The Efficiency of Dynamic, Post-Auction Bargaining: Evidence from Wholesale Used-Auto Auctions

Bradley Larsen*

Stanford University, NBER and eBay Research Labs

(MID-REVISION DRAFT—PRELIMINARY AND INCOMPLETE)

June 2014

Abstract

This study quantifies the efficiency of a real-world bargaining game with two-sided incomplete information. Myerson and Satterthwaite (1983) and Williams (1987) derived the theoretical efficient frontier for bilateral trade under two-sided uncertainty, but little is known about how well real-world bargaining performs relative to the frontier. The setting is wholesale used-auto auctions, an $80 billion industry where buyers and sellers participate in alternating-offer bargaining when the auction price fails to reach a secret reserve price. Using 270,000 auction/bargaining sequences, this study nonparametrically estimates bounds on the distributions of buyer and seller valuations and then estimates where bargaining outcomes lie relative to the efficient frontier. Findings indicate that the observed auction-followed-by-bargaining mechanism is quite efficient, achieving a large fraction of the surplus and trade volume which can be achieved on the efficient frontier.

*I thank Panle Jia Barwick, Glenn Ellison, and Stephen Ryan for invaluable help and advice throughout this project. I would also like to thank Isaiah Andrews, Victor Chernozhukov, Joachim Freyberger, Ariel Pakes, Paulo Somaini, and Juuso Toikka for extremely helpful suggestions. This paper also benefited from conversations with many MIT graduate students and faculty, as well as John Asker, Matias Cattaneo, Mingli Chen, Denis Chetverikov, Peter Cramton, Francesco Decarolis, Bernard Elyakime, Ian Gale, David Genesove, Ken Hendricks, Kate Ho, Kyoo-il Kim, Elena Krasnokutskaya, Katie Larsen, Greg Lewis, Thierry Magnac, Francesca Molinari, Christopher Palmer, Christopher Parmeter, Parag Pathak, Jimmy Roberts, Mark Satterthwaite, Alan Sorensen, Matt Shum, Jean Tirole, and many seminar and conference participants. I thank several anonymous auction houses for providing extensive data and data help. I also thank numerous employees and customers of these auction houses, as well as Freeman Dawson and Steven Lang, for providing institutional details. I acknowledge support from the National Science Foundation Graduate Research Fellowship under Grant No. 0645960.

1 Introduction

From haggling in an open-street market to negotiating a corporate takeover deal, alternating-offer bargaining between a buyer and seller is one of the oldest and most common forms of transaction. When both parties have incomplete information, it is known that equilibrium outcomes are difficult to characterize.[1] Myerson and Satterthwaite (1983) demonstrated that full efficiency is not possible, and theoretical efficiency bounds are derived in Myerson and Satterthwaite (1983) and Williams (1987), but it is unknown how well real-world bargaining performs relative to these bounds. Williams (1987) emphasized that "little is known about whether or not these limits can be achieved with 'realistic' bargaining procedures." This paper is the first attempt to bring data to this question in order to quantify the efficiency of real-world bargaining with two-sided incomplete information. I develop a framework to estimate distributions of private valuations on both sides of the market at wholesale used-auto auctions. I then map these primitives into results from the theoretical mechanism design literature in order to measure the efficiency of bargaining relative to the first-best frontier and the information-constrained, second-best efficient frontier.

The data analyzed in this paper consist of several hundred thousand sequences of back-and-forth bargaining offers between buyers and sellers at wholesale used-auto auctions, a large market where new and used-car dealers buy vehicles from other dealers as well as from rental companies and banks. This industry passes 15 million cars annually through its lanes, selling about 60% of the vehicles, worth a total of $80 billion. Throughout much of the industry, auction houses employ the following mechanism: a secret reserve price set by the seller followed by an ascending price auction, which, when the secret reserve price is not met, is followed by post-auction, alternating-offer bargaining mediated by the auction house. This setting is ideal for studying bargaining under two-sided uncertainty because all players are experienced professionals and are likely to understand well the game being played. Also, because the bargaining takes place after an ascending auction and after the seller's choice of a secret reserve price, the efficiency of bargaining can be studied while imposing only minimal assumptions on the structure or equilibrium of the bargaining game.

After a brief introduction to the industry, Section 2 discusses the data in detail. The data comes from several different auction houses from 2007 to 2010, containing detailed information on each car as well as the actions taken by sellers and buyers in each stage of the game. The data is broken down into two main samples: cars sold by used and new-car dealers (which I refer to as the dealers sample), and cars sold by large institutions, such as rental companies, banks, and fleet companies (which I refer to as the fleet/lease sample). Reserve prices exceed auction prices by an average of $1,000, with this gap being higher in the dealers sample than the fleet/lease sample. The game ends at the auction, with no recorded bargaining, approximately 76% of the time in the dealers sample and 82% of the time in the fleet/lease

[1] Fudenberg 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."

sample. When the game ends in trade at a bargained price, this price exceeds the auction price by $600 on average.

I lay out a simple model in Section 3 which describes the three stages of the game at wholesale auto auctions. The post-auction bargaining is modeled as a general alternating-offer bargaining game. The auction stage is modeled as an ascending auction with symmetric private values among bidders, where bidders' values and the seller's value are correlated through auction-level heterogeneity unobserved by the econometrician. The seller's secret reserve price is chosen optimally before the auction. I prove two preliminary results which motivate an estimation strategy: first, truth-telling is a dominant strategy for bidders in the auction, and second, the seller's secret reserve price strategy is monotone. These two properties allow for nonparametric identification of the distributions of buyer and seller types.

Sections 4 and 5 constitute the heart of the paper. Section 4 presents the approach for estimating the distributions of buyer and seller valuations. After controlling for observable covariates and auction house fees, I account for unobserved heterogeneity at the auction level through a semi-nonparametric approach, relying on deconvolution argument due to Kotlarski (1967) for identification. I then estimate the distribution of buyer valuations using an order statistics inversion. The approach for estimating the distribution of seller types is new. It exploits the property that the seller's secret reserve price will be strictly monotonic. I use bounds defined by revealed preferences arguments taken from the seller's response to the auction price. The approach is similar in spirit to Haile and Tamer (2003), using bounds implied by very basic assumptions about players' rationality to learn about model primitives without solving for the equilibrium of the game.

Section 5 presents the methods for estimating the efficient frontier and other counterfactual mechanisms from mechanism design theory. It is important to note that throughout the paper, the terms "efficient" or "second-best" refer to ex-ante incentive efficiency, taking informational asymmetries into account. To refer to full efficiency, I use the terms "ex-post efficient" or "first-best."[2] I also use the terms "surplus" and "gains from trade" interchangeably. The efficient frontier (or Pareto frontier) delineates the best possible outcomes, in terms of buyer and seller surplus, that could be achieved by any bilateral bargaining game in the presence of two-sided incomplete information. Myerson and Satterthwaite (1983) and Williams (1987) demonstrated how this frontier depends on the distributions of buyer and seller valuations. Therefore, the estimated distributions from Section 4 are crucial for solving for these mechanisms. To solve for the Pareto frontier, I adopt an approach described in Williams (1987). I also demonstrate that the direct revelation mechanism corresponding to the dynamic bargaining mechanism currently used at auto auctions is identified from information on whether trade occurred and from secret reserve and auction prices.

I combine the estimated distributions of buyer and seller valuations and the allocation functions defining the current mechanism and counterfactual mechanisms, putting theoretical bargaining and real-world bargaining on the same footing in order to quantify bargaining efficiency. First, I examine the

[2]For a more detailed taxonomy of ex-ante, interim, and ex-post efficiency under incomplete information, see Holmström and Myerson (1983).

efficiency loss due to incomplete information. Ideally, a buyer and seller should trade whenever the buyer values a good more than the seller (first-best, ex-post efficient trade). However, incomplete information on both sides gives rise to a bilateral monopoly, where each party has some market power. Myerson and Satterthwaite (1983) demonstrated that a deadweight loss occurs as each party trades off the dual incentives of increasing the probability of trade and extracting additional rent from the other party, akin to the deadweight loss in a standard one-sided monopoly pricing model. As a result, some trades fail to occur even when the buyer values the good more than the seller.[3] This deadweight loss is given by the gap between the the second-best mechanism derived in Myerson and Satterthwaite (1983) and first-best trade. I discover that incomplete information need not be a huge problem in this market: The second-best mechanism achieves 96–99% of first-best surplus.

Second, I examine the efficiency of post-auction bargaining relative to the information-constrained efficient frontier. Unlike the mechanisms discussed in Myerson and Satterthwaite (1983) and Williams (1987), alternating-offer bargaining with two-sided uncertainty has no clear equilibrium predictions due to signaling by both parties. As a result, it is unknown where alternating-offer bargaining lies within the efficient frontier. Any gap between the efficient frontier and real-world bargaining represents a deadweight loss which could theoretically be eliminated by committing to a static efficient mechanism along the frontier. Therefore, I refer to this as the deadweight loss due to mechanism choice/limited commitment.[4] Findings indicate that the post-auction bargaining lies quite close to the efficient frontier, achieving 90–95% of the efficient level of surplus. This result is true in both the dealers sample and the fleet/lease sample. The deadweight loss due to limited commitment is therefore quite small in this market.

In addition to Myerson and Satterthwaite (1983) and Williams (1987), other papers examining the theoretical efficient frontier from a mechanism design standpoint include Ausubel and Deneckere (1993), Ausubel, Cramton, and Deneckere (2002), Satterthwaite and Williams (1989), and Chatterjee and Samuelson (1983). Ausubel and Deneckere (1993) and Ausubel, Cramton, and Deneckere (2002) demonstrated theoretically that when buyer and seller distributions have monotone hazard rates and when high weights are placed on the seller or buyer payoff, some equilibria of a dynamic, alternating-offer bargaining game can reach the efficient frontier. Satterthwaite and Williams (1989) studied the k double auction game and found that generically only the $k = 0$ or $k = 1$ double auctions reach the efficient frontier.[5] Chatterjee and Samuelson (1983) demonstrated that in the symmetric uniform case the $k = 1/2$ double auction also reaches the efficient frontier.

A large theoretical literature on incomplete information bargaining has yielded valuable insights through focusing on special cases rather than the full, two-sided incomplete information setting.[6] The

[3] Formally, Myerson and Satterthwaite (1983) demonstrated that when the supports of buyer and seller types overlap, there does not exist an incentive-compatible, individually rational mechanism which is ex-post efficient and which also satisfies a balanced budget.

[4] Cramton (1992) and Elyakime, Laffont, Loisel, and Vuong (1997) also referred to this as an issue of commitment.

[5] A k double auction consists of both the buyer and seller simultaneously reporting sealed bids to an intermediary and, if the buyer's bid exceeds the seller's, trading at a price which is a convex combination of the two bids, with weight k.

