealer observations and 131,443 fleet/lease observations that can be used in the full estimation procedure. I will refer to these samples as my final samples.

I remark here briefly on these missing reserve prices and missing auction prices shown in panel C of Table A2. Missing reserve prices typically occur when the seller chooses not to report a reserve price, either planning to be present at the auction sale to accept or reject the auction price in person or planning to have the auction house call her on the phone rather than determining a reserve price a priori. Missing auction prices can indicate that a planned vehicle run actually never took place; the auction house does not always have time to run every vehicle it plans to on a given day. Missing auction prices can also arise from the descending/ascending practice of auctioneers: auctioneers do not start the bidding at zero; they start the bidding high and then lower the price until a bidder indicates a willingness to pay, at which point the ascending auction begins. If bidders are slow to participate, the auctioneer will cease to lower bids and postpone the sale of the vehicle until a later date, leaving no auction price recorded. See Lacetera, Larsen, Pope, and Sydnor (2016).

If sellers who choose not to report a reserve price have much higher valuations than other sellers, then my results may be overstating the gains from trade (and the opposite is true if these missing reserve prices correspond to low-value sellers). A similar argument can be made for missing auction prices leading to an under- or over-statement of the gains from trade. To explore this possibility, I estimate the probability of trade in the 19,265 observations in the dealers regression sample and

in the 65,393 observations in the fleet/lease regression sample in which the reserve price is missing (see panel C of Table A2. I compare these numbers to the overall probability of trade in the final samples, shown in Table 1. In the dealers sample, I find that the probability of trade conditional on a missing reserve price is 0.691, compared to 0.705 in the final sample. In the fleet/lease sample, the corresponding numbers are 0.796 and 0.768, suggesting the opposite for fleet/lease cars. In both cases, the numbers are quite close, suggesting that selection introduced by missing reserve prices may not be a major concern. For missing auction prices, no similar analysis can be performed because missing auction prices always correspond to a no-sale.

C.1.2 Observable Heterogeneity Regressions

In the regressions controlling for observable covariates, I include the following controls in x_j :

1. Fifth-order polynomial terms (all degrees of the polynomial from one through five) in the auction houses' blue-book estimate and the odometer reading
2. The number of previous attempts to sell the car; the number of pictures displayed online; a dummy for whether or not the odometer reading is considered accurate, and the interaction of this dummy with the odometer reading; the interaction of the odometer reading with car-make dummies
3. Dummies for each make-model-year-trim-age combination (where age refers to the age of the vehicle in years); dummies for condition report grade (ranging from 1-5, observed only for fleet/lease vehicles); dummies for the year-month combination and for auction house location interacted with hour of sale; dummies for 32 different vehicle damage categories recorded by the auction house; and dummies for each seller who appears in at least 500 observations
4. Dummies for discrete odometer bins: four equally sized bins for mileage in $[0, 20000)$; eight equally sized bins for mileage in $[20000, 80000)$; four equally sized for mileage in $[100000, 200000)$; one bin for mileage in $[200000, 250000)$; and one bin for mileage greater than 250000.
5. Several measures of the thickness of the market during a given sale, computed as follows: for a given car on a given sale date at a given auction house, I compute the number of remaining vehicles still in queue to be sold at the same auction house on the same day lying in the same category as the car in consideration. The six categories I consider are make, make-by-model, make-by-age, make-by-model-by-age, age, or seller identity.
6. Controls for the run numbers, which represents the order in which cars are auctioned. I include fifth-order polynomials for both the run number within an auction-house-by-day combination, and the run number within an auction-house-by-day-by-lane combination.

The market thickness measures described in bullet 5 above help control for inter-auction dynamics in a reduced-form way; they control for the level of possible alternative sales to which a buyer might substitute if he fails to win the current sale. Similarly, for the seller, these controls help account for how the existence of future market opportunities affects the current transaction. See discussion in Section 3 and Appendix B.4.

For the dealers sample, overall the right-hand side variables in the regression includes 11,285 make-model-year-trim-age category effects and 313 other covariates; for the fleet/lease sample the corresponding numbers are 7,314 and 284. The adjusted R^2 from this first-stage regression is 0.95 in the fleet/lease sample and 0.93 in the dealers sample, implying that most of the variation in auction prices and reserve prices is explained by observables.

C.1.3 Subsample Analyses

Table A6 displays bounds on the expected gains from trade using several different subsamples of the data. The first row in each panel uses only observations of cars with below-median blue book value (less-expensive cars). I find that the gains from trade and the efficiency loss are lower than in the full sample (Tables 4–5). The second row examines cars with below-median age (newer cars), where I find larger gains and efficiency shortfall than in the main estimates. The main estimates of the paper treat observable characteristics X , which includes the blue book value and car age, as independent of other random variables in the game, such as buyer and seller valuations. Table A6 demonstrates that this is indeed an abstraction, as the estimates do differ in these subsamples. For less expensive cars, the estimated gains from trade and efficiency shortfall are lower than in the full sample. For newer cars, the numbers are quite similar to those in the full sample. Importantly, the qualitative findings of the paper are robust: the ex-ante and ex-post efficient outcomes are close to one another, and the real-world bargaining falls short of efficiency. Appendix B.4 discusses the other results in Table A6.

C.2 Technical Details on Estimation for Unobserved Game-Level Heterogeneity

Nonparametric identification of f_R , f^{P^A} , and f_W follows from immediately arguments used in Li and Vuong (1998), Krasnokutskaya (2011), Evdokimov and White (2012), Freyberger and Larsen (2017), and elsewhere; I rely specifically on the assumptions of Evdokimov and White (2012):

Proposition 8. *By independence of R , P^A , and W , and by Assumption A8, the joint density of $R + W$ and $P^A + W$ nonparametrically identifies the marginal densities of R , P^A , and W .*

Proof. This result follows by Lemma 1 of Evdokimov and White (2012). □

The estimation details are as follows. For each random variable $Y \in \{W, R, P^A\}$, the density of Y is approximated by normalized orthogonal Hermite polynomials, $f_Y(y) \approx$

$\frac{1}{\sigma_Y} \left(\sum_{k=1}^K \theta_k^Y H^k \left(\frac{y - \mu_Y}{\sigma_Y} \right) \right)^2 \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{y - \mu_Y}{\sigma_Y} \right)^2}$, where K is the number of elements in θ_Y ; θ_Y , μ_Y , and σ_Y are parameters to be estimated for each $Y \in \{W, R, P^A\}$; and H^k are Hermite polynomials, defined recursively by $H^1(y) = 1$, $H^2(y) = y$, and $H^k(y) = \frac{1}{\sqrt{k}} [yH^{k-1}(y) - \sqrt{k-1}H^{k-2}(y)]$ for $k > 2$. A Hermite polynomial approximation for the densities, as advocated by Gallant and Nychka (1987), has the advantage of being flexible and parsimonious, and also yields an approximation of the CDFs F_W , F_R , and F_{P^A} as functions of the same parameters as the densities.

I maximize the likelihood in (2) subject to the constraint that, for each $Y \in \{W, R, P^A\}$, $\sum_{i=1}^K (\theta_i^Y)^2 = 1$, which ensures that each approximated function is indeed a density function, and also subject to the constraint $E[W] = 0$. I choose to include $K = 5$ terms.³¹ The location and scale parameters $\{\mu_Y, \sigma_Y\}_{Y \in \{W, R, P^A\}}$ are not required but improve the performance of the estimator and are standard in estimation with Hermite polynomials. I estimate these parameters in an initial step, maximizing (2) with each density f_Y being approximated by a $N(\mu^Y, \sigma^Y)$. I then plug in the estimated values of $\{\hat{\mu}_Y, \hat{\sigma}_Y\}_{Y \in \{W, R, P^A\}}$ into (2) and maximize the likelihood to obtain consistent estimates of $\{\theta^Y\}_{Y \in \{W, R, P^A\}}$. Consistency of this estimation approach is discussed in Gallant and Nychka (1987), Freyberger and Larsen (2017), and elsewhere.

Throughout the paper, for any estimator requiring integration against the density f_W , such as in (2), I perform this integration using Gauss-Hermite quadrature with 10 nodes. For any univariate function $g(\cdot)$, Gauss-Hermite quadrature is given by $\int_{-\infty}^{\infty} g(v)dv \approx \sum_{k=1}^{K^{GH}} g(x_k) e^{x_k^2} w_k$, where K^{GH} is the number of nodes and x_k and w_k are the Gauss-Hermite quadrature nodes and weights described in Judd (1998).

C.2.1 Estimates from Unobserved Heterogeneity Step

Figure A3 displays the resulting estimated CDFs. Table A5 shows, in panel A, the variance for each of the independent components that sum up to the raw reserve prices and auction prices: $X'\gamma$, W , and R and P^A . As implied by the high adjusted R^2 measures reported in Appendix C.1, over 93% of the variation in auction prices and reserve prices is explained by observable X . In the dealers sample, of the residual variation in reserve prices after the step 1 regression, 59% comes from the variance in W and the remainder from the variance in R (computed by $Var(W)/(Var(W)+Var(R))$ using draws of R and W from the estimated F_R and F_W). In the fleet/lease sample, the unobserved heterogeneity component W only explains 34% of the residual variance. For auction prices, the proportion of residual auction prices explained by variance in W is 53% in the dealers sample and 33% in the fleet/lease sample. These numbers suggest that unobserved heterogeneity especially plays a large role for dealer cars.

³¹The above framework can be treated as a semi-nonparametric maximum likelihood setting, letting K grow appropriately with the sample size and choosing K through cross-validation. I instead fix $K = 5$, treating this as a flexible parametric approximation, as suggested in Kim and Lee (2014) and Freyberger and Larsen (2017). I find that choosing K larger than 5 does not affect estimates noticeably.

One measure of the fit of the maximum likelihood procedure performed in step 2 is shown in panel B of Table A5, where I display the correlation of the raw reserve prices and auction prices compared with the correlation of $R + W$ and $P^A + W$ simulated from their estimated distributions F_R , F_{P^A} , and F_W . In the dealers sample, these numbers are 0.513 in the raw residuals vs. 0.557 in the simulation, and in the fleet/lease sample these numbers are 0.293 vs. 0.341. In each sample the correlations differ by less than 0.05. A graphical evaluation in a similar vein is shown in Panels A and B of Figure A5, where I plot the empirical CDF of the difference between the raw variables, $R^{raw} - P^{A,raw}$, along with the simulated CDF of $R - P^A$, with R and P^A drawn independently from the estimated F_R and F_{P^A} . In both panels A and B the two CDFs are indistinguishable, suggesting that the fit is quite good.

C.2.2 Evaluating Additive Separability

I evaluate here the appropriateness of my additively separable model of heterogeneity. The variables R , P^A , and W are assumed to be independent and are not restricted to have positive support and are added to $X'\gamma$ (which they are assumed independent of) to yield the raw, observed reserve prices and auction prices. Therefore, nothing prevents my estimates from suggesting that some reserve prices and auction prices might be negative (even though in the data all raw prices are non-negative). However, in practice, the observable heterogeneity term $X'\gamma$ has most of its mass far above zero, and the random variables R , P^A , and W have most of their mass lying in a small interval around zero, and thus the sums $X'\gamma + W + R$ and $X'\gamma + W + P^A$ have the majority of their mass being positive. I demonstrate this in panel A of Table A5, which displays the 0.1 to 0.9 quantile range of $X'\hat{\gamma}$ and of the estimated distributions of R , P^A , and W . Even at the 0.1 quantile of $X'\hat{\gamma}$, the negative 0.1 quantiles for the other components of prices are small enough in magnitude that the model does not predict large negative observed prices. To evaluate this point further, I simulate from the estimated distributions (taking draws of $X'\gamma$ from the empirical distribution of the predicted value $X'\hat{\gamma}$ and taking draws of R , P^A , and W from the estimated F_R , F_{P^A} , and F_W). I find that this exercise predicts negative values of $X'\gamma + W + R$ with probability 0.008 in the dealers sample and 0.005 in the fleet/lease sample; the corresponding predictions for negative values of $X'\gamma + W + P^A$ are slightly larger but still small (0.042 and 0.010). Thus, the additively separable model I use is at best an approximation—because in reality the raw, observed prices are never negative—but appears to be a good approximation.

C.3 Technical Details on Buyer Distribution Estimation

Proposition 9. F_B is nonparametrically identified by F_{P^A} and $\Pr(N = n)$.

Proof. This result follows immediately from the fact that the right-hand side of (3) is strictly monotonic in $F_B(\cdot)$ (see, for example, Athey and Haile 2007). \square

To estimate $F_B(\cdot)$, I first replace $F_{PA}(v)$ in (3) with the maximum likelihood estimate $\hat{F}_{PA}(v)$ from (2) evaluated on a grid of values for v . I estimate $\widehat{\Pr}(N = n)$ using the empirical probability mass function of the lower bound of the number of bidders from the bid logs subsample discussed in Section 2. This treats the distribution of the lower bound as though it is the true distribution of the number of bidders. I discuss this further below in Appendix C.3.1. I also discuss evidence below in Appendix C.3.2 consistent with the underlying assumption that N is independent of buyer valuations.

With $\widehat{\Pr}(N = n)$ and $\hat{F}_{PA}(v)$ in hand, I then obtain $\hat{F}_B(v)$ by numerically solving (via a bisection method) for the value u such that

$$0 = \hat{F}_{PA}(v) - \sum_n \widehat{\Pr}(N = n) [nu^{n-1} - (n-1)u^n] \quad (18)$$

Consistency of this type of order statistics inversion estimator is shown in Menzel and Morganti (2013). The error in this estimation procedure can be made arbitrarily small—up to machine precision and the estimation error in F_{PA} and $Pr(N = n)$ —as this procedure corresponds to numerically solving for zeros of an identity.

C.3.1 Robustness to Distribution of Number of Bidders

In order to guide the choice of $\Pr(N = n)$, the distribution of the number of bidders, I first manually collected additional data by visiting multiple auction house locations and physically observing over 200 auctions. For each auction sale, I recorded the number of bidders who appeared to be actively participating or interested in the car. The mean of these observations, conditional on cases where at least two bidders appeared to be active, was 2.62, close to the mean of the lower bound on the number of bidders from the bid log sample (2.924 and 2.973 in the dealers and fleet/lease samples).

However, I also present evidence here that the key estimates in the paper are not sensitive to how $\Pr(N = n)$ is specified. In doing so I compare several possibilities for $\Pr(N = n)$. The first is that which is used in the body of the paper, which is the empirical probability mass function of the lower bound on the number of bidders in each auction in the bid log sample. Let this random variable be denoted \underline{N} . The next is the empirical probability mass function of a (very conservative) upper bound on the number of participants in each auction in the bid log sample. This upper bound comes from adding the total number of bidders who signed in through the online portal for a given lane on a given day (I observe this number in the bid log data, whether or not these signed-in bidders placed any bids) to the total number of floor *bids* (physically present bids, as described in Section 2); thus, this treats each floor bid as having come from a distinct bidder. Let this random variable be denoted \overline{N} . The mean of this upper bound is 15.64 in the dealers bid log subsample and 25.99 in the fleet/lease bid log subsample. An additional possibility is that N is drawn from some parametric distribution, such a Poisson, Negative Binomial, etc. Here I consider

cases in which N follows a Poisson distribution with mean $\lambda \in \{3, 7, 10, 20\}$, conditional on $N \geq 2$ (thus, these latter four are truncated Poisson distributions).

I find that the qualitative findings of the paper are not sensitive to the approximation chosen for $\Pr(N = n)$, and the estimates themselves do not vary drastically. In Table A7, for each of these choices of $\Pr(N = n)$, I display bounds on the expected gains from trade and probability of trade. Panel A shows bounds for ex-post efficiency, the second-best mechanism, and the real-world bargaining, and bounds on the difference between ex-post efficiency and the real-world mechanism. The estimated expected gains are lower under specifications of $\Pr(N = n)$ that place more mass on higher N . Panel B displays similar results for the probability of trade and also demonstrates that the estimated probability of trade is lower under specifications for $\Pr(N = n)$ that place more mass on higher N .³² Both panels indicate only small changes in the estimated bounds across different specifications for $\Pr(N = n)$.

I now present a discussion of why these welfare measures are relatively insensitive to this choice. The welfare measures I evaluate depend on the estimated distribution of buyer valuations, but in a particular way: they depend primarily on the distribution of the valuation of the highest bidder (the buyer who bargains) *conditional* on the auction price, integrated against the auction price density; this yields the distribution of the *maximum order statistic* of buyer valuations. To see this, note that the maximum order statistic distribution (averaged over values of N) can be computed as

$$F_{B^{(1)}}(v) = \sum_n \Pr(N = n) F_B(v)^n, \quad (19)$$

or, alternatively, it can be computed from the density of the maximum order statistic conditional on the second order statistic integrated against the density of the second order statistic, given by

$$F_{B^{(1)}}(v) = \int_{\underline{b}}^v \int_{\underline{b}}^y = \frac{f_B(y)}{1 - F_B(p^A)} f_{P^A}(p^A) dp^A dy. \quad (20)$$

Each of the welfare measures evaluated in this paper (described in Section 5) depend on an object similar to the latter formulation, although the relationship between (20) and what I evaluate is not exact—for example, the denominator in my welfare measures is $1 - F_B(\chi^{-1}(p^A))$, and in estimating $\chi^{-1}(\cdot)$ and $g(\cdot)$ in estimation steps 5–6 terms appear that are not exactly the same as maximum order statistic distributions.

Figure A6 demonstrates that the distribution of the *maximum order statistic* is much less sensitive than the underlying distribution F_B to the choice of $\Pr(N = n)$. Panels A and B show

³²To create Table A7, I run estimation step 3 under a given specification of $\Pr(N = n)$, yielding an estimate of F_B , and then I use this new F_B in re-doing estimation steps 5–6 and in computing counterfactuals. I do not estimate bounds on the probability of trade for the second-best mechanism in Table A7, as calculating bounds on this quantity is computationally intensive, as described in Proposition 6 and Table 3.

the estimates of the underlying distribution, F_B , which differ widely as $\Pr(N = n)$ changes. Panels C and D show the distribution of the maximum order statistic, obtained by using $F_{PA}(v)$ and $\Pr(N = n)$ to obtain $F_B(v)$ by solving (3) and then using this $F_B(v)$ and $\Pr(N = n)$ to obtain $F_{B(1)}(v)$ using (19) or (20). Panels C and D demonstrate that the distribution of the maximum order statistic is quite insensitive to the choice of $\Pr(N = n)$ used in this procedure.

For certain choices of $\Pr(N = n)$, including the Poisson distribution, it is possible to prove analytically that the computed maximum order statistic distribution will be entirely insensitive to $\Pr(N = n)$. I state this as the following proposition.³³ Below, note that $F_{B(2)}$ denotes the second order statistic distribution (averaged over values of N), which is equivalent to F_{PA} in the body of the paper.

Proposition 10. *If $\Pr(N = n)$ is given by a Poisson distribution with parameter λ , and $F_{B(2)}$ is known, then, at any point v , the maximum order statistic distribution $F_{B(1)}(v)$ obtained by the following two steps will have zero derivative with respect to λ : 1) use $F_{B(2)}(v)$ and $\Pr(N = n)$ to obtain $F_B(v)$ by solving (3); 2) use this $F_B(v)$ and $\Pr(N = n)$ to obtain $F_{B(1)}(v)$ using (19) or (20).*

Proof. When $\Pr(N = n)$ is a Poisson with mean λ , it can be shown that $F_B(v) = \psi(F_{B(2)}(v); \lambda)$, where $\psi(F_{B(2)}(v); \lambda)$ is defined implicitly as the solution to

$$F_{B(2)}(v) = (1 + \lambda(1 - \psi(F_{B(2)}(v); \lambda)))e^{\lambda(\psi(F_{B(2)}(v); \lambda) - 1)} \quad (21)$$

and the maximum order statistic distribution is given by

$$F_{B(1)}(v) = e^{\lambda(\psi(F_{B(2)}(v); \lambda) - 1)} \quad (22)$$

Holding fixed v and $F_{B(2)}(v)$, implicitly differentiating (21) with respect to λ yields

$$0 = \left(1 - \psi(F_{B(2)}(v); \lambda) - \lambda \frac{d\psi(F_{B(2)}(v); \lambda)}{d\lambda}\right) F_{B(1)}(v) + (1 - \lambda(1 - \psi(F_{B(2)}(v); \lambda))) \frac{dF_{B(1)}(v)}{d\lambda} \quad (23)$$

where I take into account that $F_{B(1)}$ will depend on λ . Note that

$$\frac{dF_{B(1)}(v)}{d\lambda} = e^{\lambda(\psi(F_{B(2)}(v); \lambda) - 1)} \left(\psi(F_{B(2)}(v); \lambda) - 1 + \lambda \frac{d\psi(F_{B(2)}(v); \lambda)}{d\lambda} \right) \quad (24)$$

Plugging (24) into (23) yields

$$0 = \lambda(1 - \psi(F_{B(2)}(v); \lambda)) \frac{dF_{B(1)}(v)}{d\lambda} \quad (25)$$

³³I thank Zhaonan Qu for help with this proof.

Recall that $\psi(F_{B^{(2)}}(v); \lambda) = F_B(v)$. This expression in (25) must hold at all v , even at v where $F_B(v) \neq 1$. Thus, it must be the case that $\frac{dF_{B^{(1)}}(v)}{d\lambda} = 0$. \square

Proposition 10 should not be misinterpreted as an unconditional statement that the distribution of the maximum order statistic does not depend on λ . Rather, the result demonstrates that, conditional on $F_{B^{(2)}}$, the exercise of inverting $F_{B^{(2)}}$ to obtain F_B , and then computing $F_{B^{(1)}}$ from this F_B , is invariant to λ . It may be possible to prove this mathematical result for a larger class of $\Pr(N = n)$, for, as demonstrated in Figure A6, even at non-Poisson distributed N (such as the distributions of the upper and lower bounds on the number of bidders or the truncated Poisson distributions), $F_{B^{(1)}}$ is very insensitive to the choice of $\Pr(N = n)$.

To my knowledge, Proposition 10 is new to the literature, and may be of some independent interest, as it suggests assumptions under which one can compute the marginal distribution of the maximum order statistic solely from knowledge of the marginal distribution of the second order statistic, without knowing N or even fully specifying the distribution of N .³⁴ In turn, these two objects—the marginal distributions of the first and second order statistics—can then be used to identify many objects of interest in auction settings, such as bidder surplus, seller profits, and optimal reserve prices, as pointed out by Aradillas-Lopez, Gandhi, and Quint (2013).

C.3.2 Valuations Independent of Number of Bidders

The estimation of buyer valuations in Section 4 relies on the assumption (stated in Assumption A1) that buyer valuations are independent of N , the number of bidders. Here I examine the validity of this assumption.

For any $k \leq n$ and any v , let

$$\psi_{k:n}(v) \equiv \frac{n!}{(n-k)!(k-1)!} \int_0^v t^{k-1}(1-t)^{n-k} dt$$

Proposition 2 of Aradillas-López, Gandhi, and Quint (2016) demonstrated the following (modified to the environment of this paper). Suppose buyer values (given by the random variable $\tilde{B} = B+W$) are non-negatively correlated within a given auction, as they will be in an environment of conditionally independent private values with independent, additively separable unobserved heterogeneity. Then for any v and any $n > n'$, valuations being independent of N implies

$$\psi_{n-1:n}^{-1}(F_{\tilde{B}^{n-1:n}}(v)) \geq \psi_{n'-1:n'}^{-1}(F_{\tilde{B}^{n'-1:n'}}(v)) \quad (26)$$

where, for any n , $F_{\tilde{B}^{n-1:n}}$ represents the distribution of the auction price (including unobservable

³⁴Precisely, these assumptions are symmetric, conditionally independent, private values (or independent private values with separable unobserved heterogeneity) with valuations independent of N and with N distributed according to a Poisson distribution with unknown mean λ .

heterogeneity) when n bidders are present. Importantly, Aradillas-López, Gandhi, and Quint (2016) demonstrated that standard models of endogenous entry in auctions, such as those of Samuelson (1985) and Levin and Smith (1994), would violate (26).

I apply this result by performing the inversion in (26) for different values of n in the bid log subsample. For this exercise, I treat the lower bound on the number of bidders in a given auction in this subsample as though it represents the true number of bidders in that auction. Panels E and F of Figure A6 display estimates of $\psi_{n-1:n}^{-1}(F_{\tilde{B}^{n-1:n}}(\cdot))$ for the most prevalent values of n observed in the bid log subsample. In both the dealers and fleet/lease samples, a pattern emerges consistent with the inequality in (26) and inconsistent with the models of endogenous entry highlighted in Aradillas-López, Gandhi, and Quint (2016), lending credence to the assumption that valuations are independent of N .

C.4 Technical Details on Estimation of Seller Valuation Bounds

I begin this section with a nonparametric identification proof. Below, and throughout the appendix, for any function $f(v) : \mathbb{R}^K \mapsto \mathbb{R}$, let its Fourier transform be denoted

$$\mathcal{FT}_f(u) = \int f(v)e^{-2\pi i v \cdot u} dv$$

where $v \cdot u$ represents the inner product. When I take a Fourier transform of any function that does not necessarily have full support, I consider its support to be extended such that the function is zero outside of its original support.

Proposition 11. *The functions F_S^L and F_S^U are nonparametrically identified by f_{PA} , f_W , F_R , $\Pr(D^S = A | \tilde{P}^A = v)$, and $\Pr(D^S \neq Q | \tilde{P}^A = v)$.*

Proof. First, note that a solution exists satisfying the constraints in Φ , defined in the sentence below (4); one such solution is $F_S^L = F_S^U = F_R$. Let $(F_{S,0}^L, F_{S,0}^U)$ represent a pair of bounds in Φ such that the objective function in (4) is minimized; $q(v; F_{S,0}^L)$ and $q(v; F_{S,0}^U)$ are then $q(v; \cdot)$ evaluated at this pair of functions. I now demonstrate that these functions $q(\cdot; F_{S,0}^L)$ and $q(\cdot; F_{S,0}^U)$ uniquely determine $(F_{S,0}^L, F_{S,0}^U)$, demonstrating that $(F_{S,0}^L, F_{S,0}^U)$ is identified. Let

$$\bar{q}^L(v) = q(v; F_{S,0}^L) \left(\int M_S(v, z) dz \right) = \int F_{S,0}^L(v-w) f_{PA}(v-w) f_W(w) dw$$

Then, evaluated at $F_{S,0}^L$, the term in (4) that depends on F_S^L can be written

$$\left| \Pr(D^S = A | \tilde{P}^A = v) \left(\int M_S(v, z) dz \right) - \bar{q}(v) \right|^2 \quad (27)$$

The object $\bar{q}^L(v)$ is a convolution of $a(v) \equiv F_{S,0}^{L,0}(v)f_{PA}(v)$ against the density $f_W(w)$. Specifically, for any $u \in \mathbb{R}$,

$$\begin{aligned}\mathcal{FT}_{\bar{q}^L}(u) &= \int \left(\int a(v-w)f_W(w)dw \right) e^{-2\pi i v u} dv \\ &= \int f_W(w) \left(\int a(v-w)e^{-2\pi i v u} dv \right) dw \\ &= \int f_W(w) \left(\int a(y)e^{-2\pi i(y+w)u} dy \right) dw \\ &= \left(\int f_W(w)e^{-2\pi i w u} dw \right) \int a(y)e^{-2\pi i y u} dy \\ &= \mathcal{FT}_{f_W}(u)\mathcal{FT}_a(u)\end{aligned}$$

where $y = v - w$ and $dy = dv$. Rearranging yields

$$\mathcal{FT}_a(u) = \frac{\mathcal{FT}_{\bar{q}^L}(u)}{\mathcal{FT}_{f_W}(u)} \quad (28)$$

By Assumption A8(i), $\mathcal{FT}_{f_W}(u)$ —the characteristic function of f_W —has only isolated zeros, and hence (28) does not divide by zero except on a set of measure zero. Taking the inverse Fourier transform then yields $a(v)$. This function can then be rearranged to solve for $F_{S,0}^L$ by dividing by $f_{PA}(v)$, which is positive at every v in the support of P^A . Thus, $\bar{q}(v; F_{S,0}^L)$ uniquely identifies $F_{S,0}^L$. A similar argument applies for $F_{S,0}^U$. \square

To estimate the bounds on the distribution of seller valuations, I parameterize $F_S^L(\cdot)$ and $F_S^U(\cdot)$ using piecewise linear splines. A spline approximation—or other linear sieve approximation—to an unknown function has the advantage of being linear in parameters while remaining very flexible in fitting the function. Let $\{v_k^S\}_{k=1}^{K_S}$ represent a fixed vector of K_S knots on the support of \tilde{P}^A , and $LS_S(\cdot) : \mathbb{R} \mapsto \mathbb{R}^{K_S}$ be the piecewise linear spline transformation for this vector of knots. An approximation to F_S is then given by $F_S(\cdot; \theta^S) \equiv LS_S(v)' \theta^S$ and an approximation to the density is given by $f_S(v; \theta^S) \equiv LS_S'(v)' \theta^S$, where θ^S is a K_S -by-1 vector of parameters.³⁵ I choose the knots $\{v_k^S\}_{k=1}^{K_S}$ to be uniformly spaced between the 0.001–0.999 quantiles of \tilde{P}^A , and I choose $K_S = 200$ as the number of knots for each bound ($K_S = 250$ or $K_S = 300$ yield similar results, as shown in Panels A and B of Appendix Figure A7).

The lower and upper bounds can then be denoted $F_S^L(v) = LS_S(v)' \theta^{S,L}$ and $F_S^U(v) = LS_S(v)' \theta^{S,U}$, where $(\theta^{S,L}, \theta^{S,U})$ are each K_S -by-1 vectors of parameters to be estimated. I esti-

³⁵Specifically, let $m_k(v) = (v - v_k^S)/(v_{k+1}^S - v_k^S)$ and let $v \in [v_k^S, v_{k+1}^S]$ for some $k \in \{1, \dots, K_S - 1\}$. Then $LS_S(v)$ returns a K_S -by-1 vector with $(1 - m_k(v))$ as the k^{th} element, $m_k(v)$ as the $(k+1)^{\text{th}}$ element, and zeros elsewhere. Thus, letting $\Delta_k = (\theta_{k+1}^S - \theta_k^S)/(v_{k+1}^S - v_k^S)$, $F_S(v; \theta^S)$ is given by $\theta_k^S + (v - v_k^S)\Delta_k$. For the density, $LS_S'(v)$ returns the derivative of the vector $LS_S(v)$ with respect to v , and thus $f_S(v; \theta^S) = \Delta_k$.

mate these parameters by solving the following constrained least squares problem:

$$\min_{\theta^{S,L}, \theta^{S,U}} \sum_{k=1}^{K_S} [\Omega^L(v_k^S; \theta^{S,L})^2 + \Omega^U(v_k^S; \theta^{S,U})^2] \quad (29)$$

where

$$\begin{aligned} \Omega^L(v_k^S; \theta^{S,L}) &\equiv \widehat{\Pr}(D^S = A | \tilde{P}^A = v_k^S) \int (\hat{M}_S(v_k^S, z) dz) - \widetilde{LS}_S(v_k^S)' \theta^{S,L} \\ \Omega^U(v_k^S; \theta^{S,U}) &\equiv \widehat{\Pr}(D^S \neq Q | \tilde{P}^A = v_k^S) \int (\hat{M}_S(v_k^S, z) dz) - \widetilde{LS}_S(v_k^S)' \theta^{S,U} \end{aligned}$$

and where $\widetilde{LS}_S(v_k^S)$ is a K_S -by-1 vector with the ℓ^{th} element given by $\int (LS_S(v_k^S - w)' e(\ell)) \hat{M}_S(v_k^S, w) dw$ and $e(\ell)$ is a K_S -by-1 selection vector with 1 in the ℓ^{th} spot and zeros elsewhere. An advantage, again, of the spline approximation is that this object $\widetilde{LS}_S(v_k^S)$, which requires numerical integration against f_W , can be computed outside the optimization problem; this also simplifies the computation of analytical derivatives to supply to the optimization algorithm. I impose several constraints on the minimum distance problem in (29): (i) F_S^L lies graphically above F_R and graphically below F_S^U ; (ii) F_S^L and F_S^U lie in $[0, 1]$; (iii) F_S^L and F_S^U are weakly increasing; and (iv) $F_S^L(v)$ and $F_S^U(v)$ are equal to 0 for any $v < v_1^S$ and equal to 1 for any $v > v_{K_S}^S$. These last three constraints ensure that $F_S^L(v)$ and $F_S^U(v)$ will correspond to proper distribution functions.

Computing the objective function in (29) requires estimates of \hat{F}_R , \hat{f}_{PA} , and \hat{f}_W , as well as $\widehat{\Pr}(D^S = A | \tilde{P}^A = \tilde{p}^A)$, and $\widehat{\Pr}(D^S \neq Q | \tilde{P}^A = \tilde{p}^A)$. Estimates of the first three objects come from the maximum likelihood procedure in (2). I estimate $\widehat{\Pr}(D^S = A | \tilde{P}^A = \tilde{p}^A)$ using a local linear regression of the event $1\{D^S = A\}$ on realizations of $\tilde{P}^A = \tilde{p}^A$.³⁶ I estimate $\widehat{\Pr}(D^S \neq Q | \tilde{P}^A = \tilde{p}^A)$ analogously. In (29) the combination of a nonparametric first-stage estimate of a nuisance vector (e.g. $\widehat{\Pr}(D^S = A | \tilde{P}^A = v_k^S) \left(\int \hat{M}_S(v_k^S, z) dz \right)$) followed by a parametric second stage estimated through minimum distance falls into the class of two-step semiparametric GMM estimators, discussed in Akerberg, Chen, and Liao (2014) and references therein, and is consistent for estimating $\hat{\theta}^{S,L}, \hat{\theta}^{S,U}$.

For this optimization exercise, I find that the left tail of the bounds, in particular the left tail of F_S^U , can be sensitive to the starting values. To evaluate this sensitivity, I estimated the bounds using starting values of $F_S^L(v) = F_S^U(v) = 1$ everywhere, starting values of $F_S^L(v) = F_S^U(v) = F_R(v)$

³⁶For this local linear regression, as well as the local linear regressions run in estimating $\chi^{-1}(\cdot)$, I use a Gaussian kernel with the rule-of-thumb bandwidth presented in Ruppert, Sheather, and Wand (1995), who suggested computing the bandwidth on the sample lying between the α and $1 - \alpha$ quantiles for some $\alpha \in [0, 1]$ in order to reduce sensitivity to data in the tails. I use $\alpha = .001$. I replace any estimates of conditional probabilities lying outside of $[0, 1]$ with 0 and 1, following Frölich (2006).