[6] The incomplete-information bargaining literature focuses on settings of one-sided uncertainty (Gul, Sonnenschein, and

set of empirical papers which structurally estimated models of incomplete information bargaining is quite small, including Sieg (2000), Ambrus, Chaney, and Salitsky (2011), and Silveira (2012), who focused on settings of one-sided incomplete information bargaining, and Keniston (2011), who estimated a model of two-sided uncertainty and compared alternating-offer bargaining to a fixed-price mechanism.[7] One advantage of the current paper over previous structural papers is that, because the bargaining occurs after an auction and after the seller reports a secret reserve price, the model's primitives, namely the distributions of buyer and seller valuations, can be identified from these pre-bargaining actions without relying on much structure or a particular equilibrium notion for the bargaining game. This is particularly useful given that, unlike auction settings or complete information bargaining games, there is no canonical model of alternating-offer bargaining under incomplete information with a continuum of types.[8]

2 The Wholesale Auto Auction Industry

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. 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. The industry leaders, Manheim and Adesa, maintain a 50% and 25% market share, respectively, and the remaining auction houses are referred to as independent. Each auction house serves as a platform in a two-sided market, competing to attract both sellers and buyers. 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 new and used car

Wilson 1986; Gul and Sonnenschein 1988; Fudenberg and Tirole 1983; Sobel and Takahashi 1983; Fudenberg, Levine, and Tirole 1985; Ausubel and Deneckere 1989; Rubinstein 1985a,b; Bikhchandani 1992; Grossman and Perry 1986; Admati and Perry 1987; Cramton 1991), settings of one-sided offers (Cho 1990; Feinberg and Skrzypacz 2005; Cramton 1984; Ausubel and Deneckere 1993), 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 (Abreu and Gul 2000; Watson 1998). Two papers which modeled 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). Neither of these models fits the type of bargaining observed at wholesale auto auctions. See Binmore, Osborne, and Rubinstein (1992), Ausubel, Cramton, and Deneckere (2002), Roth (1995), and Kennan and Wilson (1993) for additional surveys of the theoretical and experimental bargaining literature.

[7] A separate strand of literature, such as Tang and Merlo (2010) and Watanabe (2009), discussed the identification and estimation of complete information rather than incomplete information games.

[8] Several papers analyzed models of auctions followed by *complete* information bargaining (Huh and Park 2010; Elyakime, Laffont, Loisel, and Vuong 1997), one-shot bargaining (Bulow and Klemperer 1996; Eklof and Lunander 2003; Menezes and Ryan 2005), or one-sided incomplete information (Wang 2000). Genesove (1991) also discussed post-auction bargaining at wholesale auto auctions. He tested several parametric assumptions for the distributions of buyer and seller valuations, finding that these assumptions performed poorly in explaining when bargaining occurred and when it was successful.

dealers.[9] Sellers may also be used or new car dealers (whome 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, or old fleet vehicles. Throughout, I will refer to the seller as "she" and the buyer as "he".

Sellers bring their cars to the auction house and most report a secret reserve price.[10] In the days preceding the auctioning of the car, 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.[11] 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. The characteristics of the car as well as the current high bid are listed on a large monitor near the auctioneer. The entire bidding process usually takes 30–90 seconds.

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, 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, at which point the seller can accept the auction price, end the negotiations, or propose a counteroffer.[12] If the seller counters, the auction house calls the buyer. Bargaining continues until one party accepts or terminates negotiations. The typical time between calls is 2–3 hours.[13] Auction house employees contacting each party take care not to reveal the other party's identity in order to prevent the buyer and seller from agreeing on a trade outside of the auction house, avoiding auction house fees.[14]

A seller accepting the auction price (or bargaining offers) below the reserve price may seem puzzling

[9] Note that the term "new" means the dealer is authorized to sell new cars from a manufacturer, but can also sell used cars. On the other hand, "used" car dealers can only sell used cars. Genesove (1993) discussed the differences of cars sold by new vs. used-car dealers and found weak evidence of adverse selection among cars sold by used-car dealers. Note also that the general public is not allowed at these auctions; Hammond and Morrill (2012) presented a model explaining this feature of auto auctions.

[10] Some sellers choose not to report a reserve as they plan to either be present during the auction or to await a phone call from the auction house informing them of the auction price prior to deciding whether to accept.

[11] According to conversations with participants and personal observations at auction houses, few buyers appear to visit the auction house prior to the day of sale.

[12] If the seller a present during the auctioning of the car, the seller may choose to accept or reject the auction price immediately. If the seller is not present but the auctioneer observes that the auction price and the reserve price far enough apart that phone bargaining is very unlikely to succeed, the auctioneer may choose to immediately reject the auction price on behalf of the seller.

[13] During the time a car is in the bargaining process, or if that bargaining has ended in no trade, interested buyers other than the high bidder may also contact the auction house and place offers on the car. If the bargaining between the original high bidder and seller ends in disagreement, bargaining commences with the next interested buyer as long as his offer is higher than previous offers the seller has rejected. This occurs for about three percent of the cars in the full dataset. This separate form of dynamics is not accounted for in the model below, and hence the observations are not included in the analysis.

[14] If the seller is an institution then her identity is revealed, and institutions tend to try to build a positive reputation of not setting excessively high reserve prices. Conversations with industry participants reveal that at some auction houses outside of the data sample studied in this paper, the identity of small sellers is also revealed.

given that the seller could have accomplished an equivalent outcome by reporting a lower secret reserve price. Industry participants explain this phenomenon as having several potential causes: sellers systemically set high reserve prices due to overly optimistic beliefs about auction prices (Treece 2013) or in attempt to influence auctioneers to achieve higher prices (Lacetera, Larsen, Pope, and Sydnor 2014), or sellers are learning something about demand after observing the auction.[15]

If the auction and/or bargaining does not result in trade the first time the vehicle is up for sale (or first "run"), the vehicle can either be taken back to the seller's business location or, more often, remain at the auction house until the next available sales opportunity, usually the following week.[16] The seller can change her reserve price before the next run of the vehicle. If trade takes place but the buyer feels he has been given a lemon, he may, under certain conditions, request arbitration, in which the auction house intervenes to either undo the sale or negotiate a lower sale price.[17] This occurs less than three percent of the time in practice.

The data used in this paper come from six auction houses, each maintaining a large market share in the region in which it operates. Between January 2007 and March 2010 these auction houses passed over 600,000 vehicles through their lanes. The data from these auction houses includes detailed information on each car, including make, model, year, trim, and odometer reading; condition report (prepared by the auction house for fleet/lease vehicles); a blue book estimate provided by the auction house; the number of pictures displayed in the online pre-sale profile of the car; disclosure codes for different types of damage reported by the seller; and the identity of the seller.

An observation in the dataset represents a run of the vehicle, that is, a distinct attempt to sell the vehicle through the auction or, if the reserve price is not met, through post-auction bargaining. The total number of runs recorded in the data is approximately 1,000,000, so on average a vehicle passes through the lanes 1.67 times. I treat each run as an independent observation and do not model dynamics between runs. 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.

I drop observations with no recorded auction house blue book estimate; cars less than one year or greater than 16 years old; cars with less than 100 miles or greater than 300,000 miles on the odometer; observations in which the auction sale timestamp is missing; and observations for which the following variables lie outside their respective 0.01 and 0.99 quantiles: auction price, reserve price, blue book price,

[15]For the estimation approach I adopt in Section 4, knowing the precise data-generating process for reserve prices is unnecessary; it is only necessary that secret reserve prices be strictly increasing in the seller's underlying true valuation, a result generated by the model in Section 3. Appendix A.2 generalizes the theoretical results of Section 3 to a demand-learning case.

[16]Genesove (1995) presented a search model to study the seller's decision to reject the auction price and take the car back to her own car lot.

[17]For a buyer to be able to request arbitration, the car's sale price must be greater than $2,500 and the alleged lemon feature of the car must fall into certain pre-determined categories, such as structural damage, which was unreported by the seller.

Table 1: Descriptive Statistics for Auto Auction Data

	Dealers sample		Fleet/lease sample	
	mean	s.d.	mean	s.d.
Trade	0.711	0.453	0.760	0.427
Reserve price	$7,380	$5,195	$10,307	$5,764
Auction price	$6,257	$4,923	$9,826	$5,846
Blue book	$6,812	$4,853	$10,963	$6,140
Age (years)	6.779	3.380	3.150	2.558
Mileage	97,920	46,511	57,178	40,320
Sample size	136,135		136,459	

Notes: Trade is an indicator for whether trade occurred between the buyer and seller. Blue book is an estimate of the market value of the car, provided by the auction house.

or the gap between the reserve and auction price. I drop observations for which fewer than ten vehicles were observed at a given make-model-year-trim-age combination or days in which fewer than 100 cars were offered for sale at a given auction house. I drop observations where the auction price or reserve price is missing or is equal to zero and incomplete bargaining sequences.[18] In the end, I am left with 136,135 runs of cars sold by used-car dealers (which I will refer to as the dealers sample), and 136,459 sold by fleet/lease sellers (which I will refer to as the fleet/lease sample).

Summary statistics are displayed in Table 1. The probability of trade is 0.71 in the dealers sample and slightly higher in the fleet/lease sample. 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, in fleet/lease cars the reserve price does not exceed the blue book price on average.

Table 2 displays the characteristics of observations in the dealers sample which end at each period of the game, and similarly in Table 3 for the fleet/lease sample. For each period of the game, the columns

[18] Some observations record a secret reserve price or an auction price but not both. These observations, or observations with incorrectly recorded bargaining sequences (such as a seller acceptance followed by a buyer counteroffer) are not suitable for my final analysis but are still useful in controling for observable heterogeneity as explained in Section 4.2. Missing secret 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 ex-ante. Missing auction prices can occur due to 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 (2014).

Table 2: Outcome of game by period: Dealers sample

		Full dealers sample					Conditional on trade occurring		
Ending period	Player's turn	# Obs	% of Sample	% Trade	Reserve price	Auction price	Reserve price	Auction price	Final price
1	(Auction)	104,505	76.766%	80.67%	$7,458	$6,447	$6,949	$6,048	$6,048
				(39.49%)	($5,243)	($4,975)	($4,933)	($4,702)	($4,702)
2	S	11,075	8.135%	62.18%	$6,622	$5,285	$6,366	$5,306	$5,306
				(48.50%)	($4,952)	($4,646)	($4,862)	($4,626)	($4,626)
3	B	14,929	10.966%	13.85%	$7,194	$5,575	$7,769	$6,544	$6,959
				(34.54%)	($4,978)	($4,619)	($5,276)	($5,030)	($5,114)
4	S	3,293	2.419%	70.88%	$7,721	$6,312	$7,563	$6,249	$6,503
				(45.44%)	($5,147)	($4,844)	($5,015)	($4,728)	($4,803)
5	B	1,893	1.391%	44.43%	$8,175	$6,627	$8,519	$7,001	$7,643
				(49.70%)	($5,198)	($4,873)	($5,430)	($5,090)	($5,247)
6	S	247	0.181%	82.59%	$8,265	$6,747	$8,307	$6,818	$7,306
				(38.00%)	($5,353)	($5,051)	($5,437)	($5,143)	($5,244)
7	B	159	0.117%	60.38%	$8,349	$6,746	$8,554	$6,992	$7,743
				(49.07%)	($5,180)	($4,958)	($5,225)	($5,031)	($5,174)
8	S	25	0.018%	76.00%	$8,968	$7,576	$8,934	$7,663	$8,197
				(43.59%)	($5,045)	($5,001)	($5,599)	($5,557)	($5,628)
9	B	6	0.004%	66.67%	$7,583	$6,225	$6,675	$5,188	$5,838
				(51.64%)	($4,362)	($4,517)	($2,871)	($2,964)	($2,800)
10	S	3	0.002%	100.00%	$8,333	$6,100	$8,333	$6,100	$7,233
				(0.00%)	($5,620)	($4,453)	($5,620)	($4,453)	($5,701)