As highlighted in Section 4, by construction, the Proposition 2 population bounds will not cross. Empirical estimates of these bounds may cross, however. For example, a local linear approximation to the conditional probabilities in Proposition 2 may cross; I find that such crossing is very minimal in my application. A kernel regression approximation of these bounds (or other local constant approach to approximating a conditional average) will not cross.

everywhere, and 20 additional cases varying uniformly between $F_S^L(v) = F_S^U(v) = 1$ and $F_S^L(v) = F_S^U(v) = F_R(v)$. Panels E and F of Figure A7 show the estimates under the $F_S^L(v) = F_S^U(v) = 1$ starting values in solid lines; the estimates under the $F_S^L(v) = F_S^U(v) = F_R(v)$ starting values in dotted lines; and, in dashed lines, the estimates under the case that achieved the smallest objective value of (29) for the 22 cases tested. Searching over a large number of cases is computationally burdensome, especially when bootstrapping the estimation procedure. I use the $F_S^L(v) = F_S^U(v) = 1$ starting values for my estimation throughout the paper because the objective value under these starting values is very close to that under the dashed-line case.

I now discuss briefly the surjectivity of these bounds in this context as well as in other alternating-offer settings. As explained in Section 3, the revealed preference bounds proposed in this paper can be used in other alternating-offer settings to identify bounds on the valuations of the player who responds to the first bargaining offer, whom I will refer to as the *first responder*. In order for these bounds to be surjective (i.e. provide bounds for the whole range of the CDF from $[0, 1]$) when evaluated on a given range of knots, the data must contain some realizations of the first bargaining offer that are extreme enough that the first responder accepts with probability close to one and some realizations of the first bargaining offer such that the first responder rejects (quits) with probability close to one. If this is not the case in a given dataset, specific institutional details may aid in identifying the support of the bounds, as I now describe.

In examining the surjectivity of the bounds in the current application, I find that there is sufficient variation in the data for the lower bound to be surjective over the range of chosen knots. The estimate of F_S^U reaches a value of 1, but does not fully reach a value of 0; in particular, I find an estimate of $\hat{\theta}_1^{S,U}$ (the lowest spline parameter) of approximately 0.075 in the dealers sample and 0.09 in the fleet/lease sample, meaning that the estimate of F_U^S does not rule out the possibility that 7.5% of sellers (or 9% of sellers in the fleet/lease sample) have values less than or equal to v_1^S , the lowest knot; as explained above, this lowest knot is equal to the 0.001 quantile (approximately the minimum) of the auction price distribution. I therefore treat F_S^U as having a mass point at v_1^S . This mass point could be placed at any arbitrary point weakly below v_1^S . I choose to place it at v_1^S based on Assumption A9, that is, $\underline{s} \geq \underline{b}$, which I believe is a conservative lower bound for the support of seller's values.³⁷

The goodness of fit of the estimated bounds can be evaluated using the integrated absolute

³⁷Surjective estimates of the bounds correspond to the case where $\hat{\theta}_1^{S,U}$ and $\hat{\theta}_1^{S,L}$ are close to 0 and the estimates of $\hat{\theta}_{K_S}^{S,U}$ and $\hat{\theta}_{K_S}^{S,L}$ are close to one. If the bounds are not surjective, the following arguments provide a complete description of $F_S^L(v)$ and $F_S^U(v)$ for points v that lie outside of the support of the chosen knots: (i) for $v > v_{K_S}^S$, $F_S^L(v)$ is given by a horizontal line at $\hat{\theta}_{K_S}^{S,L}$ from $v_{K_S}^S$ to $\hat{F}_R^{-1}(\hat{\theta}_{K_S}^{S,L})$ and by $\hat{F}_R(v)$ above that point; (ii) for $v > v_{K_S}^S$, $F_S^U(v) = 1$; (iii) for $v < v_1^S$, $F_S^L(v) = \hat{F}_R(v)$; (iv) for $v < v_1^S$, $\hat{F}_S^U(v) = 0$ (or, if the assumption that $\underline{s} \geq \underline{b}$ were to be relaxed, $\hat{F}_S^U(v)$ for $v < v_1^S$ would be given by a horizontal line at $\hat{\theta}_1^{S,L}$). If the estimates of F_S^L and F_S^U are surjective, then specifying conditions (i)–(iv) is unnecessary. The only condition that is binding in practice in my application is condition (iv).

error, which I compute as follows for F_S^L :

$$\frac{1}{\sum_{k=1}^{K_S} \hat{M}_S(v_k^S, z) dz} \sum_{k=1}^{K_S} |\Omega^L(v_k^S; \hat{\theta}^{S,L})| \quad (30)$$

I compute an analogous measure for F_S^U , with $\Omega^L(v_k^S; \hat{\theta}^{S,L})$ replaced with $\Omega^U(v_k^S; \hat{\theta}^{S,U})$. These integrated absolute error measures can be interpreted as the average percent by which the conditional probability statement estimated directly from the data fails to be fit by its corresponding model equivalent, where this average is weighted by where the data lies. I find this error to be less than 0.4% for the estimates of F_S^L and F_S^U in both the dealers and fleet/lease samples, as shown in Table A5. Panels C and D in Figure A5 display these estimated conditional probability statements and the corresponding fitted estimates. Table A5 also shows the integrated absolute error restricting to cases that constitute violations of the conditional probability statements in (29), computed as

$$\frac{1}{\sum_{k=1}^{K_S} \hat{M}_S(v_k^S, z) dz} \sum_{k=1}^{K_S} \Omega^L(v_k^S; \hat{\theta}^{S,L}) 1 \left\{ \Omega^L(v_k^S; \hat{\theta}^{S,L}) > 0 \right\}$$

for F_S^L and by

$$\frac{1}{\sum_{k=1}^{K_S} \hat{M}_S(v_k^S, z) dz} \sum_{k=1}^{K_S} -\Omega^U(v_k^S; \hat{\theta}^{S,U}) 1 \left\{ \Omega^U(v_k^S; \hat{\theta}^{S,U}) < 0 \right\}$$

for F_S^U . These quantities are 0.2% or smaller in both samples.

C.5 Technical Details on Estimation of Lower Bound of Support of Types Who Bargain

I begin this section with a nonparametric identification proof.

Proposition 12. *For any function $F_S(\cdot)$ lying in the estimated bounds $[\hat{F}_S^L(\cdot), \hat{F}_S^U(\cdot)]$, the functions $\rho(s)$ and $\rho^{-1}(r)$ are nonparametrically identified by F_R . At every point in the support of P^A and less than the upper bound of the support of R , the function $\chi^{-1}(\cdot)$ in (6) is nonparametrically identified by F_R , f_{PA} , f_W , F_B , and $\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R})$.*

Proof. For any function $F_S(\cdot)$ lying in the estimated bounds $[\hat{F}_S^L(\cdot), \hat{F}_S^U(\cdot)]$, the function $\rho(s)$ can be constructed as $\rho(s) = F_R^{-1}(F_S(s))$, with F_R replaced with the estimated \hat{F}_R from (2). Similarly, $\rho^{-1}(r)$ can be constructed as $\rho^{-1}(r) = F_S^{-1}(F_R(r))$.

Identification of χ^{-1} is more involved. Equation (6) can be written as

$$\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R}) \left(\int M_\chi(\tilde{p}^A, z) dz \right) = \int \frac{1 - F_B(\chi^{-1}(\tilde{p}^A - w)) M_\chi(\tilde{p}^A, w)}{1 - F_B(\tilde{p}^A - w)} dw \quad (31)$$

where $M_\chi(\tilde{p}^A, w) \equiv f_{P^A}(\tilde{p}^A - w)(1 - F_R(\tilde{p}^A - w))f_W(w)$ is the likelihood of the event ($P^A = \tilde{p}^A - w, \tilde{P}^A < \tilde{R}, W = w$). Now let

$$a_\chi(v) \equiv \frac{(1 - F_B(\chi^{-1}(v)))f_{P^A}(v)(1 - F_R(v))}{1 - F_B(v)} \quad (32)$$

The integrand on the right-hand side of (31) is a convolution of $a_\chi(v)$ against the density $f_W(w)$, given by

$$q_\chi(v) = \int a_\chi(v - w)f_W(w)dw$$

By the same argument as in the proof of Proposition 11, for any $u \in \mathbb{R}$,

$$\mathcal{FT}_{a_\chi}(u) = \frac{\mathcal{FT}_{q_\chi}(u)}{\mathcal{FT}_{f_W}(u)} \quad (33)$$

By Assumption A8(i), $\mathcal{FT}_{f_W}(u)$ has only isolated zeros, and hence (33) does not divide by zero except on a set of measure zero. Taking the inverse Fourier transform then yields $a_\chi(v)$. This function can then be rearranged to solve for $\chi^{-1}(\cdot)$ as

$$\chi^{-1}(v) = F_B^{-1} \left(1 - \frac{a_\chi(v)(1 - F_B(v))}{(1 - F_R(v))f_{P^A}(v)} \right)$$

Note that this expression does not divide by zero at any v that is in the support of P^A and that is less than the upper bound of the support of R ; at such v , $(1 - F_R(v))f_{P^A}(v)$ is positive. Also, $F_B(\cdot)$ is invertible (because $f_B(\cdot)$ is positive by Assumption A5). Thus $\chi^{-1}(\cdot)$ is nonparametrically identified. \square

Note that Proposition 12 only demonstrates identification of $\chi^{-1}(p^A)$ at each p^A that is less than the upper bound of the support of R . This is because at each point beyond this upper bound, the term $1 - F_R(p^A)$ is equal to zero. However, knowing χ^{-1} beyond this point is unnecessary because at such a p^A the game will always end at the auction (because the auction price is larger than the largest R).

For estimation, I let $h_\chi(\cdot) \equiv 1 - F_B(\chi^{-1}(\cdot))$, and I approximate this function as a flexible piecewise linear spline, denoted $h_\chi(\cdot; \theta^\chi) \equiv LS_\chi(\cdot)' \theta^\chi$, where θ^χ is a vector of spline parameters and $LS_\chi(\cdot) : \mathbb{R} \mapsto \mathbb{R}^{K_\chi}$ is the linear spline transformation for a fixed vector of knots $\{v_k^\chi\}_{k=1}^{K_\chi}$, defined analogously to the linear spline transformation described in Appendix C.4 for estimating seller valuation bounds. I choose the knots $\{v_k^\chi\}_{k=1}^{K_\chi}$ to be uniformly spaced between the 0.001–0.999 quantiles of \tilde{P}^A , and set $\chi^{-1}(v) = v$ for any v outside of these knots. I choose $K_\chi = 25$. I estimate

θ^x by solving the following constrained least squares problem:

$$\min_{\theta^x} \sum_{k=1}^{K_\chi} \left[\widehat{\Pr}(D_1^B = 0 | \tilde{P}^A = v_k^x, \tilde{P}^A < \tilde{R}) \left(\int \widehat{M}_\chi(v_k^x, z) dz \right) - \widetilde{LS}_\chi(v_k^x)' \theta^x \right]^2 \quad (34)$$

where $\widetilde{LS}_\chi(v_k^x)$ is a K_χ -by-1 vector with the ℓ^{th} element given by

$$\int \frac{(LS_\chi(v_k^x - w)' e(\ell)) \widehat{M}_\chi(v_k^x, w)}{1 - \widehat{F}_B(v_k^x - w)} dw$$

and $e(\ell)$ is a K_χ -by-1 selection vector with 1 in the ℓ^{th} spot and zeros elsewhere. I impose the constraints that $h_\chi(\cdot; \theta^x)$ is decreasing and $0 \leq h_\chi(v_k^x; \theta^x) \leq 1 - F_B(v_k^x)$ for each k . Computing the objective function in (34) requires the estimates of \widehat{F}_B , \widehat{F}_R , \widehat{f}_{PA} , and \widehat{f}_W from above. The object $\widehat{M}_\chi(\cdot)$ comes from evaluating $M_\chi(\cdot)$ using these estimated distributions. To estimate $\widehat{\Pr}(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R})$, I use a local linear regression of the event $1\{D_1^B = 0\}$ on realizations of $\tilde{P}^A = \tilde{p}^A$ using observations where the auction price falls below the reserve ($\tilde{P}^A < \tilde{R}$). As with the estimator described in (29), the estimator in (34) falls into the class of two-step semiparametric GMM estimators and will be consistent for $\hat{\theta}^x$. From estimates of $h_\chi(\cdot)$, an estimate of $\chi^{-1}(\cdot)$ is given by $\chi^{-1}(\cdot) = F_B^{-1}(1 - h_\chi(\cdot))$.³⁸

Panels A and B of Figure A4 display the estimates for $\hat{\rho}^{-1}(\cdot)$ evaluated at the upper and lower bounds for the seller CDF. Panels C and D display the estimates for $\hat{\chi}^{-1}(\cdot)$. The goodness of fit of the estimates for $\hat{\chi}^{-1}(\cdot)$ can be evaluated using the integrated absolute error, as in (30), comparing the conditional probability statement $\widehat{\Pr}(D_1^B = 0 | \tilde{P}^A = v_k^x, \tilde{P}^A < \tilde{R})$ estimated from the data to its fitted equivalent from (34). I find this error to be less than 0.1% in both the dealers and fleet/lease samples, as shown in Table A5. Panels E and F in Figure A5 display this estimated conditional probability and the corresponding fitted estimates.

C.6 Technical Details on Estimation of the Direct Mechanism Corresponding to Real-World Bargaining

I begin this section with a nonparametric identification proof.

Proposition 13. *The function $g(\cdot)$ in (7) is nonparametrically identified by f_R , f_{PA} , f_W , F_B , and $\Pr(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A)$.*

³⁸The estimation of χ^{-1} , ρ , and ρ^{-1} each require constructing inverse functions (of F_B , F_R , and F_S); I construct these via linear interpolation of their respective CDF estimates. I do the same for inverse CDFs discussed in Appendix C.6.

Proof. Equation (7) can be written as

$$\Pr(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}^A) \left(\int M_g(\tilde{r}, \tilde{p}^A, z) dz \right) = \int \frac{1 - F_B(g(\tilde{r} - w, \tilde{p}^A - w)) M_g(\tilde{r}, \tilde{p}^A, w)}{1 - F_B(\tilde{p}^A - w)} dw \quad (35)$$

where $M_g(\tilde{r}, \tilde{p}^A, w) \equiv f_R(\tilde{r} - w) f_{P^A}(\tilde{p}^A - w) f_W(w)$. Now let

$$a(v_1, v_2) \equiv \frac{1 - F_B(g(v_1, v_2)) f_R(v_1) f_{P^A}(v_2)}{1 - F_B(v_2)} \quad (36)$$

and let $v = (v_1, v_2)$. The integrand on the right-hand side of (35) is a convolution of $a(v)$ against the density $f_W(w)$, given by

$$q_g(v) = \int a(v_1 - w, v_2 - w) f_W(w) dw$$

Identification of $a(v)$ follows from a convolution theorem argument. Taking the Fourier transform of q_g yields the following, where $y = (y_1, y_2)$, $y_1 = v_1 - w$, $y_2 = v_2 - w$, and hence $dy = du$:

$$\begin{aligned} \mathcal{FT}_{q_g}(u) &= \int_{\mathbb{R}^2} \left(\int_{\mathbb{R}} a(v_1 - w, v_2 - w) f_W(w) dw \right) e^{-2\pi i(v_1 u_1 + v_2 u_2)} dv \\ &= \int_{\mathbb{R}} f_W(w) \left(\int_{\mathbb{R}^2} a(v_1 - w, v_2 - w) e^{-2\pi i(v_1 u_1 + v_2 u_2)} dv \right) dw \\ &= \int_{\mathbb{R}} f_W(w) \left(\int_{\mathbb{R}^2} a(y_1, y_2) e^{-2\pi i((y_1 + w)u_1 + (y_2 + w)u_2)} dy \right) dw \\ &= \left(\int_{\mathbb{R}} f_W(w) e^{-2\pi i w(u_1 + u_2)} dw \right) \int_{\mathbb{R}^2} a(y_1, y_2) e^{-2\pi i(y_1 u_1 + y_2 u_2)} dy \\ &= \mathcal{FT}_{f_W}(u_1 + u_2) \mathcal{FT}_a(u) \end{aligned}$$

Because \mathcal{FT}_{f_W} has only isolated zeros (Assumption A8(i)), it follows that

$$\mathcal{FT}_a(u) = \frac{\mathcal{FT}_{q_g}(u)}{\mathcal{FT}_{f_W}(u_1 + u_2)}$$

Taking the inverse Fourier transform then yields $a(v_1, v_2)$. This function can then be rearranged to solve for $g(\cdot)$ as

$$g(v_1, v_2) = F_B^{-1} \left(1 - \frac{a(v_1, v_2)(1 - F_B(v_2))}{f_R(v_1) f_{P^A}(v_2)} \right)$$

where $f_R(v_1) f_{P^A}(v_2)$ is positive at every v_1 in the support of R and v_2 in the support of P^A ; and $F_B(\cdot)$ is invertible (because $f_B(\cdot)$ is positive by Assumption A5). Thus $g(\cdot)$ is nonparametrically identified. \square

For estimation, I approximate the function $h_g(r, p^A) \equiv \frac{1-F_B(g(r, p^A))}{1-F_B(p^A)}$ using a flexible bilinear spline defined on a two-dimensional grid for \tilde{r} and \tilde{p}^A , with $K_{\mathcal{A}} = 25$ grid points in each dimension. Let these grid points be denoted $\{v_k^R\}_{k=1}^{K_{\mathcal{A}}}$ and $\{v_j^{P^A}\}_{j=1}^{K_{\mathcal{A}}}$. This yields $K_g = K_{\mathcal{A}}^2 = 625$ parameters to be estimated. I chose these knots to be uniformly spaced between the 0.001 and 0.999 quantiles of \tilde{R} and \tilde{P}^A , respectively. I denote this approximation $h_g(r, p^A; \theta^g) \equiv BL(r, p^A)' \theta^g$, where $BL(\cdot) : \mathbb{R}^2 \mapsto \mathbb{R}^{K_g}$.³⁹ Like the linear spline approximation used to estimate χ^{-1} and bounds on F_S , this bilinear approximation has several advantages: it is very flexible; it allows me to perform the integration against f_W in a first step, outside of the optimization over θ_g ; and it eases the calculation of analytical derivatives of the objective function to feed to the optimization routine.

I estimate θ^g by solving the following constrained least squares problem:

$$\min_{\theta^g} \sum_{k=1}^{K_{\mathcal{A}}} \sum_{j=1}^{K_{\mathcal{A}}} \left[\widehat{\Pr}(\mathcal{A} | \tilde{R} = v_k^R, \tilde{P}^A = v_j^{P^A}) \left(\int \widehat{M}_g(v_k^R, v_j^{P^A}, z) dz \right) - \widehat{BL}_g(v_k^R, v_j^{P^A})' \theta_g \right]^2 \quad (37)$$

where $\widehat{BL}_g(v_k^R, v_j^{P^A})$ is a K_g -by-1 vector with the ℓ^{th} element given by $\int (BL_g(v_k^R - w, v_j^{P^A} - w)' e(\ell)) \widehat{M}_g(v_k^R, v_j^{P^A}, w) dw$ and $e(\ell)$ is a K_g -by-1 selection vector with 1 in the ℓ^{th} spot and zeros elsewhere. The constraints I impose are that $h_g(r, p^A) \in [0, 1]$; $h_g(r, p^A)$ is decreasing in r ; and $g(r, p^A) \geq \underline{g}(r, p^A) \equiv \max\{p^A, \rho^{-1}(r)\} \Rightarrow h(r, p^A) \leq \frac{1-F_B(g(r, p^A))}{1-F_B(p^A)}$. I enforce this latter constraint by evaluating $\rho^{-1}(r) = F_S^{-1}(F_R(r))$ at $F_S = \hat{F}_S^L$. This ensures that the estimated real-world mechanism is consistent with the estimated valuation distributions in that it does not imply trade occurring for cases where $S > B$.

In (37), the objects \hat{F}_B , \hat{f}_R , \hat{f}_{P^A} , and \hat{f}_W are those estimated above, and $\widehat{M}_g(\cdot)$ comes from evaluating $M_g(\cdot)$ using these estimated distributions. To estimate the conditional probability $\widehat{\Pr}(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}^A)$ I use a tensor product of two univariate cubic b-spline functions, with fifteen knots in each dimension (uniformly spaced between the 0.001 and 0.999 quantiles of \tilde{R} and \tilde{P}^A , respectively).⁴⁰ As with (29) and (34), this estimator falls into the class of two-step semiparametric GMM estimators, which will be consistent for estimating $\hat{\theta}^g$. With estimates of $\hat{\theta}^g$,

³⁹For $r \in [v_k^R, v_{k+1}^R]$, $p^A \in [v_j^{P^A}, v_{j+1}^{P^A}]$, $BL(r, p^A)$ returns a vector of K_g elements, with one element corresponding to each combination of grid points from $\{v_k^R\}_{k=1}^{K_{\mathcal{A}}}$ and $\{v_j^{P^A}\}_{j=1}^{K_{\mathcal{A}}}$ (and hence one element corresponding to each element of θ_g). Specifically, $BL(r, p^A)$ returns zeros everywhere except for the following locations, where $m(r, p^A) = (v_{k+1}^R - v_k^R)(v_{j+1}^{P^A} - v_j^{P^A})$: it returns $(v_{k+1}^R - r)(v_{j+1}^{P^A} - p^A)/m(r, p^A)$ for the element in θ_g corresponding to $(v_k^R, v_j^{P^A})$; $(r - v_k^R)(v_{j+1}^{P^A} - p^A)/m(r, p^A)$ for the element in θ_g corresponding to $(v_{k+1}^R, v_j^{P^A})$; $(v_{k+1}^R - r)(p^A - v_j^{P^A})/m(r, p^A)$ for the element in θ_g corresponding to $(v_k^R, v_{j+1}^{P^A})$; and $(r - v_k^R)(p^A - v_j^{P^A})/m(r, p^A)$ for the element in θ_g corresponding to $(v_{k+1}^R, v_{j+1}^{P^A})$.

⁴⁰This yields 121 parameters to be estimated in the estimation of $\widehat{\Pr}(\mathcal{A} = 1 | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}^A)$. Given a set of I knots, $\{t_i\}_{i=1}^I$, a univariate b-spline transformation of degree d evaluated at a point v is defined as follows. $BS_1(v)$ returns an $(I-d-1)$ -by-1 vector with 1 as its i^{th} element if $t_i \leq v < t_{i+1}$ and 0 otherwise. $BS_k(v)$, for $k \in \{2, \dots, d\}$, is defined recursively, returning a vector with the i^{th} element given by $BS_{i,k}(v) = \frac{v-t_i}{t_{i+k}-t_i} BS_{i,k-1}(v) + \frac{t_{i+k+1}-v}{t_{i+k+1}-t_{i+1}} BS_{i+1,k-1}(v)$. A cubic b-spline is given by $BS_3(v)$, and with $I = 15$ knots, yields $I - d - 1 = 11$ parameters to be estimated. The tensor product of two such univariate b-splines thus yields 121 parameters to be estimated. In estimating this conditional probability, I replace any estimates lying outside of $[0, 1]$ with 0 and 1.

I then obtain $g(r, p^A) = F_B^{-1}(1 - (1 - F_B(p^A))h_g(r, p^A; \hat{\theta}_g))$.

Panels E and F of Figure A4 display the estimates for $\hat{g}(\cdot)$ evaluated at $P^A = 0$. The goodness of fit of the estimates for \hat{g} can be evaluated using the integrated absolute error, just as in (30) for the seller distribution bounds. This compares the conditional probability statement $\widehat{\Pr}(\mathcal{A} | \tilde{R} = v_k^R, \tilde{P}^A = v_j^{P^A})$ estimated from the data to its fitted equivalent from (37). I find this error to be 6.0% in the dealers sample and 2.47% in the fleet/lease sample, as shown in Table A5. This error is somewhat larger than that of the estimates of F_S^L , F_S^U , and χ^{-1} , as the estimation of g requires fitting a two-dimensional flexible function rather than just a one-dimensional function.

C.7 Technical Details on Computation of Ex-ante Efficient Mechanisms and Welfare Measures

C.7.1 Ex-ante Efficient Mechanisms

The focus of this paper is the efficiency of the bilateral bargaining between the seller and the high bidder. To analyze this, I integrate welfare measures over the support of buyer and seller types who bargain in a given instance of the game, $([b(p^A), \bar{b}], [\underline{s}(p^A), \bar{s}])$, and then integrate over realizations of the auction price p^A . Because valuations and transfers are additively separable in game-level heterogeneity, W will play no role in evaluating welfare measures.

Allocation functions x^η corresponding to mechanisms along the ex-ante efficient frontier can be computed as follows. For a given realization of the auction price, p^A , let $f_S(s|p^A) = \frac{f_S(s)}{1 - F_S(\underline{s}(p^A))}$ and $f_B(b|p^A) = \frac{f_B(b)}{1 - F_B(\bar{b}(p^A))}$ represent the densities of seller and buyer valuations conditional on p^A at the beginning of the bargaining (i.e. these are the Bayesian updated beliefs about valuations after the auction stage). Let $F_S(s|p^A)$ and $F_B(b|p^A)$ represent their corresponding distributions. Let q represent the expected utility of the lowest seller type ($\underline{s}(p^A)$). The allocation function x^η is the solution to the linear programming problem,

$$\begin{aligned} \max_x \left\{ \eta[q(x; p^A) + \bar{U}_S(x; p^A)] + (1 - \eta)[\Gamma(x; p^A) - q(x; p^A) + \bar{U}_B(x; p^A)] \right\} \quad (38) \\ \text{subject to } x(s, b; p^A) \in [0, 1] \quad \forall (s, b) \in [\underline{s}(p^A), \bar{s}] \times [\underline{b}(p^A), \bar{b}] \\ \bar{x}_B(b; p^A) \equiv \int_{\underline{s}(p^A)}^{\bar{s}} x(s, b; p^A) f_S(s|p^A) ds \quad \text{weakly increasing in } b \\ \bar{x}_S(s; p^A) \equiv \int_{\underline{b}(p^A)}^{\bar{b}} x(s, b; p^A) f_B(b|p^A) db \quad \text{weakly decreasing in } s \\ \Gamma(x; p^A) \geq 0 \end{aligned}$$

See Section 2 of Williams (1987). The values $\bar{U}_S(x)$ and $\bar{U}_B(x)$ are the ex-ante expected utilities

(prior to knowing their valuations) of the seller and buyer, given by

$$\bar{U}_S(x; p^A) = \int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} x(s, b; p^A) f_B(b|p^A) F_S(s|p^A) ds db \quad (39)$$

$$\bar{U}_B(x; p^A) = \int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} x(s, b; p^A) (1 - F_B(b|p^A)) f_S(s|p^A) ds db \quad (40)$$

The objects $\Gamma(x; p^A)$ and $q(x; p^A)$ are given by

$$\Gamma(x; p^A) = \int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} (b - s) x(s, b; p^A) f_S(s|p^A) f_B(b|p^A) ds db - [\bar{U}_S(x; p^A) + \bar{U}_B(x; p^A)]$$

$$q(x; p^A) = \begin{cases} 0 & \text{if } \eta \leq \frac{1}{2} \\ \Gamma(x; p^A) & \text{if } \eta > \frac{1}{2} \end{cases}$$

The monotonicity constraints on \bar{x}_B and \bar{x}_S in (38) ensure that the mechanism will be incentive compatible. Williams (1987) referred to these constraints as the “troublesome constraints,” and he demonstrated that these constraints can be ignored if the distributions F_S and F_B are *regular*. To define regularity, let

$$\phi_S(s, \alpha_1; p^A) \equiv s + \alpha_1 \frac{F_S(s|p^A)}{f_S(s|p^A)} \quad \text{and} \quad \phi_B(b, \alpha_2; p^A) \equiv b - \alpha_2 \frac{1 - F_B(b|p^A)}{f_B(b|p^A)} \quad (41)$$

Regularity corresponds to $\phi_S(\cdot, 1; p^A)$ and $\phi_B(\cdot, 1; p^A)$ being increasing functions. When this is the case, the mechanisms along the ex-ante efficient frontier can be written as follows:

$$x^{\alpha_1(\eta), \alpha_2(\eta)}(s, b; p^A) = 1 \{ \phi_B(b, \alpha_2(\eta); p^A) \geq \phi_S(s, \alpha_1(\eta); p^A) \} \quad (42)$$

where the parameters $(\alpha_1(\eta), \alpha_2(\eta))$ can be solved for at any given η as described in Theorem 3 of Williams (1987). This simplifies the linear programming problem in (38) considerably to merely solving for a scalar.

When distributions are not regular, however, as in my application, I am required to solve the full linear programming problem (38) subject to all of the original constraints. Plugging in estimates of densities, CDFs, and lower support functions from steps 2–5 of the estimation procedure, I use linear programming to solve the problem in (38) separately at each value of p^A (i.e. each realization of the lower bound of the support of bargaining types) on a grid of points for the s and b dimensions. Appendix C.7.4 discusses this grid choice and also discusses numerical integration. I also solve this problem separately at different F_S lying in the estimated CDF bounds, as described in the Section C.7.2.

C.7.2 Computing Bounds on Welfare

Each of the welfare measures described in Section 5 can be computed by plugging in the corresponding empirical estimates from Section 4 and using *any* seller valuation CDF lying between the estimated seller CDF bounds. Let the space of such CDFs be given by $\mathcal{F}^* = \{F_S \in \mathcal{F} : F_S(v) \in [F_S^L(v), F_S^U(v)] \forall v \in [\underline{s}, \bar{s}]\}$. As in the proof of Proposition 6, let the notation $\mathcal{W}(x, F_S)$ refer to a welfare measure under a given allocation function x and at a given seller valuation distribution F_S (suppressing dependence on F_B given that I have point estimates of F_B). In the spirit of Reguant (2016), bounds on any welfare measure $\mathcal{W}(x, F_S)$, denoted $[\underline{\mathcal{W}}(x), \overline{\mathcal{W}}(x)]$, can be computed by

$$\left[\min_{F_S \in \mathcal{F}^*} \mathcal{W}(x, F_S), \max_{F_S \in \mathcal{F}^*} \mathcal{W}(x, F_S) \right] \quad (43)$$

Numerically computing these bounds is extremely burdensome, as it requires searching over a high-dimensional parameter vector F_S and solving a linear programming problem at each realization of the lower bound of the support of types at each guess of F_S . To reduce this computational complexity, in computing bounds on welfare I reduce the number of knots in the spline approximation to $K'_S = 25$, with these knots uniformly spaced over the range of the original $K_S = 200$ knots. Panels C and D of Appendix Figure A7 demonstrate that this number of knots still approximates the bounds well. Let the subset of these uniformly spaced indices be denoted \mathcal{K} and let the vector $\tilde{\theta}^{S,L} = \{\hat{\theta}_{k'}^{S,L} : k' \in \mathcal{K}\}$ and the vector $\tilde{\theta}^{S,U} = \{\hat{\theta}_{k'}^{S,U} : k' \in \mathcal{K}\}$; that is, these vectors are subsets of the original estimated spline coefficients from Section 4 corresponding to the uniformly spaced K'_S knots. The search to find bounds on welfare can then be performed over the set $\Theta^S \equiv \{\theta^S : \theta_k^S \in [\tilde{\theta}_k^{S,L}, \tilde{\theta}_k^{S,U}] \text{ and } \theta_k^S \leq \theta_{k+1}^S \forall k < K'_S\}$ following (43).

To ease the computational burden further, I derive the monotonicity results presented in Proposition 6 and summarized in Table 3. Each cell in Table 3 marked with a \downarrow signifies that the specified measure will decrease given a first-order stochastically dominating shift in the distribution of seller valuations. Each cell marked with an asterisk indicates that there is no mathematical proof of a monotonicity result and that the bounds must be determined numerically using (43). The buyer and seller gains in the column labeled *second-best* are marked with “–” because, in the analysis below, I will report bounds on these quantities that correspond to the respective player’s surplus under the F_S that leads to the maximum and minimum bounds on the total gains from trade. Similarly, for the buyer-optimal column, bounds on the total gains from trade and seller gains correspond to the F_S that leads to the maximum and minimum bounds on the buyer gains from trade, and similarly for the seller-optimal column (although it turns out that some of these welfare measures also satisfy the monotonicity property, as discussed in the proof of Proposition 6).

C.7.3 Bounding Bargaining Costs

For the real-world mechanism, bargaining costs will enter into the computation of the expected gains from trade. For the lower bounds on expected gains from trade in the real-world bargaining ($\underline{\mathcal{W}}_{EG}(x^{RW})$, $\underline{\mathcal{W}}_B(x^{RW})$, and $\underline{\mathcal{W}}_S(x^{RW})$ in the notation of the proof of Proposition 6), I incorporate these costs by subtracting an upper bound on expected bargaining costs, described below. To obtain an upper bound on gains from trade, bargaining costs are set to zero.

Bounds on the parameters c_B and c_S can be derived from cases in which a player chooses to make a counteroffer. A necessary condition for a party to choose to counter is that the payoff in the state where the opponent accepts with probability one must exceed the player's payoff from accepting the current offer on the table. That is, $p_2^S - c_S \geq p_1^B$ for a seller offer and $b - p_3^B - c_B \geq b - p_2^S$ for a buyer offer. Rearranging yields

$$p_2^S - p_1^B \geq c_S \quad (44)$$

$$p_2^S - p_3^B \geq c_B \quad (45)$$

Thus, an upper bound on c_S is given by the minimum gap between period 2 and period 1 offers and an upper bound on c_B is given by the minimum gap between period 2 and period 3 offers (in cases where such offers took place). Rather than use the minimum over all observations, I follow Chernozhukov, Lee, and Rosen (2013) to obtain a bias-corrected, one-sided 95% confidence bound for c_S and c_B and treat these as upper bounds on c_S and c_B , which I denote \bar{c}_S and \bar{c}_B .⁴¹

Let the random variable \mathcal{T} be the period in which the game ends. The buyer's and seller's ex-ante expected disutility due to bargaining costs are then given $c_B E[\lfloor \mathcal{T}/2 \rfloor]$ and $c_S E[\lfloor (\mathcal{T} - 1)/2 \rfloor]$, respectively, because by round t of the game the buyer has made a total of $\lfloor \frac{t}{2} \rfloor$ offers (where $\lfloor \cdot \rfloor$ is the floor function), and similarly for the seller.