Notes: For each period (period 1 = auction and immediately following, period 2 = seller's first turn in post-auction bargaining, period 3 = buyer's first turn, etc.), table reports the number of observations ending in that period, percent of total sample ending in that period, and percent of time 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.

in these tables display the number of observations ending in that period, the percent of the total sample which this number represents, the percent of cases which ended in trade, as well as the unconditional reserve price and auction price (whether or not trade occurred) and the reserve price, auction price, and final price estimated only using cases which did end in trade. Period 1 is the auction. Observations ending in period 1 represent cases which ended with the auction or immediately thereafter, and hence

Table 3: Outcome of game by period: Fleet/lease sample

			Full fleet/lease sample				Conditional on trade occurring		
Ending period	Player's turn	# Obs	% of Sample	% Trade	Reserve price	Auction price	Reserve price	Auction price	Final price
1	(Auction)	112,410	82.376%	80.73%	$10,818	$10,559	$10,716	$10,655	$10,655
				(39.44%)	($5,883)	($5,886)	($5,993)	($5,990)	($5,990)
2	S	11,016	8.073%	90.99%	$7,412	$6,163	$7,386	$6,231	$6,231
				(28.64%)	($4,593)	($4,387)	($4,588)	($4,391)	($4,391)
3	B	10,793	7.909%	13.89%	$8,246	$6,464	$8,457	$7,257	$7,710
				(34.58%)	($4,211)	($3,961)	($4,256)	($4,132)	($4,199)
4	S	1,013	0.742%	83.51%	$8,619	$7,053	$8,608	$7,080	$7,404
				(37.12%)	($4,476)	($4,236)	($4,473)	($4,242)	($4,316)
5	B	1,117	0.819%	46.73%	$8,928	$7,338	$9,091	$7,643	$8,280
				(49.92%)	($4,573)	($4,321)	($4,612)	($4,367)	($4,482)
6	S	47	0.034%	87.23%	$9,083	$7,246	$8,780	$6,967	$7,587
				(33.73%)	($4,451)	($4,423)	($3,986)	($3,881)	($4,092)
7	B	56	0.041%	57.14%	$10,876	$8,915	$10,802	$8,680	$9,697
				(49.94%)	($6,129)	($5,628)	($6,893)	($6,223)	($6,613)
8	S	4	0.003%	50.00%	$11,250	$9,650	$6,250	$4,800	$5,475
				(57.74%)	($7,263)	($6,730)	($1,768)	($1,131)	($1,591)
9	B	3	0.002%	33.33%	$11,633	$10,150	$11,400	$10,600	$10,900
				(57.74%)	($5,754)	($5,489)	.	.	.

Notes: For each period (period 1 = auction and immediately following, period 2 = seller's first turn in post-auction bargaining, period 3 = buyer's first turn, etc.), table reports the number of observations ending in that period, percent of total sample ending in that period, and percent of time 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.

cases for which no bargaining actions are recorded.[19]

The remaining periods are labeled with even numbers for seller turns and odd numbers for buyer

[19] When the reserve is not met, the game may still end immediately (and very frequently does–61.26% of the time in the dealers sample and 45.65% of the time in the fleet/lease sample (not shown in Tables 2–3)—when one of the following is true: 1) the seller is present at the auction house and can immediately accept or reject the auction price; 2) the auctioneer rejects the auction price on behalf of the seller, knowing that alternating-offer bargaining is unlikely to occur; or 3) the high bidder walks away from bargaining before the seller is contacted. The third case is observable in the data, and occurs 1.32% of the time in the dealers sample and 0.59% of the time in the fleet/lease sample (not shown in Tables 2–3). The first two cases are indistinguishable in the data, and therefore I assume that any time the auctioneer chooses to reject the auction price on behalf of the seller this choice coincides with what the seller would have done given the chance.

turns. Period 2 is the seller's first turn in the post-auction bargaining game, and observations ending in this period represent cases in which the seller accepted the auction price or quit. Period 3 is the buyer's first turn in bargaining, and is reached only if the seller chose to counter in period 2. Play continues back and forth between the buyer and seller until one party accepts or quits.

Table 2 demonstrates that in 77% of the dealers sample the game ends at the auction, and in these cases the final price when trade happens (which occurs 81% of the time) is naturally the auction price. Consider now the fifth period of the game. Only 1.4% of the full sample reaches the this period, but this still consists of nearly 2,000 observations. In the fifth period, when trade does occur, it occurs at an average final price ($7,643) which is over $600 above the average auction price ($7,001), but still does not reach as high as the average reserve price ($8,519). Overall, Table 2 suggests that observations which ended in later periods had somewhat higher reserve prices than those ending in earlier periods, consistent with Coasian dynamics. Only 3 buyer-seller pairs endured ten periods of the game, all of them coming to agreement in the end, at an average price over $1,000 above the average auction price.

Table 3 displays similar patterns for the fleet/lease sample. 82% of the sample ended the game at the auction. Less than one percent ended in the fifth period, but this still consists of over 1,000 observations. Buyer-seller pairs who traded in this round did so at an average of $600 above the auction price. Other than observations ending in the first period, observations with phone bargaining ending in later periods had higher average reserve prices than those ending in earlier periods, again consistent with Coasian dynamics. Three pairs remained for nine periods of the game, with only one of the three ending in trade.

3 Model of Post-Auction Bargaining with Secret Reserve Price

This section presents a model of the auction-followed-by-bargaining mechanism used at wholesale auto auctions. I first discuss the timing of the mechanism and set up some general assumptions. I discuss each stage of the game, starting from the end with the post-auction bargaining stage. I then present a model of the ascending auction stage, demonstrating that truth-telling is a weakly dominant strategy, and the stage in which the seller chooses a secret reserve price, demonstrating that the seller's strategy is strictly increasing.

The timing of the game at wholesale auto auctions is as follows:

1. Seller sets a secret reserve price.
2. N bidders bid in an ascending auction.
3. If the auction price 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.[20]

[20]At wholesale auto auctions, some large institutional sellers are given the option to elect to eliminate step 4 above, implying that when the auction price does not meet the secret reserve price, the high bidder is not allowed to immediately

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) The ascending auction follows a button auction model with $N \geq 2$ risk-neutral bidders participating. For $i = 1, ..., N$, each buyer i has a private valuation $\tilde{B}_i = W + B_i$, with $W \sim F_W$, $B_i \sim F_B$, and with B_i independent of W.

(A2) The risk-neutral seller has a private valuation $\tilde{S} = W + S$, where $S \sim F_S$, with S independent of W and B_i for all i.

(A3) In bargaining players face a per-offer disutility, $(c_B, c_S) > 0$, as well as discount factor, $\delta \in [0, 1)$, where $1 - \delta$ represents the probability that bargaining will break down exogenously.

I further assume F_B, F_S, and F_W have corresponding atomless densities f_B, f_S, and f_W, with supports $[\underline{B}, \overline{B}]$, $[\underline{S}, \overline{S}]$, and $(-\infty, \infty)$. Let realizations of these random variables be denoted b, s, and w.

The assumption of private values implies that the random variable W is observed by all buyers and the seller. In the estimation framework in Section 4, I allow W to be unobservable to the econometrician but observed by all players, thus allowing for all players' values to be correlated.[21] Conditional on W, buyers and sellers have independent private values (IPV). A motivation for this framework is that buyer, or dealer sellers, have valuations arising primarily from their local demand and inventory needs.[22] Also, seller valuations can depend on the value at which the car was assessed as a trade-in, or, for a bank or leasing company, valuations can arise from the size of the defaulted loan.[23] The button auction model is a natural choice given that jump bidding is rare as it is the auctioneer who calls out bids, and bid increments are small.[24] The assumption of symmetry is not strong in this setting given that buyer identities are unknown to the seller in bargaining and given the assumption of a private values button auction, implying

walk away from bargaining but must wait until the seller responds to the auction price. This situation is referred to as a "binding-if auction." It can be shown that in a binding-if auction, the seller's secret reserve price strategy is only guaranteed to be weakly increasing, rather than strictly as in the non-binding-if case. It can also be shown that bidders will not necessarily drop out of bidding at their valuations but may instead drop out at a price slightly below their valuation to account for the possibility of paying bargaining costs. In the data, there is no way to know if a sale took place in a binding-if setting. I treat all auctions as non-binding-if auctions.

[21] I also introduce characteristics observable to *both* the econometrician and all players and I control for these observables.

[22] Conversations with buyers, sellers, and auction house employees support this assumption: buyers claim to decide upon their willingness to pay before bidding begins, often having a specific retail customer lined up for a particular car. See Lang (2011). Studying Korean auto auctions, Kim and Lee (2008) tested and failed to reject the IPV assumption, while Roberts (2010) found evidence of unobserved auction-level heterogeneity (analogous to W).

[23] These 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.

[24] Bid increments lie between \$25 and \$100.

that bidders' auction strategies will not depend on the identities of other participants.[25] The parameter δ captures the feature that bargaining may end through an auction house employee failing to follow up on a bargaining sequence, occurring in 1–2% of bargaining interactions.[26]

In what follows, I demonstrate two properties which prove useful for estimation: 1) a buyer's auction strategy is to drop out at his value, as in a standard ascending auction, and 2) the seller's secret reserve price strategy, $\rho(S)$, is strictly increasing in her type S. I model the game conditional on a realization of W and thus omit W for notational simplicity and return to it in Section 3.4).

3.1 Bargaining Stage

This section describes a simple model of the dynamic, post-auction bargaining game. The game begins with an offer by the buyer in period $t = 1$. At wholesale auto auctions, this offer is the auction price at the auction. The seller then chooses between accepting (A), quitting (Q)—meaning terminating the negotiations—or making a counteroffer (C). Accepting ends the game, with trade taking place at the accepted price. Quitting also ends the game, with no trade taking place. After a counteroffer by the seller, play returns to the buyer, who then chooses between accepting, quitting, and counter offering. Thus, at t even it is the seller's turn, and at t odd it is the buyer's turn. Below, I refer to period "t" as being the seller's turn and period "$t + 1$" as being the buyer's turn. Where useful for clarification, I also include the superscripts "S" or "B" in notation to denote an action taken by the seller or buyer respectively.

Suppose it is the seller's turn at time t. Let $H^t \equiv \{P_\tau\}_{\tau=1}^{t-1}$ represent the set of offers made from period 1 up through period $t - 1$. 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 time t is given by the following. Conditional on the history of offers being $H^t = h^t \equiv \{p_\tau\}_{\tau=1}^{t-1}$, which includes the buyer's most recent offer $P_{t-1}^B = p_{t-1}^B$, a seller of type $S = s$,

[25] See Coey, Larsen, and Sweeney (2014) for a formal treatment of the implications of asymmetries in ascending auctions.