Applying this approach yields estimates of an upper bound of less than \$25 for both c_S and c_B in the dealers sample and \$50 for both c_S and c_B in the fleet/lease sample. The upper bound on the total expected loss due to bargaining costs is \$23.3 for buyers and \$4.1 for sellers in the dealers sample, and \$33.6 for buyers and \$5.2 for sellers in the fleet/lease sample.

C.7.4 Numerical Integration

The integrals in the welfare measures I evaluate involve integration in three dimensions, p^A , b , and s . For integration in the p^A dimension, this integral is computed using Gauss-Chebyshev quadrature with 25 nodes. For any univariate function $g(\cdot)$ to be integrated over $[\underline{v}, \bar{v}]$, Gauss-

⁴¹Chernozhukov, Lee, and Rosen (2013) explained that taking the minimum over all observations using inequalities such as those in those in (44) and (45) will be biased downward, precisely because they are derived from taking a minimum. The authors suggested a bias-corrected $1 - \alpha$ confidence bound (which the authors refer to as being half-median unbiased) given, in this case, by taking the $1 - \alpha$ quantile of bootstrapped estimates of the minimum. I perform this using 200 bootstrap replications.

Chebyshev quadrature is given by $\int_{\underline{v}}^{\bar{v}} g(v)dv \approx \frac{\pi(\bar{v}-\underline{v})}{2K^{GC}} \sum_{k=1}^{K^{GC}} g(x_k)w_k$, where K^{GC} is the number of nodes; $x_k = (1/2)(z_k + 1)(\bar{v} - \underline{v}) + \underline{v}$; $w_k = (1 - z_k^2)^{1/2}$; and $z_k = \cos(\pi(2k - 1)/(2K^{GC}))$ (see Judd 1998).

The integration in the s and b dimension is required not only for evaluating welfare measures but also for solving the linear programming problem in (38); the p^A dimension, on the other hand, is only involved in computing a simple average. For the s and b dimensions, therefore, I choose a larger number of nodes (50 in each dimension) to achieve a high degree of accuracy in solving for the ex-ante efficient mechanisms. I choose these nodes to be evenly spaced quantiles of F_S and F_B , which works particularly well here for numerical integration due to the fact that the seller valuation distribution bounds can be multi-modal (corresponding to a CDF that is nearly flat over large portions of the seller support). I construct f_S and f_B for these numerical integrals by finite differences of F_S and F_B .

C.8 Computing Confidence Sets for Welfare Bounds

Confidence sets for the bounds on welfare can be computed by bootstrapping. Specifically, for any estimated bounds on a welfare measure, which constitute an interval, $[\underline{\mathcal{W}}(x), \overline{\mathcal{W}}(x)]$, the lower 95% bootstrapped confidence band about $\underline{\mathcal{W}}(x)$ and the upper 95% bootstrap confidence band about $\overline{\mathcal{W}}(x)$ will provide a conservative 95% confidence interval for the set $[\underline{\mathcal{W}}(x), \overline{\mathcal{W}}(x)]$. This claim follows by a simple Bonferroni inequality argument: For a fixed $\alpha \in [0, 1]$, let $\underline{c}_{\alpha/2}$ be the lower critical value for the $1 - \alpha$ confidence band for $\underline{\mathcal{W}}(x)$ and $\bar{c}_{1-\alpha/2}$ be the upper critical value for the $1 - \alpha$ confidence band for $\overline{\mathcal{W}}(x)$. Also, let \underline{A} be the event that $\underline{\mathcal{W}}(x) \geq \underline{c}_{\alpha/2}$ and let \bar{A} be the event that $\overline{\mathcal{W}}(x) \leq \bar{c}_{1-\alpha/2}$. Therefore, $\Pr(\underline{A}) = \Pr(\bar{A}) = 1 - \alpha/2$. Bonferroni inequalities imply

$$\Pr(\underline{A} \cap \bar{A}) \geq \Pr(\underline{A}) + \Pr(\bar{A}) - \Pr(\underline{A} \cup \bar{A}) \geq \Pr(\underline{A}) + \Pr(\bar{A}) - 1 = 1 - \alpha,$$

thus completing the argument. To compute these confidence sets I use 200 bootstrap replications of the full estimation procedure and counterfactual computations: for each replication, I take a sample of observations the same size as the original sample, drawn with replacement from the original sample, and use this bootstrap sample to perform all estimation steps 1–6 and all counterfactual computations, including the large numerical search for welfare bounds where required by Table 3/Proposition 6.

This approach yields very similar results to the method proposed by Chernozhukov, Hong, and Tamer (2007) (CHT) but has the advantage of yielding an asymmetric confidence set that is guaranteed to lie within the range of bootstrapped estimates, i.e. the confidence set will naturally be contained within the minimum and maximum bootstrapped estimates, as it is composed of quantiles that are on the interior of these estimates. This is not the case with the symmetric confidence set of CHT, for example, where it is possible for one end of the estimated confidence

set to lie outside the extremes of the bootstrapped estimates. For example, even if the estimated gap between the expected gains from trade in the second-best and real-world mechanisms is positive in every bootstrap sample, the CHT confidence set can contain zero. Given this feature, and given that the above confidence sets are easier to compute, I adopt the above approach throughout the paper.

D Auction House Fees

At wholesale auto auctions, when a buyer and seller agree on a price, both parties pay a fee to the auction house to consummate the trade. These fees introduce a trade-reducing wedge: if a buyer values the car more than the seller, but that difference in valuations is less than the sum of the buyer and seller fees, the parties will not trade. The analysis in the body of the paper does not remove this wedge, but rather lets these same fees be paid in counterfactual mechanisms just as they are in the real-world mechanism. In this section, I compute the ex-post efficient outcome absent auction house fees. I also compute the mechanism that implements the optimal broker fees, also discussed in Myerson and Satterthwaite (1983).

From results in Myerson and Satterthwaite (1983), when an intermediating broker is paid by the buyer and seller whenever trade occurs, and when buyer and seller distributions are regular, then the mechanism that would maximize expected revenue for this broker is given by $x^{\alpha_1(\eta), \alpha_2(\eta)}(s, b; p^A)$ from (42) with $\alpha_1 = \alpha_2 = 1$. Thus, in the regular case, no optimization problem need be solved to determine the broker-optimal mechanism. When distributions are not regular—as is the case with my estimates—it is much more complex to solve for the broker-optimal mechanism. The mechanism is given by the solution to

$$\begin{aligned}
 & \max_x \Gamma(x; p^A) && (46) \\
 & \text{subject to } x(s, b; p^A) \in [0, 1] \quad \forall (s, b) \in [\underline{s}(p^A), \bar{s}] \times [\underline{b}(p^A), \bar{b}] \\
 & \bar{x}_B(b; p^A) \equiv \int_{\underline{s}(p^A)}^{\bar{s}} x(s, b; p^A) f_S(s|p^A) ds \quad \text{weakly increasing in } b \\
 & \bar{x}_S(s; p^A) \equiv \int_{\underline{b}(p^A)}^{\bar{b}} x(s, b; p^A) f_B(b|p^A) db \quad \text{weakly decreasing in } s
 \end{aligned}$$

As with the ex-ante efficient mechanisms, I solve for the broker-optimal mechanism by solving the linear program in (46) separately for each value of p^A on a grid of points for the s and b dimensions, and I do so separately for different F_S lying in the estimated CDF bounds, as described in Section C.7.2.

I now describe how I compute the ex-post efficient outcome absent auction house fees. In practice, buyer and seller fees can consist of both a fixed fee and a percentage commission, but

the latter makes up only a small portion of the overall fee. Here I consider these fees to be solely a fixed component, which I approximate using the average of seller fees and the average of buyer fees observed in the data, denoted h^S and h^B here. This structure implies that the distribution of buyers' willingness to pay and sellers' willingness to sell *absent* auction house fees is simply a mean-shift of the buyer and seller valuation distributions estimated in Section 4; these shifted distributions can be denoted F_{B^*} and F_{S^*} , where $B^* \equiv B + h^B$ and $S^* \equiv S - h^S$.⁴² After performing this mean-shift to the estimated distributions of buyer and seller valuations, I compute the ex-post efficient outcome exactly the same as in the main counterfactuals (which use the non-adjusted valuation distributions).

Table A8 displays the results of my analysis of auction house fees, with panel A showing results for the dealers sample and panel B for the fleet/lease sample. The first column displays the expected gains from trade and probability of trade under ex-post efficiency in the absence of auction house fees. The second column reports the outcome of the broker-optimal mechanism, along with the revenue this would raise for the auction house. The final column is similar to the real-world bargaining welfare results from Tables 4–5, but here I also report the auction house revenue itself, and I include this revenue in the calculation of gains from trade.⁴³

The expected auction house fees in the real bargaining is \$199 in both samples (the last row, last column in panels A and B). The expected gains from trade under ex-post efficiency with these fees removed (the first column in panels A and B) are naturally higher than those reported in Tables 4–5. Part of this is mechanical due to removing the \$199 fee. But Table A8 shows that removing auction house fees increases the ex-post efficient gains from trade by more than just this amount; this is because, in an ex-post efficient world, removing fees would allow additional buyers and sellers (those whose valuations differ by less than the fee amount) to profitably trade.

I find that the optimal revenue for the auction house is \$1,270–3,652 in the dealers sample and \$1,134–2,464 in the fleet/lease sample. These numbers are far above the revenue numbers in the real mechanism. Both buyer gains and seller gains would be much lower in the broker-optimal mechanism than in the real mechanism, as the majority of surplus would be taken by the auction house. The probability of trade would also fall due under the broker-optimal mechanism to the broker's rent-extraction behavior, decreasing from 0.646 in the real-world to 0.271–0.560 in the dealers sample, and from 0.658 to 0.303–0.501 in the fleet/lease sample. This optimal revenue could only be achieved if the auction house were a *monopolist* provider of wholesale dealer-to-dealer

⁴²The following is an example of how to interpret these adjusted distributions relative to the original F_B and F_S distributions referred to elsewhere in the paper: if a buyer has a valuation of $B = b$, this buyer is willing to pay a price b to the seller plus the auction house fee. Absent auction house fees, the buyer would be willing to pay more— $b^* = b + h^B$ —directly to the seller. Similarly, a seller with valuation $S = s$ is willing to accept a price of s from the buyer and also pay h^S to the auction house. Absent auction house fees, the seller would be willing to accept a lower price— $s^* = s - h^S$ —from the buyer.

⁴³Auction house revenue is the sum of the buyer and seller fee when trade occurs and 0 when trade fails. The buyer and seller fee conditional on trade is reported in Appendix Table A3.

trading. In practice, auction houses owned by different companies compete in most major cities, and buyers and sellers have other outlets for buying and selling cars. This competition prevents auction houses from fully exploiting the broker-optimal mechanism.

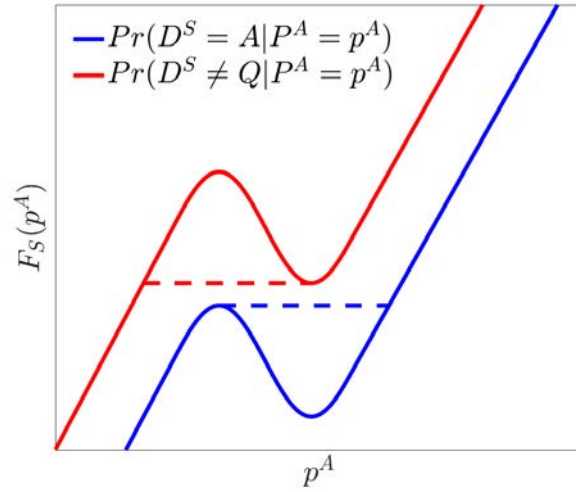
Appendix References

- Abreu, D. and F. Gul (2000): “Bargaining and reputation,” *Econometrica*, 68(1), 85–117.
- Ackerberg, D., J. Chen and Z. Liao (2014): “Asymptotic efficiency of semiparametric two-step GMM,” *Review of Economic Studies*, 81(3), 919–943.
- Admati, A. R. and Perry, M. (1987). Strategic delay in bargaining. *Review of Economic Studies*, 54(3), 345.
- Aradillas-Lopez, A., A. Ghandi, and D. Quint (2013): “Identification and inference in ascending auctions with correlated private values,” *Econometrica*, 81(2), 489–534.
- Ausubel, L. M. and R. J. Deneckere (1989): “Reputation in bargaining and durable goods monopoly,” *Econometrica*, 511–531.
- Bikhchandani, S. (1992). A bargaining model with incomplete information. *Review of Economic Studies*, 59(1), 187.
- Binmore, K., M. J. Osborne, and A. Rubinstein (1992): “Noncooperative models of bargaining,” *Handbook of Game Theory*, 1, 179–225.
- Chatterjee, K. and L. Samuelson (1987): “Bargaining under two-sided incomplete information: An infinite horizon model with alternating offers,” *Review of Economic Studies*, 54(2), 175–192.
- Chatterjee, K. and L. Samuelson (1988): “Bargaining under two-sided incomplete information: The unrestricted offers case,” *Operations Research*, 36(4), 605–618.
- Chernozhukov, V., H. Hong, and E. Tamer (2007): “Estimation and confidence regions for parameter sets in econometric models,” *Econometrica*, 75(5), 1243–1284.
- Chernozhukov, V., S. Lee, and A. M. Rosen (2013): “Intersection bounds: Estimation and inference,” *Econometrica*, 81(2), 667–737.
- Cho, I. (1990). Uncertainty and delay in bargaining. *Review of Economic Studies*, 57(4), 575.
- Coey, D., Larsen, B., and K. Sweeney (2014): “The bidder exclusion effect,” NBER Working Paper 20523.

- Coe, D., Larsen, B., Sweeney, K., and C. Waisman (2018): “The simple empirics of optimal online auctions,” NBER Working Paper 24698.
- Compte, O. and P. Jehiel (2002): “On the role of outside options in bargaining with obstinate parties,” *Econometrica*, 70(4), 1477–1517.
- Cramton, P. C. (1984). Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *Review of Economic Studies*, 51(4), 579.
- Cramton, P. C. (1992). Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59, 205–225.
- Feinberg, Y. and A. Skrzypacz (2005): “Uncertainty about uncertainty and delay in bargaining,” *Econometrica*, 73(1), 69–91.
- Frölich, M. (2006): “Non-parametric regression for binary dependent variables,” *Econometrics Journal*, 9(3), 511–540.
- Fudenberg, D. and J. Tirole (1983): “Sequential bargaining with incomplete information,” *Review of Economic Studies*, 50(2), 221.
- Gantner, A. (2008): “Bargaining, search, and outside options,” *Games and Economic Behavior*, 62(2), 417–435.
- Grossman, S.J. and Perry, M. (1986). Sequential bargaining under asymmetric information. *Journal of Economic Theory*, 39(1), 120–154.
- Gul, F., H. Sonnenschein, and R. Wilson (1986): “Foundations of dynamic monopoly and the Coase conjecture,” *Journal of Economic Theory*, 39(1), 155–190.
- Kennan, J. and R. Wilson (1993): “Bargaining with private information,” *Journal of Economic Literature*, 31(1), 45–104.
- Kim, Dong-Hyuk (2013). Optimal choice of a reserve price under uncertainty,” *International Journal of Industrial Organization*, 31(5), 587–602.
- Levin, D. and J. L. Smith (1994): “Equilibrium in auctions with entry,” *American Economic Review*, 84(3), 585–599.
- Menzel, K. and P. Morganti (2013): “Large sample properties for estimators based on order statistics approaches in auctions,” *Quantitative Economics*, 4(2), 329–375.
- Reguant, M. (2016): “Bounding outcomes in counterfactual analysis,” Working paper, Northwestern University.

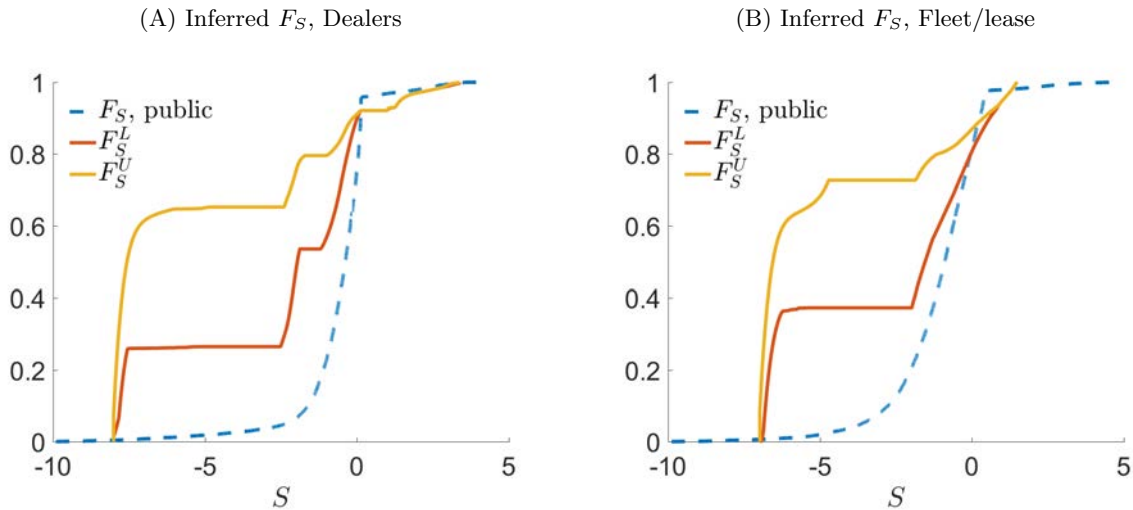
- Roth, A. E. (1995): “Bargaining experiments,” in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth, Princeton University Press, Princeton, NJ.
- Rubinstein, A. (1985a). A bargaining model with incomplete information about time preferences. *Econometrica*, 1151–1172.
- Rubinstein, A. (1985b): “Choice of conjectures in a bargaining game with incomplete information,” in *Game Theoretic Models of Bargaining*, ed. by A. Roth, Cambridge Univ Pr, 99–114.
- Ruppert, D., S. J. Sheather, and M. P. Wand (1995): “An effective bandwidth selector for local least squares regression,” *Journal of the American Statistical Association*, 90(432), 1257–1270.
- Samuelson, W. F. (1985): “Competitive bidding with entry costs,” *Economics Letters*, 17(1-2), 53–57.
- Sobel, J. and I. Takahashi (1983): “A multistage model of bargaining,” *Review of Economic Studies*, 50(3), 411.
- Watson, J. (1998): “Alternating-offer bargaining with two-sided incomplete information,” *Review of Economic Studies*, 65(3), 573–594.
- Zhang, A. J. (2017): “Surplus monotonicity in second-best bargaining,” Working paper, Stanford University.