[26] This number is based on the percent of bargaining sequence records in which trade failed and the sequence of offers was incomplete, ending with an counteroffer. A key reason for modeling discounting (either time discounting or, in this case, a probability of exogenous breakdown) lies in the results of Perry (1986), who showed in a game with two-sided uncertainty that if 1) there is no discounting and 2) bargaining costs take the form of an additive cost common to all buyers and an additive cost common to all sellers then the unique equilibrium is for bargaining to end immediately. Cramton (1991) discusses how allowing for discounting overcomes this feature.

chooses to accept (A), quit (Q), or counter (C), yielding the following payoffs:

$$
\begin{aligned}
&\text{A} : p^B_{t-1} \\
&\text{Q} : s \\
&\text{C} : V^S_t\left(s|h^t\right) \\
&\quad = \max_p \Big\{ p\delta \Pr\left(D^B_{t+1} = A|h^{t+1}\right) + s\left(\delta \Pr\left(D^B_{t+1} = Q|h^{t+1}\right) + 1 - \delta\right) \\
&\quad + \delta \Pr\left(D^B_{t+1} = C|h^{t+1}\right)\left(\delta E_{P^B_{t+1}}\left[\max\left\{P^B_{t+1}, s, V^S_{t+2}\left(s|H^{t+2}\right)\right\}\Big| h^{t+1}\right] + s(1-\delta)\right)\Big\} - c_S
\end{aligned}
$$

where p is the counteroffer chosen by the seller.[27] The per-period bargaining disutility ($c_S > 0$) is assumed to be common across sellers, and the probability of not terminating exogenously ($\delta < 1$) is assumed to be common across sellers as well as buyers. 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. Exogenous breakdown may occur in any period, in which case the seller receives s as a payoff.[28]

Similarly, the buyer's payoff at time $t+1$ is given by the following. Conditional on the history of offers being $H^{t+1} = h^{t+1} \equiv \{p_\tau\}^t_{\tau=1}$, which includes the seller's most recent offer $P^S_t = p^S_t$, a buyer of type $B = b$ chooses to accept (A), quit (Q), or counter (C), yielding the following payoffs:

$$
\begin{aligned}
&\text{A} : b - p^S_t \\
&\text{Q} : 0 \\
&\text{C} : V^B_{t+1}\left(b|h^{t+1}\right) \\
&\quad = \delta\Bigg(\max_p \Big\{ (b-p)\Pr\left(D^S_{t+2} = A|h^{t+2}\right) \\
&\quad + \delta \Pr\left(D^S_{t+2} = C|h^{t+2}\right) E_{P^S_{t+2}}\left[\max\left\{b - p^S_{t+2}, 0, V^B_{t+3}\left(b|H^{t+3}\right)\right\}\Big| h^{t+2}\right]\Big\}\Bigg) - c_B
\end{aligned}
$$

where p is the counteroffer chosen by the buyer and $c_B > 0$ represents the buyer's per-period bargaining disutility, assumed to be common across buyers. The buyer's outside option is normalized to zero.

3.2 Ascending Auction Stage

This section discusses bidders' strategies in the ascending auction stage of the mechanism. Bidder i's strategy is the price, β_i, at which he stops bidding as a function of his type, $B_i = b_i$, which represents his

[27] The realization h^{t+1} includes all of h^t as well as the seller's choice of counteroffer in this period, and the random variable H^{t+2} includes the seller's choice in this period, as well as the buyer's choice in the next period, the random variable P^B_{t+1}.

[28] In reality, the outside option of both the buyer and seller is a complicated object that cannot be estimated in the scope of this data, as a buyer who exits bargaining has the choice to obtain vehicles from a variety of sources, such as other sales at the same auction house, competing auction houses, online markets, or trade-ins, and a seller who exits bargaining may choose to leave the car at the auction house, return the car to her car lot, or sell it through another source.

valuation for the car. Let $R = \rho(S)$ be a random variable representing the secret reserve price of seller of type S who uses reserve price strategy $\rho(\cdot)$. Let

$$\beta = \max_{k \neq i} \beta_k(B_k)$$

Bidder i will be the highest bidder if and only if $\beta_i > \beta$. The expected payoff of bidder i from following bidding strategy $\beta_i(b_i)$ is given by

$$M(b_i, \beta) = \begin{cases} (b_i - \beta)\Pr(\beta > R) & \\ \qquad +\pi^B(\beta, b_i)\Pr\left(\beta < R, \pi^B(\beta, b_i) > 0\right), & \text{if } \beta_i > \beta. \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

Buyer i would decide to enter bargaining if $\pi_B(\beta, b_i) > 0$, where $\pi^B(\beta, b_i)$ represents the buyer's expected payoff from entering bargaining, equivalent to the counteroffer payoff in the previous section but with the auction price as the buyer's counteroffer, rather than the maximizing counteroffer. In this setup, the following property holds:

Proposition (1). *If in the bargaining game the seller never accepts offers below the auction price, truth-telling is weakly dominant for bidders in the auction stage. That is,* $\beta_i(b_i) = b_i$.

All proofs are found in Appendix A.

This result implies that the winning bid will be the second order statistic from the distribution of buyer valuations, as in standard ascending or second price auctions.[29] Intuitively, the assumption that the seller never accepts bargaining offers below the auction price ensures that buyers will not be tempted to bid beyond their valuations in the ascending auction stage in hopes of bargaining to a lower price in the post-auction bargaining stage. This is supported by the data: the bargained price is not lower than the auction price. Moreover, bidders are not tempted to drop out before the bidding reaches their valuations because if the high bidder learns that the auction price did not meet the secret reserve, he can always opt out of bargaining. And because the auction price is the second-highest valuation of the bidders, the seller cannot infer anything about the valuation of the winner other than learning the point at which the buyer distribution is truncated, eliminating any incentive of buyers to shade bids downward.

3.3 Secret Reserve Price Stage

In this section, I discuss the seller's choice of a secret reserve price, chosen before the beginning of the auction to maximize the seller's expected revenue. Applying Proposition 1, the auction price will be $B^{(2)}$,

[29] Huh and Park (2010) found the same result in a theoretical model of second price auctions with complete information (rather than incomplete information) post-auction bargaining: bidders' strategies were unaffected by the presence of bargaining.

the second order statistic of buyer valuations. In choosing her secret reserve price, $\rho(S)$, a seller of type $S = s$ wishes to maximize her ex-ante payoff, given by

$$E_{B^{(2)},B}\left[B^{(2)} * 1\left\{B^{(2)} > \rho(s)\right\} + s * 1\left\{B^{(2)} < \rho(s), \pi^B(B^{(2)}, B) \leq 0\right\}\right. \tag{2}$$

$$\left. + \pi^S\left(B^{(2)}, s\right) * 1\left\{B^{(2)} < \rho(s), \pi^B(B^{(2)}, B) > 0\right\}\right] \tag{3}$$

This term consists of three pieces: 1) the auction high bid, which the seller receives if it exceeds the reserve; 2) the seller's outside option, her type s, which the seller receives if the auction high bid is below the reserve and the buyer opts out of bargaining; and 3) the seller's bargaining payoff, $\pi^S\left(B^{(2)}, s\right) = \max\left\{B^{(2)}, s, V_2^S\left(s|B^{(2)}\right)\right\}$, which the seller receives when the high bid is below the reserve and bargaining occurs. I apply a monotone comparative statics result from Edlin and Shannon (1998), a special case of Topkis's Theorem, to obtain the following:

Proposition (2)**.** *The seller's optimal secret reserve price, $\rho^*(s)$, is strictly increasing in s.*

The intuition behind Proposition 2 is that the secret reserve price is never revealed and hence the seller can use a separating strategy without perfectly signaling her type. To prove this result, I first show that bargaining payoffs are weakly increasing in players' types. The strict monotonicity relies on bargaining being costly to buyers (Assumption A3), such that some buyers will choose to opt out of bargaining when informed they did not meet the secret reserve. Without costly bargaining, Topkis's Theorem can be used to show that $\rho^*(s)$ will be weakly increasing.

3.4 Auction-level Heterogeneity

Sections 3.1–3.3 derived results conditional on a given realization of auction-level heterogeneity. The independence of W, S, and B in the model described above yields the following result:

Lemma (1)**.** *Suppose, when $W = 0$, the secret reserve is r 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 when $W = w$ the secret reserve will be $r + w$, the period t offer will be $p_t + w$, and the period t decision will be d_t.*

Lemma 1 is similar to results used elsewhere in the empirical auctions literature (Haile, Hong, and Shum 2003) but is a generalization specific to this setting of a secret reserve price auction followed by bargaining. The result makes it feasible to apply empirical approaches accounting for auction-level heterogeneity, both observed and unobserved, as described in Sections 4.2–4.3. Lemma 1 is also crucial in identifying the region of the buyer and seller type space in which trade occurs, as described in Section 5.2.

4 Estimating Distributions of Buyer and Seller Valuations

In this section, I exploit the model properties derived above in order to estimate the distribution of buyer and seller valuations. The estimation consists of several steps:

1. Accounting for auction house fees
2. Controlling for observed heterogeneity (auction-level characteristics observable to the econometrician)
3. Controlling for unobserved heterogeneity (auction-level heterogeneity observed by the players but not by the econometrician)
4. Estimating the distribution of buyer valuations through an order statistic inversion
5. Estimating bounds on the distribution of seller valuations using revealed preferences arguments

Let $j = 1...J$ represent observations in the data, where each observation contains a complete set of actions for one play of the game, i.e. a reserve price, auction price, and any bargaining actions. Let r_j^{raw} and $b_j^{(2),raw}$ represent the reserve price and auction price in the raw data, prior to any adjustments for auction house fees or heterogeneity.

4.1 Auction House Fees

If trade occurs for car j, both the buyer and seller pay a fee, h_j, to the auction house, and the fee schedule is approximately a linear function of the transaction price, p_j.[30] Let

$$h(p) = \alpha_0 + \alpha_1 p.$$

Let $h_j^r \equiv h(r_j^{raw})$ and $h_j^{b^{(2)}} \equiv h(b^{(2),raw})$ represent the fee which would be charged if r_j^{raw} or $b^{(2),raw}$ were the final price. Using an estimate of these objects, the auction price for car j, can be adjusted upward to account for the fact that a buyer would be required to pay the auction price as well as the fee, and the reserve price an similarly be adjusted downward. Let

$$r_j^h \equiv r_j^{raw} - h_j^r$$
$$b_j^{(2),h} \equiv b_j^{(2),raw} + h_j^{b^{(2)}}$$

Thus r_j^h and $b_j^{(2),h}$ represent the reserve price and auction price which would have occurred absent auction house fees.[31]

Fees are observed in the data whenever a transaction occurs. Therefore, the parameters α_0 and α_1 can be estimated with a linear regression using the sample of data in which trade occurs. In practice,

[30] The fee paid by the buyer and seller are not necessarily equivalent for each car, but the fee schedules do not differ drastically and approximating them as equivalent simplifies the fee adjustment.

[31] Note that this is simply an approximation and abstracts away from any effect which the presence of auction house fees has on equilibrium prices.

fees can vary by auction house and can change from year to year. Therefore, I perform the above steps by estimating α_0 and α_1 separately for each auction house and year.

4.2 Observed Heterogeneity

To account for auction-level characteristics, x_j, which are observed to the econometrician as well as to the players, I apply Lemma 1. I assume x_j is independent of the unobserved heterogeneity of car j, w_j (unobserved by the econometrician but observed by all players). Let the total auction-level heterogeneity be given by $\Gamma(x_j, w_j) = x_j'\gamma + w_j$. Lemma 1 implies that auction prices and reserve prices can be "homogenized" (Haile, Hong, and Shum 2003) by jointly regressing reserve prices and auction prices on observables as follows:

$$\begin{bmatrix} r_j^h \\ b_j^{(2),h} \end{bmatrix} = \begin{bmatrix} x_j'\gamma + \gamma_r \\ x_j'\gamma \end{bmatrix} + \begin{bmatrix} \tilde{r}_j - \gamma_r \\ \tilde{b}_j^{(2)} \end{bmatrix},$$

where $\tilde{r}_j = r_j + w_j$, $\tilde{b}_j^{(2)} = b_j^{(2)} + w_j$, and the parameter γ_r captures the difference in means between reserve prices and auction prices. In the vector x_j I include fifth-order polynomial terms (all degrees of the polynomial from one through five) in the auction houses' blue-book estimate, the odometer reading, and the run number of the vehicle.[32] x_j also contains the number of previous auction attempts for the car; the number of pictures displayed online; a dummy for whether or not the odometer reading is considering accurate, and the interaction of this dummy with the odometer reading; the interaction of the odometer reading with a car-make dummies; 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 auction house location interacted with date of sale and auction house location interacted with hour of sale; dummies for 32 different vehicle damage categories recorded by the auction house; dummies for each seller who offers cars for sale at least 500 times; and dummies for discrete odometer bins.[33] Finally, x_j also includes six measures of the thickness of the market during a given auction sale.[34] The R^2 from this first-stage regression is 0.96 in the fleet/lease sample and 0.95 in the dealers sample, implying that most of the variation in auction prices and reserve prices is explained by observables.[35] An estimate of $\tilde{r}_j$ is then given by subtracting $x_j'\hat{\gamma}$ from r_j^h, and similarly for $\tilde{b}_j^{(2)}$.

[32]The run number 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.

[33]Odometer bins are as follows: 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.

[34]I compute these market thickness measures as follows: for a given car on a given sale date at a given auction house, I compute the number of remaining vehicles in still in queue to be sold at the same auction house on the same day which lie 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.

[35]In order to improve estimates of γ, these regressions include observations for which the reserve price is recorded but the auction price is missing and vice versa, as well as observations with incorrectly recorded bargaining sequences, as explained

4.3 Unobserved Heterogeneity

To account for heterogeneity w_j in the value of car j which is observed by the players but not by the econometrician, I apply a result due to Kotlarski (1967), which implies that observation of $\tilde{r}_j = r_j + w_j$ and $\tilde{b}_j^{(2)} = b_j^{(2)} + w_j$, along with an assumption that $E[b^{(2)}] = 0$, is sufficient to identify the densities f_W, f_R, and $f_{B^{(2)}}$. This identification result has been applied in estimation elsewhere in the auctions literature by directly computing characteristic functions and applying Fourier inversions to perform a deconvolution. I adopt a simpler approach using semi-nonparametric maximum likelihood (SNP).[36] The likelihood function be given by

$$\mathcal{L} = \prod_j \left[\int f_{B^{(2)}}(\tilde{b}_j^{(2)} - w) f_R(\tilde{r}_j - w) f_W(w) dw \right] \tag{4}$$

I approximate each density using normalized Hermite polynomials. For each random variable Y,

$$f_Y(y) \approx \frac{1}{\sigma} \left(\sum_{k=0}^{K} \theta_k^Y H_k \left(\frac{y - \mu_Y}{\sigma_Y} \right) \right)^2 \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{y-\mu_Y}{\sigma_Y}\right)^2}$$

where K is a smoothing parameter; θ_Y, μ_Y, and σ_Y are parameters; 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 $k > 2$.

I maximize the likelihood in (4) subject to the constraints $\sum_{i=1}^{K} (\theta_i^Y)^2 = 1$ for each random variable Y, which ensures each approximated function is indeed a density function, and also subject to the constraint $E[b^{(2)}] = 0$. For the smoothing parameter I choose $K = 5$. The location and scale parameters $\{\mu_Y, \sigma_Y\}_{Y=W,R,B^{(2)}}$ are not required for identification but improve the performance of the estimator and are standard in SNP estimation with Hermite polynomials. I estimate them in an initial step, maximizing (4) with the vectors θ^Y set to zero for each Y, that is, each density f_Y is approximated by a $N(\mu^Y, \sigma^Y)$. I then plug in the estimated values of $\{\hat{\mu}_Y, \hat{\sigma}_Y\}_{Y=W,R,B^{(2)}}$ into (4) and maximize over $\{\theta^Y\}_{Y=W,R,B^{(2)}}$. I perform the integration in (4) by Gauss-Hermite quadrature (See Appendix B.1).

4.4 Distribution of Buyer Valuations

Identification of the underlying distribution of buyer valuations, F_B, follows from Proposition 2, which implies that the auction price will be the second order statistic of buyer valuations. Given knowledge of the distribution of the second order statistic of buyer valuations, $F_{B^{(2)}}$ (which is identified by by the previous section), and assuming a known distribution for N, the number of bidders, it is known that

$$F_{B^{(2)}}(b^{(2)}) = \sum_n \Pr(N = n) \left[nF_B(b^{(2)})^{n-1} - (n-1)F_B(b^{(2)})^n \right] \tag{5}$$

See, for example, Athey and Haile (2007). I assume N follows a Poisson distribution with mean λ, truncated below at $N = 2$, yielding

$$F_{B^{(2)}}(b^{(2)}) = \frac{e^{-\lambda}}{1 - e^{-\lambda}(1+\lambda)} \left(e^{\lambda F_B(b^{(2)})}(1 + \lambda(1 - F_B(b^{(2)}))) - 1 - \lambda \right)$$

in Section 2.

[36] See Freyberger and Larsen (2014a) for details on SNP estimation in auction settings with unobserved heterogeneity.

See Lemma 5 of Matsuki (2013).[37] Given λ, and given an estimate of $F_{B^{(2)}}$ from Section 4.3, the function $F_B(b^{(2)})$ can be solved for at any point $b^{(2)}$ using standard nonlinear equation approaches.[38] I adopt a bisection method. Based on personal observation of many used-car auctions, I choose $\lambda = 7$ for dealer sales and $\lambda = 10$ for fleet/lease sales.[39]

4.5 Distribution of Seller Valuations

To identify the distribution of seller valuations, I invoke 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 observation-level bounds are not available. Instead, I obtain bounds on the distribution of $\tilde{S} = S+W$ relying on probability statements formed from many observations of sellers' initial responses to the auction price, $\tilde{B}^{(2)} = B^{(2)} + W$. I then translate these bounds to bounds on F_S by applying Proposition 2.

As in Section 3, let $D_2^S \in \{A, Q, D\}$ represent the seller's choice in period 2 of the bargaining game to accept, quit, or counter. Also, for purposes of identification, in cases where the auction price exceeded the reserve price, or in cases where the auction price fell below the reserve price but there was immediate trade, I consider D_2^S to have a value of accept (A); and in cases where there was immediate disagreement, I consider D_2^S to have a value of quit (Q). Therefore, a realization of $D_2^S = d_2^S$ is available for each observation in the data.

Observe that

$$\Pr(D_2^S = A|\tilde{B}^{(2)} = \tilde{b}^{(2)}) \leq \Pr(\tilde{S} \leq \tilde{b}^{(2)})$$
$$\Pr(D_2^S = Q|\tilde{B}^{(2)} = \tilde{b}^{(2)}) \leq \Pr(\tilde{S} \geq \tilde{b}^{(2)})$$

Both observations follow from revealed preference arguments. Intuitively, if a seller accepts an auction price of $\tilde{b}^{(2)}$, it must be the case that the seller values the car less than $\tilde{b}^{(2)}$, and if the seller quits when the auction price is $\tilde{b}^{(2)}$, it must be the case that the seller values keeping the car herself more than $\tilde{b}^{(2)}$.[40] Define $\tilde{L}(v) \equiv \Pr(D_2^S = A|\tilde{B}^{(2)} = v)$ and $\tilde{U}(v) \equiv \Pr(D_2^S \neq Q|\tilde{B}^{(2)} = v)$.

Combining these two observations yields

$$\tilde{L}(v) \leq F_{\tilde{S}}(v) \leq \tilde{U}(v) \tag{6}$$

[37] I modify Lemma 5 of Matsuki (2013) to apply to the case where N is always at least 2.

[38] Differentiating (5) yields the density:

$$f_B(b^{(2)}) = \frac{f_{B^{(2)}}(b^{(2)})(1 - e^{-\lambda}(1+\lambda))}{\lambda^2 e^{\lambda(F_B(b^{(2)})-1)}(1 - F_B(b^{(2)}))}$$

An estimate of $f_{B^{(2)}}$ comes from the sieve MLE approach of Section 4.3.

[39] Many more bidders are physically present at auction houses during auction sales; these numbers correspond approximately to the mean number of bidders showing active interest in the auction. Similar numbers are found in Genesove (1991).

[40] The equivalent of these two statements in the Haile and Tamer (2003) English auction setting, for buyer valuations, is that a buyer never bids more than the buyer's value and never lets a competitor win at a price the buyer would have been willing to pay. Applying similar bounds for buyer choices in the bargaining game is more complicated than for seller choices given that some bargaining games end quickly, before the buyer's turn, and hence the bargaining games for which a buyer's choice is observable suffer from selection. See Freyberger and Larsen (2014b).

In order to obtain bounds on F_S rather than on $F_{\tilde{S}}$, recall that $F_{\tilde{R}}$ and F_R are both identified (the former from observations of $\tilde{r}$ after controlling for observed heterogeneity and auction house fees, and the latter from the identification argument in Section 4.3), and note that

$$\begin{aligned} F_{\tilde{S}}(v) &= \Pr(\tilde{S} \leq v) \\ &= \Pr(\rho(\tilde{S}) \leq \rho(v)) \\ &= F_{\tilde{R}}(\rho(v)) \end{aligned} \tag{7}$$

The first equality follows by Proposition 2, that the function $\rho(\cdot)$ is strictly increasing, and the second by Lemma 1, which implies that the reserve price in an auction where the seller's value is given by the random variable $\tilde{S} = S + W$ and the buyers' values are random variables given by $\tilde{B} = B + W$ will be $\tilde{R} = \rho(S) + W$. Similarly, note that $F_S(v) = F_R(\rho(v))$. Inverting equation (7) implies

$$\rho(v) = F_{\tilde{R}}^{-1}(F_{\tilde{S}}(v)).$$

Therefore,

$$F_S(v) = F_R(F_{\tilde{R}}^{-1}(F_{\tilde{S}}(v))) \tag{8}$$

Applying the operation $F_R(F_{\tilde{R}}^{-1}(\cdot))$ to the upper and lower bounds on $F_{\tilde{S}}$ yields upper and lower bounds on F_S. Let F_S^U and F_S^L represent these upper and lower bounds.