Figure A1: Weakly Increasing Property of Seller Distribution Bounds (Illustration Only)



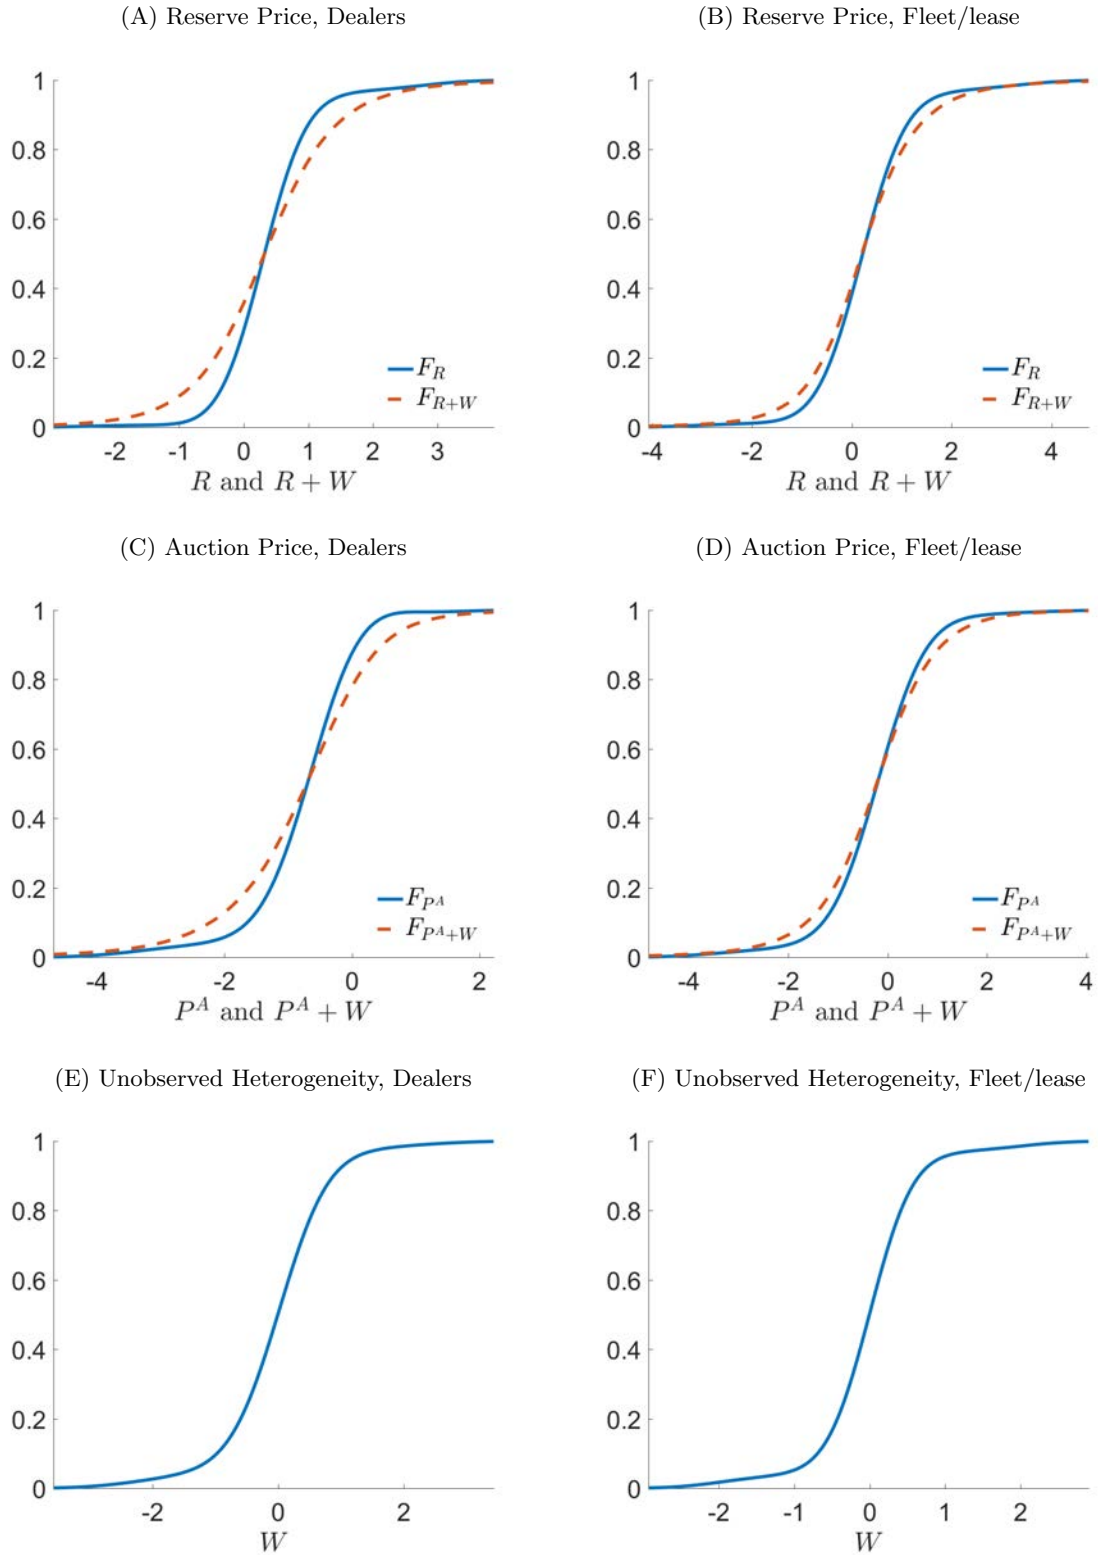
Notes: This figure illustrates an example of conditional probabilities from Proposition 2 (displayed as solid lines) that are not monotone. The dashed lines displays the monotonicized version of the seller CDF bounds from equation (16), but this alternative version does not provide any tighter bounds on the seller distribution because CDFs are monotone by definition. This figure is an illustration only; it does not display estimates.

Figure A2: Comparison of Public vs. Secret Reserve Prices



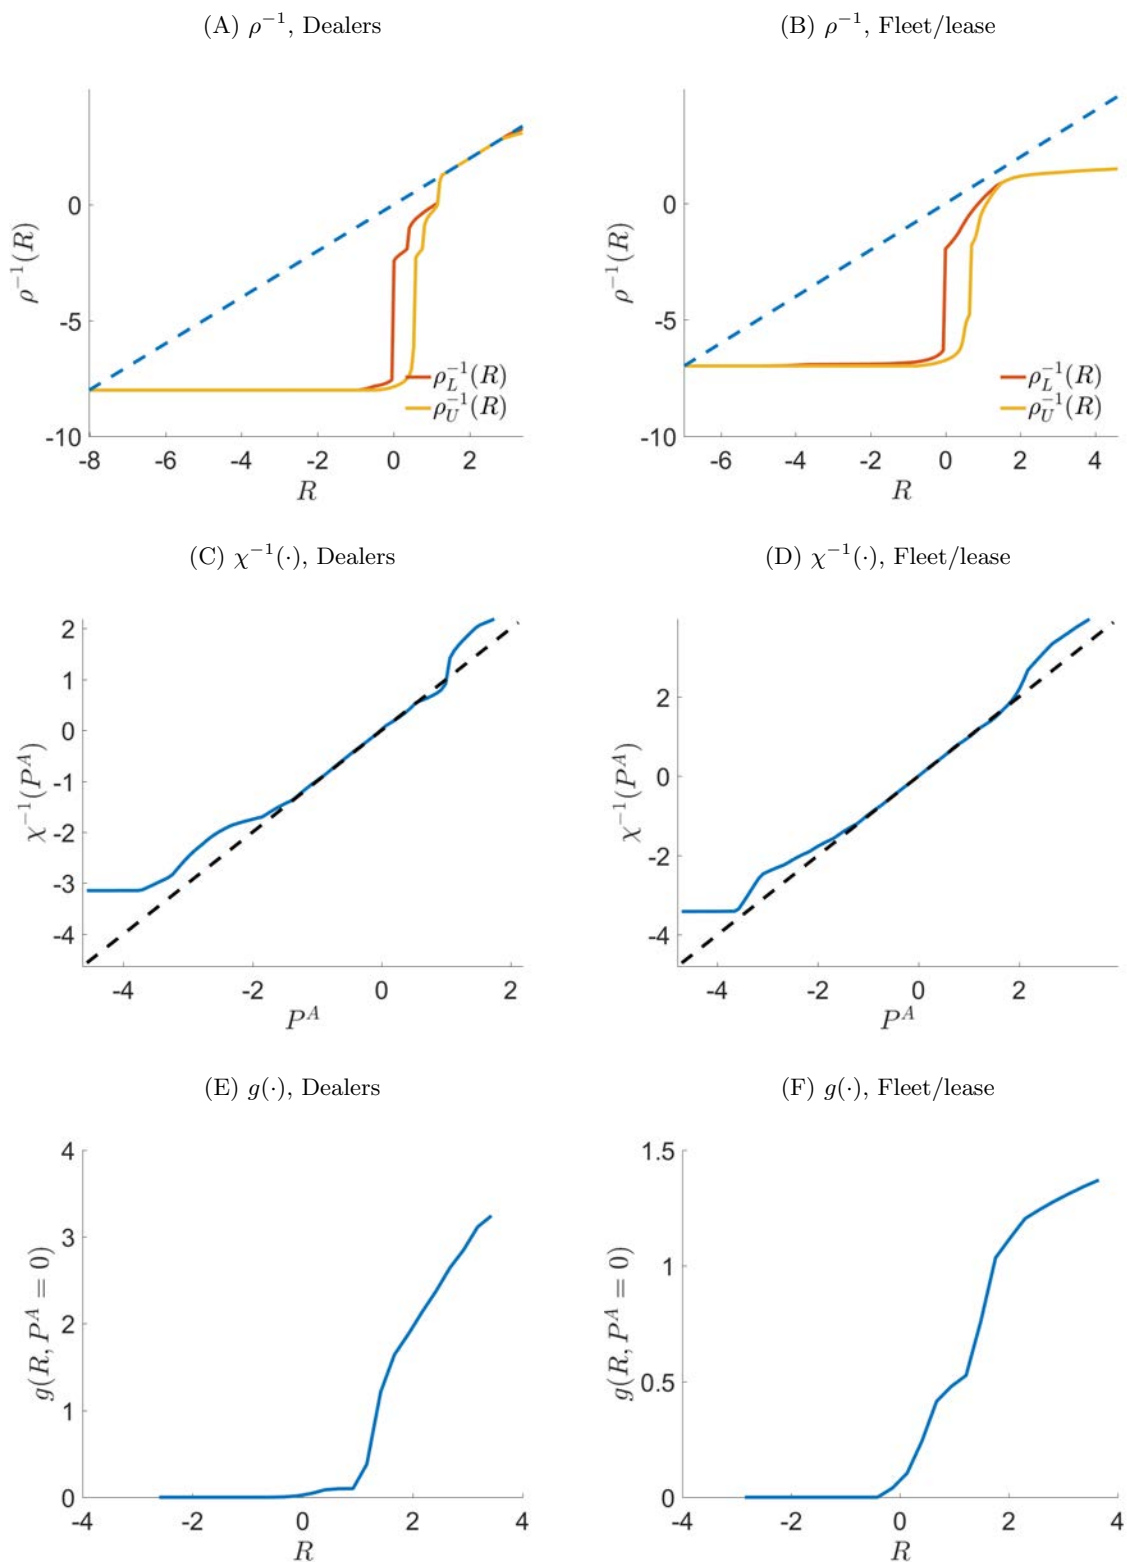
Notes: Panels A and B display the bounds on F_S (solid lines) along with CDF of seller valuations inferred from naively treating reserve prices as optimal public reserves (dashed line). Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Figure A3: Distributions of Reserve Prices, Auction Prices, and Unobserved Heterogeneity



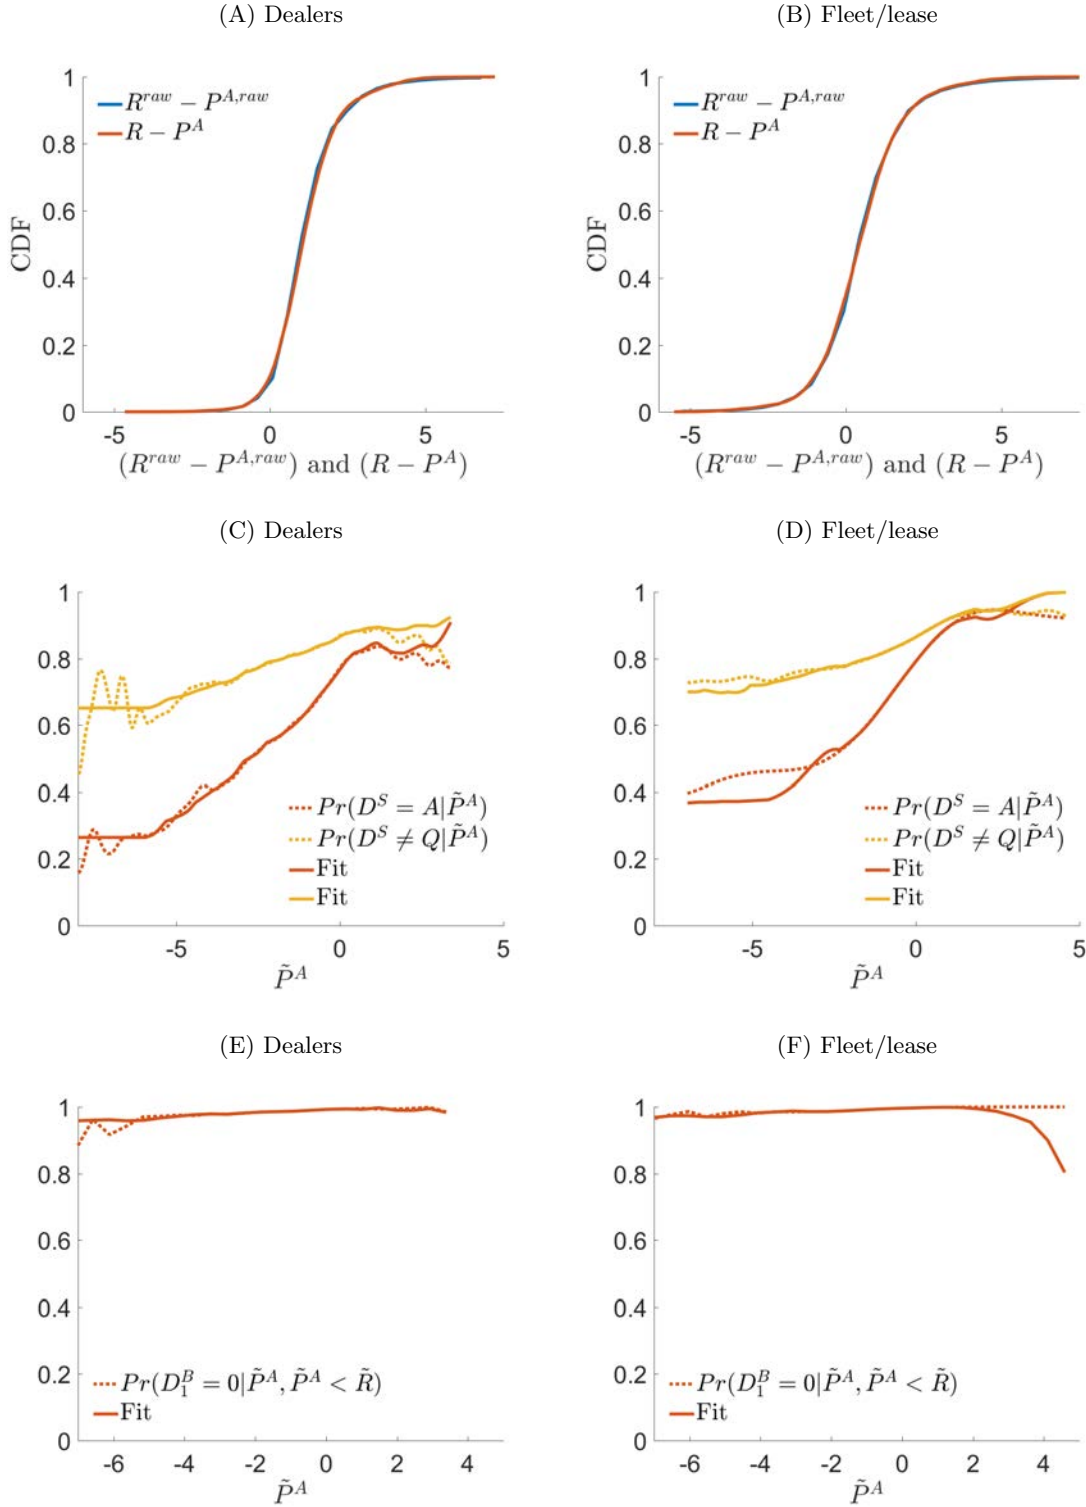
Notes: Panels A–D display distributions of reserve prices and auction prices prior to removing unobserved heterogeneity (dashed lines) and after removing unobserved heterogeneity (solid lines). Panels E and F display the estimated distribution of unobserved heterogeneity. Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Figure A4: Estimates of $\rho^{-1}(\cdot)$, $\chi^{-1}(\cdot)$, and $g(\cdot)$