In practice, I perform several steps to improve the estimates of the upper and lower bounds on $F_{\tilde{S}}$ prior to applying the operation $F_R(F_{\tilde{R}}^{-1}(\cdot))$. I first estimate $\tilde{L}(\cdot)$ and $\tilde{U}(\cdot)$ using a Nadayara-Waton kernel regression.[41] I then incorporate information on $\tilde{r}$ to improve estimates of the upper and lower bounds on $F_{\tilde{S}}$ and to ensure that these bounds correspond to distribution functions. Specifically, note that

$$\tilde{R} > \tilde{S} \Rightarrow F_{\tilde{R}}(v) \leq \tilde{L}(v) \leq \tilde{U}(v) \tag{9}$$

Therefore, I enforce that both the upper and lower bounds on $F_{\tilde{S}}$ lie above $F_{\tilde{R}}$. I also impose that the estimates of the lower and upper bounds on $F_{\tilde{S}}$ be weakly increasing through the rearrangement approach of Chernozhukov, Fernandez-Val, and Galichon (2009).[42] Let $\tilde{L}^*$ and $\tilde{U}^*$ represent the bounds on $F_{\tilde{S}}$

[41] The Nadayara-Watson estimator for $\tilde{L}(\cdot)$ is given by

$$\widehat{\tilde{L}}(u) = \frac{\sum_j K_h(u - \tilde{b}_j^{(2)}) 1\{d_{2,j}^S = A\}}{\sum_j K_h(u - \tilde{b}_j^{(2)})}$$

where $1\{d_{2,j}^S = A\}$ is an indicator for whether the seller of car j chose to accept the auction price $b_j^{(2)}$. K_h is a Gaussian kernel with bandwidth h set to the asymptotically optimal bandwidth for kernel density estimation of $\tilde{b}^{(2)}$. I follow the same procedure for $\tilde{U}(\cdot)$

[42] In this setting, rearrangement is performed by simply sorting the lower bounds (estimated on a uniformly spaced grid) and reassigning them to the original grid points, and similarly for the upper bounds. Chernozhukov, Fernandez-Val, and Galichon (2009) demonstrated that, when estimating a monotone function, a rearranged estimate is always an improvement, in terms of estimation error, over an original, nonmonotonic estimate. In practice, the approach is sensitive to extreme outliers, and so for this esimation I do not include the top and bottom 0.1% of observations of $\tilde{b}^{(2)}$, where the kernel regression estimates are noisy. In practice, the rearrangement has only little impact on the estimates, smoothing out small nonmonotonic portions of the bounds.

after enforcing (9) and after enforcing monotonicity. Finally, these estimates are not necessarily onto; that is, there may not exist values of $\tilde{b}^{(2)}$ at which sellers accept or quit with probability very close to 0 or 1. I fill in these missing parts of the upper and lower bounds by assuming a constant gap between quantiles $F_{\tilde{R}}$ and those of the upper or lower bounds on $F_{\tilde{S}}$.[43]

After estimating F_S^U and F_S^L, upper and lower bounds on F_S, I obtain draws from these distributions by inverting the distributions at random uniform draws. I estimate the corresponding densities using kernel density estimation of these draws.[44]

4.6 Distribution Estimates

Results displayed in Figures 1–4.

5 The Pareto Frontier and Real-World Bargaining

This section describes how the Pareto frontier and others counterfactual mechanisms can be solved for once the distributions of seller and buyer valuations are known. I also describe identification and estimation of the direct mechanism corresponding to the real-world, dynamic bargaining mechanism used at auto auctions. I then bring together the estimates from Section 4 to analyze the efficieny of bargaining.

The counterfactuals hold fixed the distribution of buyer and seller types. In reality, changing the mechanism could change the distribution of types.[45] Also, in counterfactuals I only change bargaining—I do not change the auction. Each counterfactual mechanism is a mechanism for bilateral trade between the seller and high bidder after a no-reserve ascending auction has occurred. The auction selects the highest-value bidder, and the lower bound for the buyer's support in the post-auction bargaining game becomes $b^{(2)}$, the high bid at the auction.[46] Finally, the counterfactual comparisons all consider surplus *after* removing auction-level heterogeneity, both observed and unobserved.

[43] To fill in these values of $\tilde{L}^*(\cdot)$ and $\tilde{U}^*(\cdot)$ close to zero and one, I do the following. Let $\{v_m\}_{m=1,\ldots,M}$ represent the grid of points on which $\tilde{L}^*(\cdot)$ is evaluated. Let $\underline{v}^L = \arg\min_m \tilde{L}^*(v_m)$ and $\overline{v}^L = \arg\max_m \tilde{L}^*(v_m)$. Then

$$F_{\tilde{S}}^L(v) \equiv \begin{cases} F_{\tilde{R}}(v + r(\underline{v}^L)) & \text{if } v < \underline{v}^L \\ \tilde{L}^*(v) & \text{if } v \in [\underline{v}^L, \overline{v}^L] \\ F_{\tilde{R}}(v + r(\overline{v}^L)) & \text{if } v > \overline{v}^L \end{cases}$$

F_S^L is then given by $F_S^L(v) \equiv F_R(F_{\tilde{R}}^{-1}(F_{\tilde{S}}^L(v)))$. I follow the same steps for $\tilde{U}^*(\cdot)$.

[44] As above, I use a Gaussian kernel with the asymptotically optimal bandwidth.

[45] For example, the buyer and seller types choosing to attend the auction house could change if the mechanism were more or less favorable for certain types. Also, the distribution of seller types could change, because embedded in the seller valuations is the option to attempt to sell the car the following week, and the payoff from doing so would change with the mechanism.

[46] Note that I do not work with a direct mechanism in which N buyers and one seller simultaneously report types to a mechanism designer, primarily because this mechanism would be starkly different from mechanisms applied in practice and because I wish to focus on the efficiency of bilateral trade in particular. I also do not consider a secret reserve price in these counterfactual mechanisms.

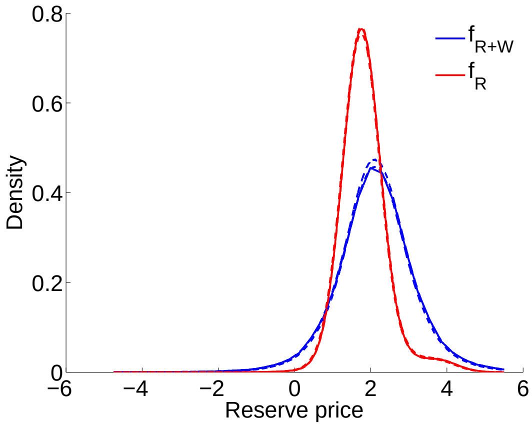

(a) Dealers

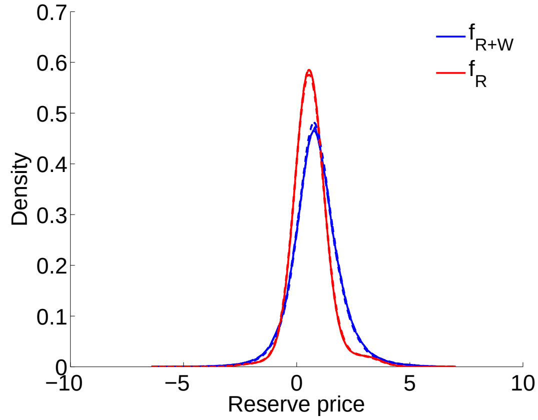

(b) Fleet/lease

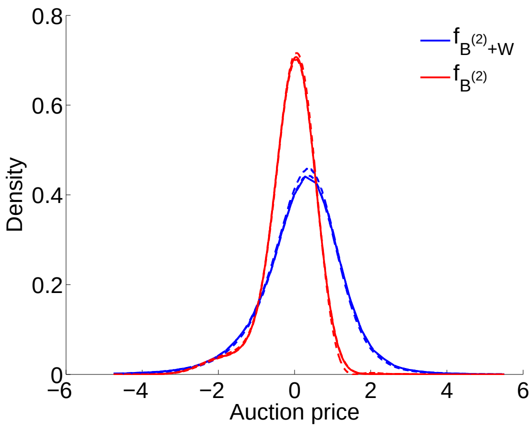

(c) Dealers

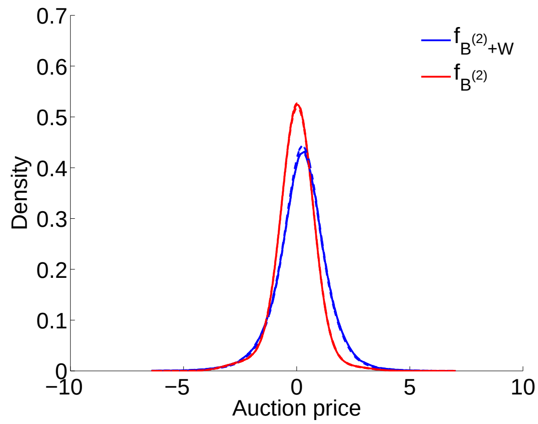

(d) Fleet/lease

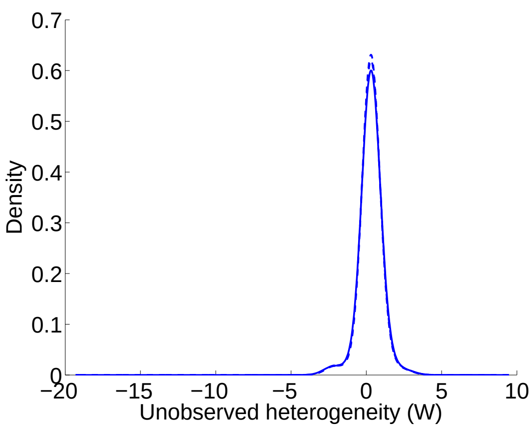

(e) Dealers

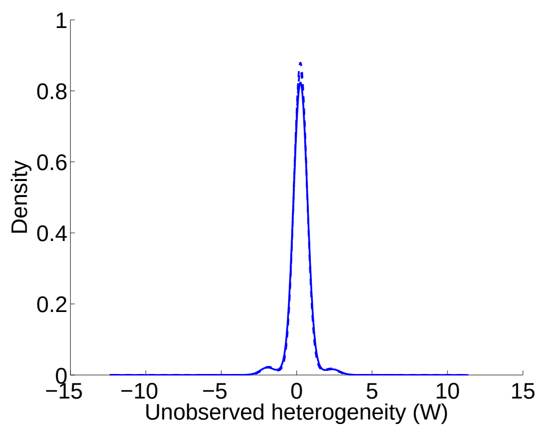

(f) Fleet/lease

Figure 1: Densities of reserve price and auction price with and without unobserved heterogeneity, and density of unobserved heterogeneity.