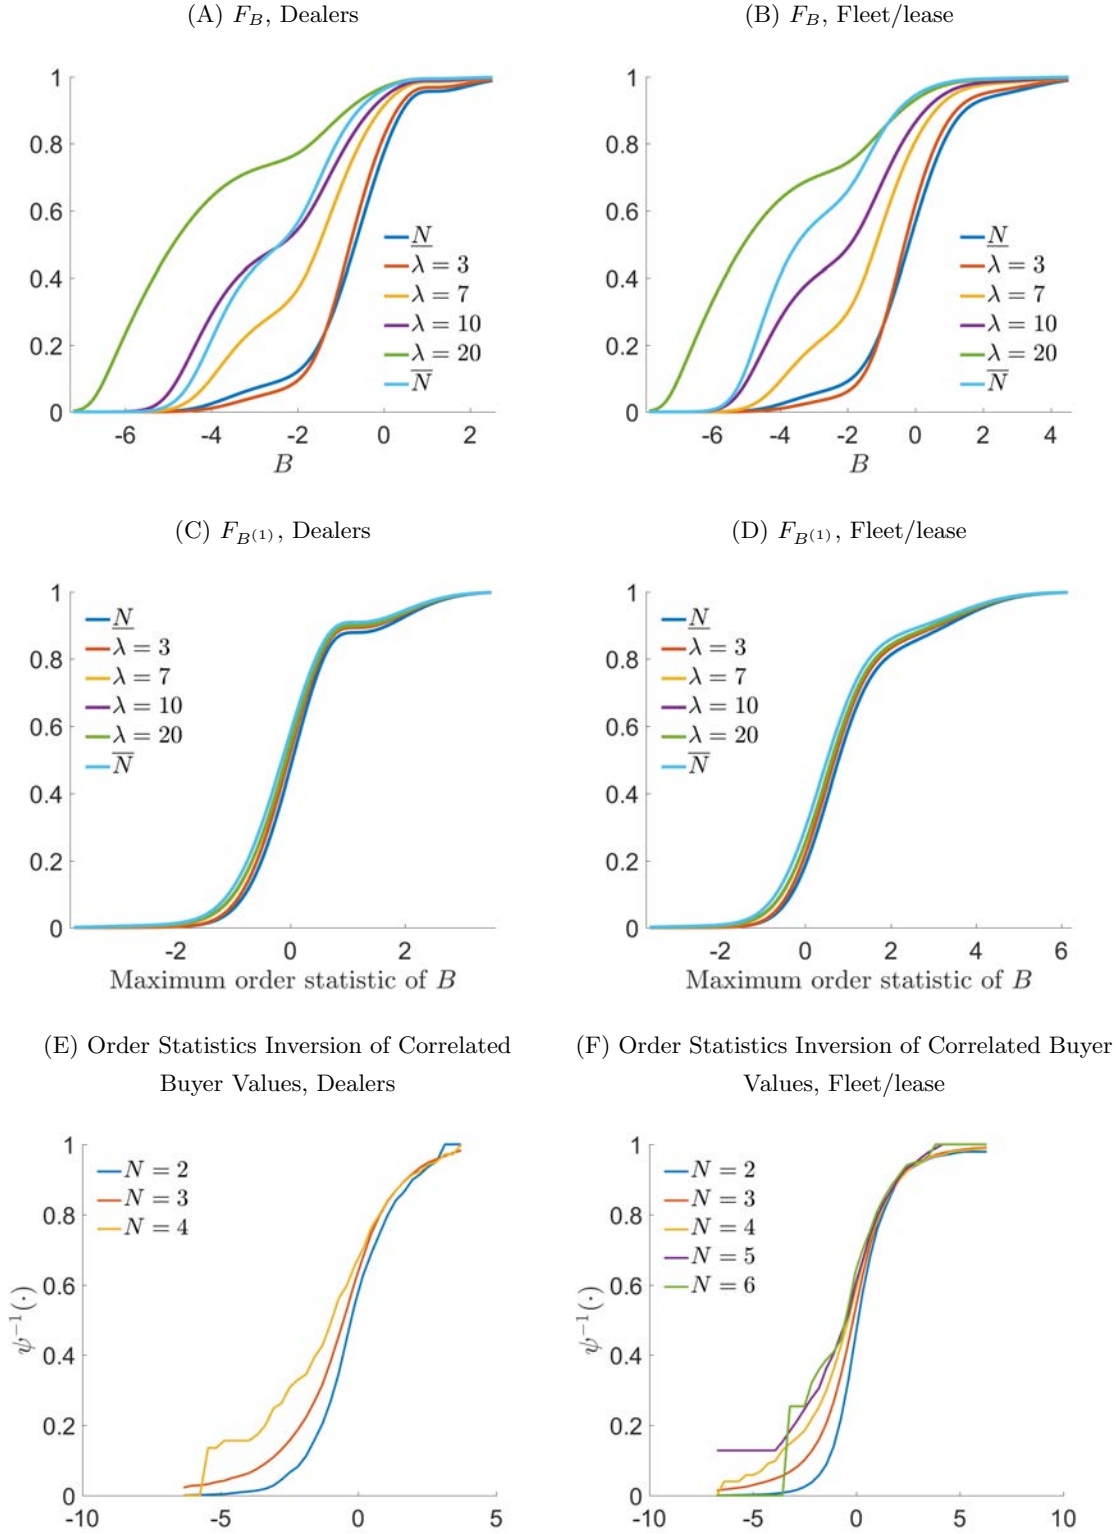
Notes: Panels A and B display estimates of $\rho^{-1}(R)$ using the upper and lower bound on the distribution of seller values (solid lines) as well as the 45 degree line (dashed line). Panels C and D display the estimates of $\chi^{-1}(P^A)$ (solid line) and the 45 degree line (dashed line). Panels E and F display estimates of $g(\cdot, 0)$; that is, the $g(R, P^A)$ function evaluated at $P^A = 0$. Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Figure A5: Fit of Estimates



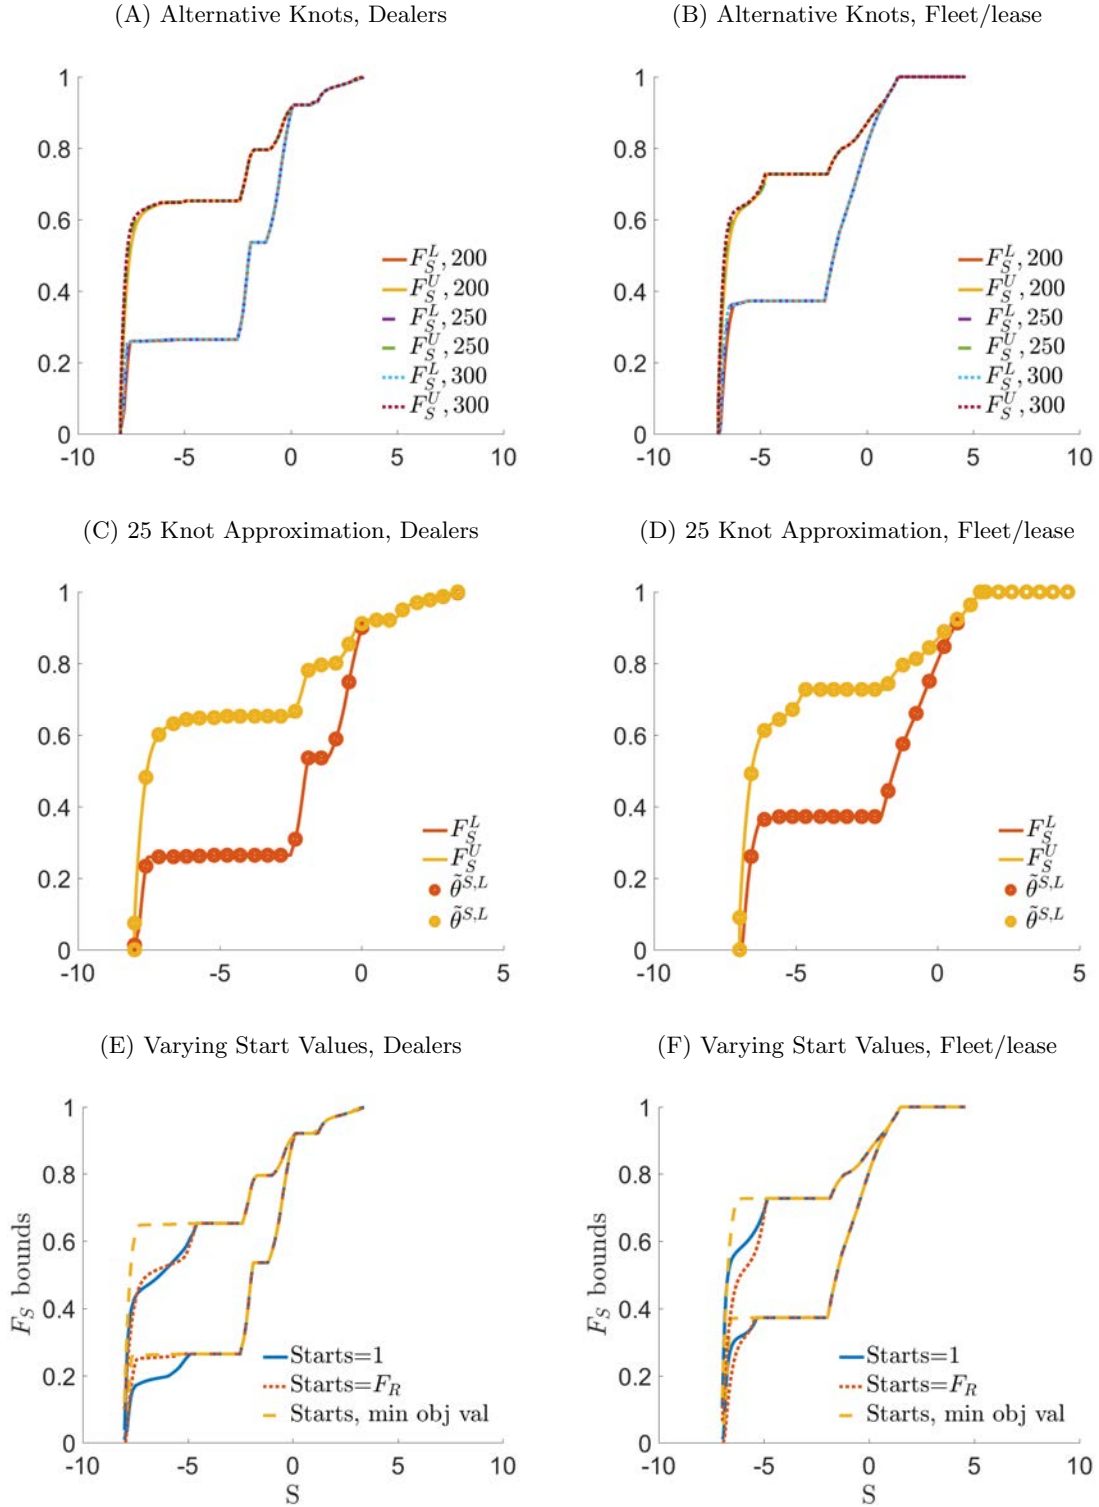
Notes: Panels A and B display the CDF of $R - P^A$ based on draws from the estimated F_R and F_{P^A} distributions compared to the CDF of the difference between the raw values in the data ($R^{raw} - P^{A,raw}$). Panels C and D display the estimated conditional probability statements, $\Pr(D^S = A | \tilde{P}^A)$ and $\Pr(D^S \neq Q | \tilde{P}^A)$, compared to the fitted estimates (from the right-hand-side quantities in (5) used to estimate F_S^L and F_S^U). Panels E and F display the estimated conditional probability statement $\Pr(D_1^B = 0 | \tilde{P}^A, \tilde{P}^A < \tilde{R})$ compared to the fitted estimate (from the right-hand-side of (6) used to estimate χ^{-1}). Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Figure A6: Robustness of Buyer Distribution Estimates



Notes: Panels A and B display estimates of F_B under different distributions for the number of bidders. Values for λ represent the mean number of bidders under a Poisson distribution. \bar{N} and \underline{N} represent distribution of the upper bound and lower bound on the number of bidders derived from bid log data. Panels C and D display the distribution of the maximum order statistic, $B^{(1)}$. Panels E and F display estimates of $\psi_{n-1:n}^{-1}(F_{B^{n-1:n+W}}(\cdot))$ for varying n , following the logic proposed in Aradillas-López, Gandhi, and Quint (2016) using the bid log subsample. Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Figure A7: Robustness of Seller Distribution Bound Estimates



Notes: Panels A and B display bounds on seller distribution using linear splines with 200 knots (solid lines), as in the main results displayed in Figure 1; 250 knots (dashed lines); 300 knots (dotted lines). Panels C and D display 25 uniformly spaced knots from the main 200 knots and their corresponding estimated coefficients. Panels E and F display the main bounds on F_S (solid lines), obtained using start values of $F_S^L(v) = F_S^U(v) = 1$ everywhere, as well as estimates under two other start values (dotted and dashed lines) discussed in Appendix C.4. Panels on the left use dealers sample and on the right use fleet/lease sample. Units = \$1,000.

Table A1: Theoretical Incomplete-Information Bargaining Literature

		One-sided offers	Alternating offers
One-sided incomplete information	2-types	Fudenberg and Tirole (1983) Sobel and Takahashi (1983) Admati and Perry (1987)	Rubinstein (1985, 1985b) Bikhchandani (1992)
	Cont. types	Fudenberg, Levine, and Tirole (1985, 1987) Gul, Sonnenschein, and Wilson (1986) Ausubel and Deneckere (1989)	Grossman and Perry (1986) Gul and Sonnenschein (1988) Cramton (1991)
Two-sided incomplete information	2-types	Fudenberg and Tirole (1983) Feinberg and Skrzypacz (2005)	Chatterjee and Samuelson (1987,1988) Compte and Jehiel (2002) Gantner (2008)
	Cont. types	Cramton (1984) Cho (1990) Ausubel and Deneckere (1992, 1993)	Cramton (1992) Perry (1986) Watson (1998) Abreu and Gul (2000)

Notes: While by no means exhaustive, this table presents a list of a variety of theoretical papers on incomplete-information bargaining settings, demonstrating that most of these papers do not model such games as being cases of two-sided incomplete information with a continuum of buyer/seller valuations, where both parties can make offers. Rather, the literature focuses primarily on settings of one-sided uncertainty (Fudenberg and Tirole 1983; Sobel and Takahashi 1983; Fudenberg, Levine, and Tirole 1985; Rubinstein 1985a; Rubinstein 1985b; Gul, Sonnenschein, and Wilson 1986; Grossman and Perry (1986); Admati and Perry 1987; Gul and Sonnenschein 1988; Ausubel and Deneckere 1989; Cramton 1991; Bikhchandani 1992), settings of one-sided offers (Cramton 1984; Cho 1990; Ausubel and Deneckere 1993; Feinberg and Skrzypacz 2005), settings with two-types rather than a continuum of types (Chatterjee and Samuelson 1988; Compte and Jehiel 2002), or settings with uncertainty not being about valuations (Watson 1998; Abreu and Gul 2000). Two papers that did model bargaining as an alternating-offer game and a continuum of types with two-sided incomplete information, where the incomplete information is about players' valuations, are Perry (1986), which predicted immediate agreement or disagreement, and Cramton (1992), which modeled the bargaining game as beginning with a war of attrition and consisting of players signaling their valuations through the length of delay between offers, as in Admati and Perry (1987). An additional line of research considers *static* bargaining games with two-sided incomplete information referred to as k double auctions (see Chatterjee and Samuelson 1983 and Satterthwaite and Williams 1989), discussed in the body of the paper. See Binmore, Osborne, and Rubinstein (1992), Kennan and Wilson (1993), Roth (1995), and Ausubel, Cramton, and Deneckere (2002) for additional surveys of the theoretical and experimental bargaining literature.