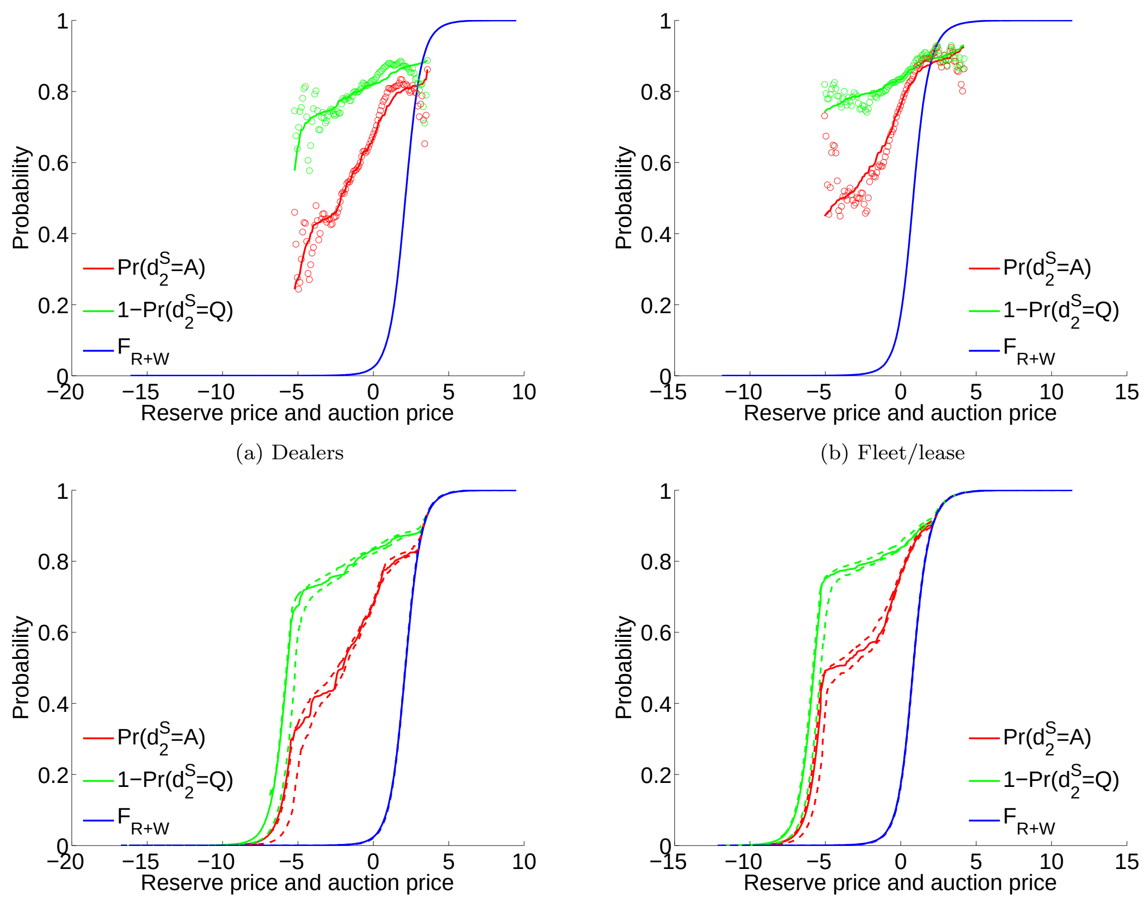

(a) Dealers (b) Fleet/lease

(c) Dealers (d) Fleet/lease

Figure 2: Bounds on distribution of $\tilde{S}$. Panels (a) and (b) display bounds prior to additional estimation steps. Panels (c) and (d) display final bounds on $\tilde{S}$.

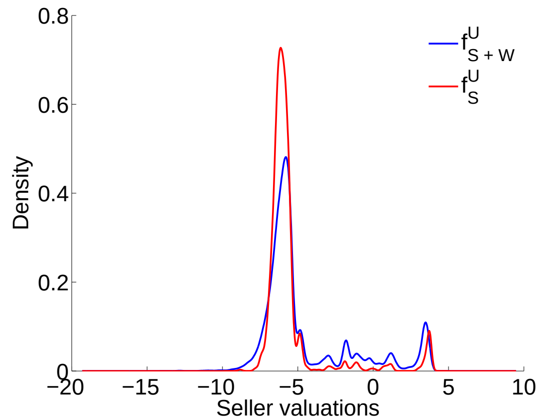

(a) Dealers, upper bound

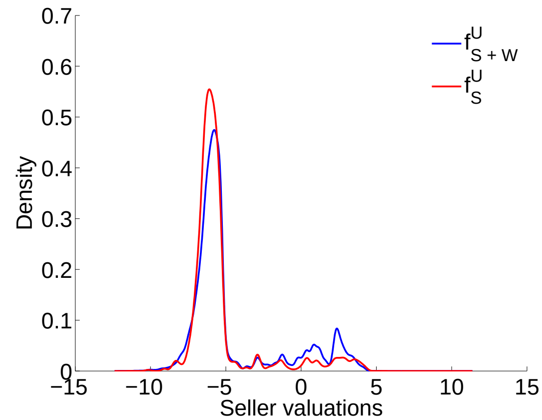

(b) Fleet/lease, upper bound

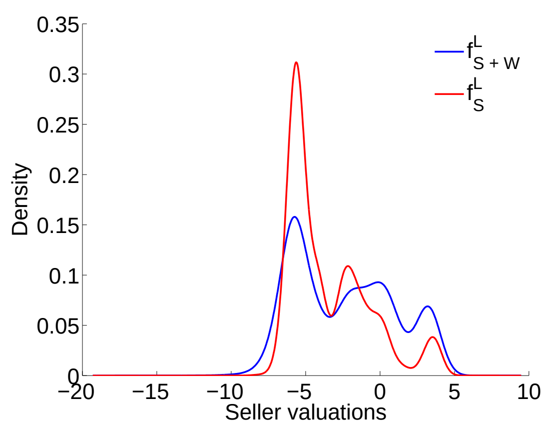

(c) Dealers, lower bound

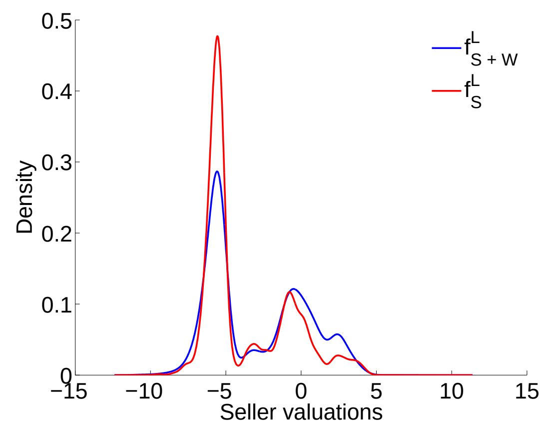

(d) Fleet/lease, lower bound

Figure 3: Densities of seller valuations (upper and lower bound) with and without unobserved heterogeneity.

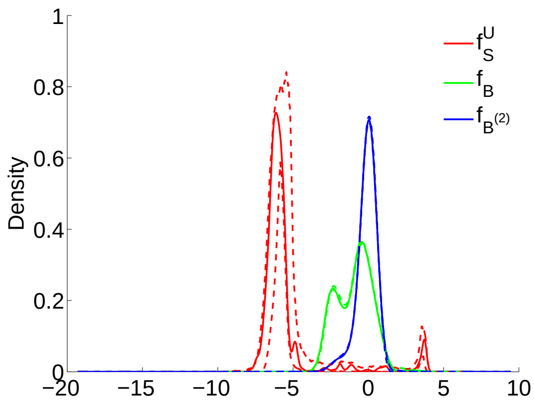

(a) Dealers

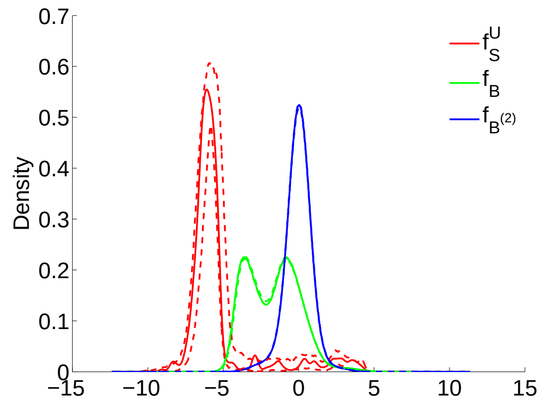

(b) Fleet/lease

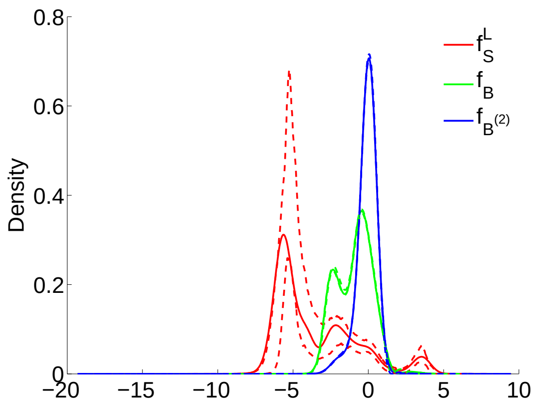

(c) Dealers

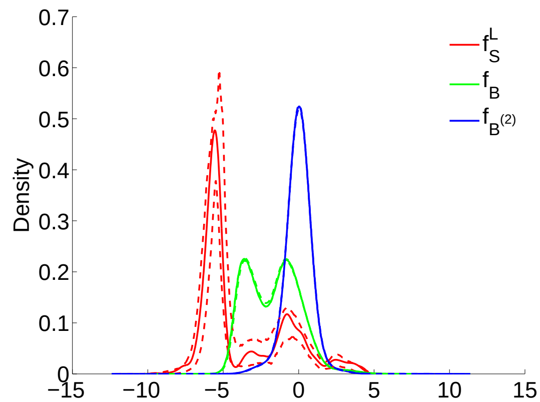

(d) Fleet/lease

Figure 4: Densities of seller valuations (upper or lower bound), buyer valuations, and auction price.