Table A2: Observations Dropped Through Sample Restrictions

A. Drops in combined sample (original combined sample size = 1,008,847)		# Obs.
Missing blue book or odometer reading		158,520
Missing timestamp for run		21,246
Misrecorded VIN		126
Duplicate record at location-VIN-seller-time level		4
Car switched from seller A to seller B than back to A		1,578
Missing auction price and reserve price		27,027
Extreme age (<0 or >16 years) or odometer reading (<100 or >300,000 miles)		5,855
Inconsistent data on whether trade occurred or not		18,309
Bargaining begins with non-high-bidder offer		10,269
Bargaining sequence includes “back-up” offer(s)		2,595
Inexplicable bargaining sequence		179
Misrecorded prices (obvious mistakes of extra (or missing) zeros)		150
Final price not equal to auction price or final negotiated price		1,423
Final price below auction price		9
Total observations remaining		761,557
Dealers observations remaining		427,607
Fleet/lease observations remaining		333,950

B. Drops within dealers vs. fleet/lease samples	Dealers		Fleet/lease	
	# Obs.	Cutoff	# Obs.	Cutoff
Blue book $\leq q(.01)$	4,328	375	3,496	550
Blue book $\geq q(0.99)$	4,285	25,550	3,351	30,825
Final price $\leq q(.01)$	1,567	300	2,175	400
Final price $\geq q(0.99)$	760	24,300	795	29,500
Auction price $\leq q(.01)$	5,661	0	5,554	1
Auction price $\geq q(0.99)$	292	25,000	291	29,000
Reserve price $\leq q(.01)$	3,279	750	5,175	0
Reserve price $\geq q(0.99)$	898	27,500	204	31,000
Counteroffers $\leq q(.01)$ or $\geq q(0.99)$ of auction price	70		42	
Day with < 100 cars sold	9,890		20,737	
Make-model-year-trim-age combo obs < 10 times	95,837		80,474	
Total observations remaining	300,740		211,656	

C. Obs that can be used only in Step 1 regression	Dealers	Fleet/lease
Reserve price recorded, auction price missing	148,075	15,034
Auction price recorded, reserve price missing	19,142	65,179
Total observations remaining for full estimation	133,523	131,443

Notes: Panel A displays the number of observations that each sample restriction (enforced in the order shown) drops from the original sample. Appendix C.1.1 describes these restrictions. Panel B displays the number of observations dropped from additional sample restrictions enforced (in the order shown) separately for the dealers and fleet/lease samples. For sample restrictions involving trimming at the 0.01 or 0.99 quantiles (denoted $q(0.01)$ and $q(0.99)$), Panel B also lists values of those quantiles. Panel C displays the number of observations for which the reserve price or auction price is missing. These observations can be used in the step 1 regression but not in the subsequent steps. The final sample size is shown in the last row.

Table A3: Additional Descriptive Statistics

	Dealers		Fleet/lease	
A. Trade Sample	Mean	S.D.	Mean	S.D.
Reserve price	\$6,978	\$4,926	\$10,401	\$5,943
Auction price	\$6,057	\$4,673	\$10,246	\$5,996
Final price	\$6,075	\$4,681	\$10,255	\$5,994
Buyer fee	\$162	\$36	\$209	\$45
Seller fee	\$146	\$54	\$94	\$15
Number of periods	1.957	0.587	1.591	0.579
Blue book	\$6,435	\$4,603	\$11,173	\$6,335
Age (years)	7.021	3.278	3.216	2.539
Odometer (miles)	100,875	45,690	56,001	39,637
Sample size	94,170		100,983	
B. No-trade Sample	Mean	S.D.	Mean	S.D.
Reserve price	\$8,427	\$5,663	\$9,996	\$5,232
Auction price	\$6,720	\$5,316	\$8,337	\$5,104
Number of periods	2.431	0.767	2.382	0.681
Blue book	\$7,740	\$5,215	\$10,212	\$5,397
Age (years)	6.166	3.503	3.051	2.590
Odometer (miles)	90,909	47,471	62,390	42,419
Sample size	39,353		30,460	

Notes: Mean and standard deviation of variables in dealers and fleet/lease samples. Number of periods is 1 if game ends through auction price exceeding reserve, 2 if seller accepts at her first bargaining turn, etc. Blue book is an estimate of the market value of the car, provided by the auction house. Panel A displays subsample where trade occurs and panel B displays subsample where no trade occurs.

Table A4: Outcomes of Game By Period: Fleet/lease Sample

Ending period	Player's turn	# Obs	% of Sample	% Trade	<u>Full Sample</u>		<u>Conditional on Sale</u>		
					Reserve price	Auction price	Reserve price	Auction price	Final price
1	(Auction)	45,204	34.391%	98.35%	\$10,202 (\$6,044)	\$10,969 (\$6,154)	\$10,246 (\$6,053)	\$11,063 (\$6,139)	\$11,063 (\$6,139)
2	S	73,718	56.084%	74.01%	\$10,694 (\$5,764)	\$9,647 (\$5,700)	\$10,586 (\$5,885)	\$9,683 (\$5,833)	\$9,683 (\$5,833)
3	B	10,630	8.087%	10.86%	\$8,345 (\$4,395)	\$6,426 (\$3,986)	\$8,533 (\$4,260)	\$7,357 (\$4,113)	\$7,817 (\$4,189)
4	S	896	0.682%	55.25%	\$8,500 (\$4,615)	\$6,819 (\$4,214)	\$8,845 (\$4,824)	\$7,219 (\$4,383)	\$7,539 (\$4,472)
5	B	920	0.700%	31.20%	\$8,884 (\$4,650)	\$7,187 (\$4,268)	\$9,385 (\$4,923)	\$7,888 (\$4,572)	\$8,509 (\$4,687)
6	S	36	0.027%	52.78%	\$8,775 (\$4,196)	\$6,968 (\$4,334)	\$8,726 (\$3,565)	\$7,124 (\$3,801)	\$7,711 (\$3,836)
7	B	35	0.027%	31.43%	\$10,454 (\$6,339)	\$8,587 (\$5,790)	\$12,991 (\$9,732)	\$10,632 (\$8,854)	\$11,727 (\$9,237)
8	S	2	0.002%	0.00%	\$16,250 (\$7,425)	\$14,500 (\$6,364)	.	.	.
9	B	2	0.002%	0.00%	\$11,750 (\$8,132)	\$9,925 (\$7,743)	.	.	.

Notes: Fleet/lease sample. For each period (period 1 = auction, period 2 = seller's first turn in bargaining, period 3 = buyer's turn, etc.), table reports the number of observations ending in that period, percent of total sample ending in that period, and percent of cases in which trade occurred. Table also reports reserve price and auction price for observations ending in a given period and, for those observations ending in trade, the reserve price, auction price, and final price conditional on trade. Corresponding statistics for the dealers sample are found in Table 2.

Table A5: Model Dispersion and Fit

A. Dispersion	Dealers			Fleet/lease		
	Variance	0.1 Quantile	0.9 Quantile	Variance	0.1 Quantile	0.9 Quantile
$X'\hat{\gamma}$	24.760	2.007	14.593	32.381	3.289	18.072
W	0.728	-0.979	0.885	0.478	-0.698	0.656
R	0.515	-0.379	1.102	0.908	-0.731	1.242
P^A	0.653	-1.642	0.074	0.967	-1.302	0.836

B. $\text{Corr}(R + W, P^A + W)$	Dealers	Fleet/lease
$\text{Corr}(R + W, P^A + W)$ from Step 1 residuals	0.513	0.293
$\text{Corr}(R + W, P^A + W)$ simulated	0.557	0.341

C. Probability of Negative Prices	Dealers	Fleet/lease
$\Pr(X'\gamma + W + R < 0)$	0.008	0.005
$\Pr(X'\gamma + W + P^A < 0)$	0.042	0.010

D. Integrated Absolute Error	Dealers	Fleet/lease
$F_S^L(\cdot)$	0.0039	0.0030
$F_S^U(\cdot)$	0.0027	0.0017
$F_S^L(\cdot)$ violations	0.0017	0.0020
$F_S^U(\cdot)$ violations	0.0018	0.0011
$\chi^{-1}(\cdot)$	0.0005	0.0003
$g(\cdot)$	0.0600	0.0247

Notes: Panel A displays the variance, 0.1 quantile, and 0.9 quantile of $X\gamma$, W , R , and P^A , simulated by drawing with replacement J draws from the estimated distribution of each of these objects, where J is the sample size ($J = 133,523$ in the dealers sample and 131,443 in the fleet/lease sample). In the case of $X\gamma$, I sample directly from the estimates of $X\gamma$ from the step 1 regression. Panel B compares the correlation in the residuals from the regression from estimation step 1 to the correlation in simulated draws from the estimated distributions for R , P^A , and W . Panel C shows the model's predicted probability of observing a raw reserve price or auction price that is not positive. Panel D displays the integrated absolute error from the constrained least squares problems in steps 4–6 of the estimation. Rows labeled *violations* show this same measure but where the only errors included in the integration are those that constitute violations of the conditional probability bounds in (4).

Table A6: Expected Gains From Trade Using Alternative Sample Restrictions

A. Dealers Sample	Ex-post	Second-best	Real bargaining	Ex-post minus real
Less Expensive Cars	[1.711, 2.933] (1.666, 3.053)	[1.694, 2.926] (1.645, 3.044)	[1.421, 2.400] (-0.049, 2.482)	[0.241, 0.583] (0.225, 3.009)
Newer Cars	[2.492, 5.808] (2.142, 6.104)	[2.434, 5.796] (2.110, 6.093)	[1.988, 4.397] (1.744, 4.632)	[0.471, 1.445] (0.324, 1.549)
More Cars Remaining of Same Make-Model	[2.565, 5.261] (2.169, 5.498)	[2.527, 5.238] (2.125, 5.483)	[2.135, 4.210] (1.819, 4.444)	[0.399, 1.081] (0.295, 1.128)
More Cars Remaining of Same Seller	[2.137, 5.228] (2.016, 5.597)	[2.083, 5.209] (1.979, 5.577)	[1.664, 3.871] (1.579, 4.173)	[0.445, 1.386] (0.350, 1.476)
First Run of Car	[2.605, 4.804] (1.997, 5.123)	[2.569, 4.792] (1.965, 5.110)	[2.314, 4.220] (1.786, 4.485)	[0.241, 0.633] (0.163, 0.711)
B. Fleet/lease Sample				
Less Expensive Cars	[1.896, 3.604] (1.843, 3.741)	[1.832, 3.574] (1.779, 3.710)	[1.622, 2.846] (1.550, 2.905)	[0.232, 0.800] (0.227, 0.897)
Newer Cars	[2.583, 4.316] (2.382, 4.604)	[2.524, 4.291] (2.328, 4.596)	[2.223, 3.422] (1.958, 3.722)	[0.321, 0.934] (0.311, 1.092)
More Cars Remaining of Same Make-Model	[2.997, 4.136] (2.762, 4.389)	[2.951, 4.121] (2.708, 4.375)	[2.616, 3.549] (2.369, 3.685)	[0.314, 0.654] (0.289, 0.801)
More Cars Remaining of Same Seller	[2.800, 4.178] (2.403, 4.443)	[2.748, 4.155] (2.355, 4.416)	[2.442, 3.612] (2.083, 3.747)	[0.290, 0.634] (0.255, 0.827)
First Run of Car	[2.663, 3.997] (2.385, 4.101)	[2.614, 3.978] (2.333, 4.081)	[2.369, 3.550] (2.126, 3.614)	[0.231, 0.510] (0.201, 0.575)

Notes: Table displays expected gains from trade as in Tables 4–5 but using specific subsamples of the data, described in Appendix B.4 and C.1.3. Panel A contains dealer sellers and panel B fleet/lease sellers. Estimated bounds are in square braces and 95% confidence set is in parentheses. Units are \$1,000.

Table A7: Expected Gains and Probability of Trade Under Different $\Pr(N = n)$

A. Expected Gains	Ex-post	Second-best	Real bargaining	Ex-post minus real
Dealers Sample				
\underline{N}	[2.442, 5.045]	[2.397, 5.029]	[1.993, 3.933]	[0.422, 1.139]
$\lambda = 3$	[2.367, 4.965]	[2.328, 4.952]	[1.946, 3.892]	[0.394, 1.101]
$\lambda = 7$	[2.321, 4.910]	[2.286, 4.899]	[1.923, 3.853]	[0.370, 1.084]
$\lambda = 10$	[2.318, 4.905]	[2.284, 4.894]	[1.922, 3.851]	[0.369, 1.082]
$\lambda = 20$	[2.317, 4.904]	[2.283, 4.893]	[1.921, 3.851]	[0.368, 1.081]
\overline{N}	[2.251, 4.832]	[2.222, 4.823]	[1.874, 3.780]	[0.350, 1.080]
Fleet/lease Sample				
\underline{N}	[2.408, 4.195]	[2.342, 4.165]	[2.080, 3.370]	[0.289, 0.864]
$\lambda = 3$	[2.305, 4.089]	[2.248, 4.063]	[1.984, 3.291]	[0.282, 0.837]
$\lambda = 7$	[2.229, 4.006]	[2.179, 3.983]	[1.923, 3.238]	[0.267, 0.807]
$\lambda = 10$	[2.225, 4.002]	[2.176, 3.979]	[1.924, 3.242]	[0.262, 0.798]
$\lambda = 20$	[2.223, 3.999]	[2.174, 3.976]	[1.921, 3.236]	[0.264, 0.802]
\overline{N}	[2.114, 3.883]	[2.073, 3.864]	[1.828, 3.147]	[0.247, 0.774]
B. Probability of Trade	Ex-post	Real bargaining	Ex-post minus real	
Dealers Sample				
\underline{N}	[0.818, 0.871]	[0.646, 0.646]	[0.172, 0.225]	
$\lambda = 3$	[0.804, 0.867]	[0.642, 0.642]	[0.162, 0.225]	
$\lambda = 7$	[0.793, 0.862]	[0.637, 0.637]	[0.155, 0.225]	
$\lambda = 10$	[0.792, 0.862]	[0.637, 0.637]	[0.155, 0.225]	
$\lambda = 20$	[0.792, 0.862]	[0.638, 0.638]	[0.154, 0.225]	
\overline{N}	[0.777, 0.857]	[0.628, 0.628]	[0.149, 0.229]	
Fleet/lease Sample				
\underline{N}	[0.857, 0.893]	[0.658, 0.658]	[0.199, 0.235]	
$\lambda = 3$	[0.844, 0.885]	[0.655, 0.655]	[0.188, 0.230]	
$\lambda = 7$	[0.830, 0.878]	[0.651, 0.651]	[0.179, 0.227]	
$\lambda = 10$	[0.829, 0.878]	[0.652, 0.652]	[0.177, 0.225]	
$\lambda = 20$	[0.829, 0.878]	[0.651, 0.651]	[0.178, 0.226]	
\overline{N}	[0.812, 0.869]	[0.644, 0.644]	[0.168, 0.225]	

Notes: This table displays bounds on the expected gains from trade (panel A) and probability of trade (panel B) estimated by altering the specification of $\Pr(N = n)$ in estimation step 3, yielding different estimates of F_B . These F_B estimates are then used in re-doing estimation steps 5–6 and in computing counterfactuals. \underline{N} denotes the lower bound based on bid log data and \overline{N} the upper bound. The λ rows denote truncated Poisson approximations (i.e. restricted to have $N \geq 2$) with mean $\lambda = 3, 7, 10, \text{ or } 20$.

Table A8: Auction House Fees and the Broker-Optimal Mechanism

A. Dealers Sample	Ex-post without fee	Broker- optimal	Real bargaining
Expected gains from trade	[2.702, 5.317] (2.451, 5.517)	[1.806, 4.371] (1.503, 4.613)	[2.192, 4.133] (2.012, 4.285)
Buyer gains		[0.215, 0.475] (0.181, 0.482)	[0.822, 0.845] (0.779, 0.913)
Seller gains		[0.321, 0.245] (0.169, 0.394)	[1.171, 3.088] (0.947, 3.202)
Probability of trade	[0.862, 0.891] (0.849, 0.922)	[0.271, 0.560] (0.226, 0.567)	[0.646, 0.646] (0.638, 0.672)
Auction house revenue		[1.270, 3.652] (1.001, 3.799)	[0.199, 0.199] (0.196, 0.207)
B. Fleet/lease Sample			
Expected gains from trade	[2.675, 4.470] (2.435, 4.639)	[1.711, 3.382] (1.594, 3.682)	[2.279, 3.569] (2.042, 3.668)
Buyer gains		[0.341, 0.596] (0.326, 0.649)	[1.158, 1.192] (1.073, 1.219)
Seller gains		[0.236, 0.322] (0.202, 0.504)	[0.922, 2.178] (0.754, 2.302)
Probability of trade	[0.903, 0.924] (0.891, 0.929)	[0.303, 0.501] (0.289, 0.556)	[0.658, 0.658] (0.653, 0.677)
Auction house revenue		[1.134, 2.464] (0.965, 2.597)	[0.199, 0.199] (0.198, 0.205)

Notes: Bounds on expected gains from trade and auction house revenue (in \$1,000 units) and on probability of trade from the ex-post efficient mechanism in the absence of auction house fees, the broker-optimal mechanism, and the real-world bargaining. Real-world bargaining estimates come from panel A of Tables 4–5 with one exception: auction house revenue in this table is included in the total expected gains from trade. Panel A displays dealers sample and panel B fleet/lease. Estimated bounds are in square braces and 95% confidence set is in parentheses.