5.1 Solving for the Pareto Frontier and other Direct Bargaining Mechanisms

In this section I discuss how I solve for the Pareto frontier and other direct, efficient mechanisms conditional on a realization of $B^{(2)} = b^{(2)}$. By the Revelation Principle (Myerson 1979), any static, incentive-compatible, individually rational, bilateral trade mechanism can be written as a direct revelation mechanism where player's truthfully report their valuations to a broker and then trade occurs with probability $x(s,b)$, with the buyer paying $p(s,b)$ to the seller.[47] Williams (1987) demonstrated that a bilateral bargaining mechanism can alternatively be summarized by the two objects (x,q), rather than (x,p), where q is the expected utility for the type $\overline{S}$.[48] The ex-ante expected utility of the buyer and seller in a mechanism (x,q) is given by

$$\overline{U}_S(x,q) = q + \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} x(s,b) F_S(s) \frac{f_B(b)}{1-F_B(b^{(2)})} ds db \tag{10}$$

$$\overline{U}_B(x,q) = G(x) - q + \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} x(s,b) \frac{(1-F_B(b))}{1-F_B(b^{(2)})} f_S(s) ds db \tag{11}$$

where

$$G(x) = \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} \left(\phi_B(b) - \phi_S(s)\right) x(s,b) f_S(s) \frac{f_B(b)}{1-F_B(b^{(2)})} ds db \tag{12}$$

and

$$\phi_S(s) \equiv \phi_S(s,1) \qquad \qquad \phi_B(b) \equiv \phi_B(b,1)$$

$$\phi_S(s,\alpha_1) = s + \alpha_1 \frac{F_S(s)}{f_S(s)} \qquad \text{and} \qquad \phi_B(b,\alpha_2) = b - \alpha_2 \frac{1-F_B(b)}{f_B(b)}$$

Williams (1987) demonstrated further that the Pareto frontier, that is, the maximized value of

$$\eta \overline{U}_S + (1-\eta)\overline{U}_B \tag{13}$$

for $\eta \in [0,1]$, can be traced out by the class of mechanisms with trading rules, $x(s,b)$, defined by

$$x^{\alpha_1(\eta),\alpha_2(\eta)}(s,b) = 1\left\{\phi_B(b,\alpha_2(\eta)) \geq \phi_S(s,\alpha_1(\eta))\right\}$$

The parameters $(\alpha_1(\eta), \alpha_2(\eta))$ can be solved for at each η using an approach developed in Williams (1987) and described in Appendix B.3. Intuitively, the approach maximizes (13) subject to $G(x^{\alpha_1(\eta),\alpha_2(\eta)}) \geq 0$, where $G(x)$ is defined in (12). This constraint implies that the worst types—the lowest buyer type and highest seller type—must receive a non-negative surplus in order to be willing to participate in the mechanism.

Existence of these mechanisms and the success of the solution method in Williams (1987) is guaranteed as long as $\phi_S(s)$ and $\phi_B(b)$ are weakly increasing. This assumption is common in the mechanism

[47] Note that the notation here is the reverse of Myerson and Satterthwaite (1983) and Williams (1987), in which p represented the probability of trade and x represented the transfer.

[48] The transfer function, p, is not essential for the results here. I report it in Appendix B.3.

design literature. I impose this condition on the estimated $\phi_S(s)$ and $\phi_B(b)$ before solving the counterfactual mechanisms. To do so, I follow the rearrangement approach of Chernozhukov, Fernandez-Val, and Galichon (2009). The monotonic estimates of $\phi_S(s)$ and $\phi_B(b)$ can then be used to re-solve for the implied densities and distributions as described in Appendix B.2.[49] Overall, the rearrangement has little effect on the densities and distributions other than smoothing out small deviations from monotonicity. Figure 8 in Appendix B.2 displays the original and rearranged densities.

Several mechanisms of interest fit into this framework, such as the following

1. First-best trade (infeasible mechanism where trade occurs whenever buyer values the car more than seller): $\alpha_1 = \alpha_2 = 0$.

2. Second-best trade (the mechanism maximizing the gains from trade): $\eta = 1/2$, $\alpha_1 = \alpha_2 = \alpha^*$, where α^* solves $G(x^{\alpha^*,\alpha^*}) = 0$.

3. Seller-optimal: $\eta = 1$, $\alpha_1 = 0$, $\alpha_2 = 1$.

4. Buyer-optimal: $\eta = 0$, $\alpha_1 = 1$, $\alpha_2 = 0$.

5. Pareto frontier: mechanisms maximizing (13) subject to $G(x^{\alpha_1(\eta),\alpha_2(\eta)}) \geq 0$ for $\eta \in [0, 1]$.

Note that an auction followed by the seller-optimal mechanism is equivalent to a public reserve auction.[50] An additional mechanism with $\alpha_1 = \alpha_2 = 1$ is discussed in Myerson and Satterthwaite (1983) and would maximize the gains to a broker (auction house) with market power. This mechanism is discussed in Appendix C.

I derive an additional direct mechanism which maximizes the probability of trade rather than the gains from trade. This result is a corollary to Theorem 2 of Myerson and Satterthwaite (1983) and the proof follows the same line of reasoning as in Myerson and Satterthwaite (1983).[51] The proof of existence relies on strict monotonicity of $\phi_S(s)$ and $\phi_B(b)$, but in practice I am able to solve for it while only imposing weak monotonicity as described above.

Corollary (1). *Suppose $\phi_S(s)$ and $\phi_B(b)$ are both strictly increasing. Then the direct mechanism maximizing the probability of trade has allocation rule $x^\kappa(s, b) = 1\{\phi_S(s) - (2\kappa)/(1-\kappa) \leq \phi_B(b)\}$, where $\kappa \in [0, 1)$ is the solution to $G(x^\kappa(s, b)) = 0$.*

Once the first-stage auction is taken into account, the ex-ante probability of trade in any of these direct mechanisms is given by

$$\Pr(trade) = \int_{\underline{B}}^{\overline{B}} \int_{B^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} x(s, b; b^{(2)}) f_S(s) \frac{f_B(b)}{1 - F_B(b^{(2)})} f_{B^{(2)}}(b^{(2)}) ds db db^{(2)} \tag{14}$$

[49]This use of Chernozhukov, Fernandez-Val, and Galichon (2009) for imposing monotonicity may be of independent interest in empirical auctions work. For example, it provides a method for imposing monotone bidding function estimates in first price auction settings, an alternative to the approach of Henderson, List, Millimet, Parmeter, and Price (2011).

[50]See Menezes and Ryan (2005).

[51]Note that the expected transfer functions for the mechanism in Corollary 1 are given by (20) and (21) in Appendix B.3

and the expected gains from trade is given by

$$\int_{\underline{B}}^{\overline{B}}\int_{B^{(2)}}^{\overline{B}}\int_{\underline{S}}^{\overline{S}}(b-s)x(s,b;b^{(2)})f_S(s)\frac{f_B(b)}{1-F_B(b^{(2)})}f_{B^{(2)}}(b^{(2)})dsdbdb^{(2)} \tag{15}$$

To perform this integration, I use Gauss-Chebyshev quadrature, as described in Appendix B.1, with 200 nodes in the s and b dimensions, and 25 nodes in the $b^{(2)}$ dimension.[52]

5.2 Estimating the Dynamic Mechanism

This section describes how I solve for surplus in the currently used mechanism by backing out a direct mechanism corresponding to the mechanism used at auto auctions. As with the above mechanisms, this direct mechanism can be characterized by functions x and p determining whether or not trade will occur and at what price. x and p will in general depend on the realization of the high bid at the auction, $b^{(2)}$, as this is the lower bound of the support of buyer types when bargaining takes place.

Let the allocation function for the dynamic mechanism be written

$$x^D(s,b;b^{(2)}) \equiv 1\left\{b \geq g\left(\rho(s), b^{(2)}\right)\right\} \tag{16}$$

where $g(\cdot)$ is an unknown function.[53] Here I again rely on the result of Proposition 2 that ρ is strictly increasing in s and hence the allocation function can be considered to be a function of s or of $r = \rho(s)$; the latter is more convenient computationally.

As (16) is currently written, $g(\cdot)$ is not identified given the available data: I observe whether or not trade occurred, but I do not observe realizations of b, or even of r or $b^{(2)}$ because of unobserved heterogeneity. However, for each observation in the data, I do have an estimate of $\tilde{r} = r + w$ and $\tilde{b}^{(2)} = b^{(2)} + w$, which implies an estimate of $\psi \equiv \tilde{r} - \tilde{b}^{(2)} = r - b^{(2)}$. To exploit this result, note that Lemma 1 implies that the allocation function can be simplified to

$$x^D(r,b;b^{(2)}) = 1\left\{b - b^{(2)} \geq g\left(r - b^{(2)}, b^{(2)} - b^{(2)}\right)\right\} \tag{17}$$

$$= 1\left\{b - b^{(2)} \geq g_0\left(r - b^{(2)}\right)\right\} \tag{18}$$

where $g_0(\psi) = g(\psi, 0)$. Recall that Lemma 1 demonstrates that the probability of trade in a setting where realizations of the reserve price, auction price, and bargaining buyer's type were given by $(r, b^{(2)}, b)$ would be the same as in a setting where each of these objects were reduced by a common amount. The probability of trade at any value of ψ, a realization of the random variable $\Psi = R - B^{(2)}$, can then be written

$$\Pr(trade|\Psi = \psi) = \int_{\underline{B}}^{\overline{B}} \frac{1 - F_B\left(g_0(\psi) + b^{(2)}\right)}{1 - F_B\left(b^{(2)}\right)} \frac{f_R\left(\psi + b^{(2)}\right) f_{B^{(2)}}\left(b^{(2)}\right)}{\int_{\underline{B}}^{\overline{B}} f_R(\psi + v) f_{B^{(2)}}(v) dv} db^{(2)} \tag{19}$$

[52] Increasing the number of nodes in any dimension did not change the results. A greater degree of accuracy in the b and s dimensions than in the $b^{(2)}$ dimension is useful, as each mechanism is solved conditional on $b^{(2)}$.

[53] Although the counterfactual mechanisms discussed in Section 5.1 all had binary allocation functions, this is not necessarily the case for the dynamic mechanism, but without assuming an indicator function the dynamic allocation function would not be identified.

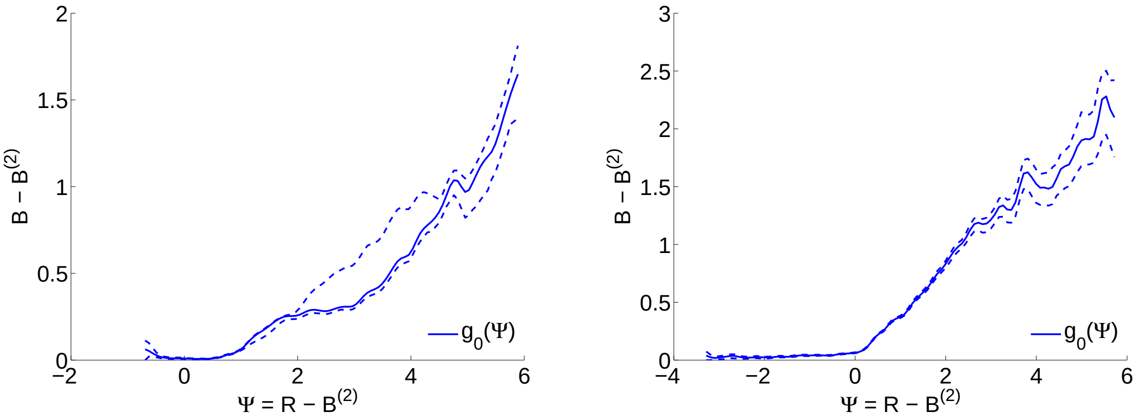

(a) Dealers

(b) Fleet/lease

Figure 5: Estimate of $g_0(R - B^{(2)})$ defining region where trade occurs.

I use (19) to solve for $g_0(\cdot)$ on a grid of ψ using a bisection method. To estimate $\Pr(trade|\Psi = \psi)$, I use a Nadayara-Watson kernel regression.[54] To estimate the right hand side of (19), I use Gauss-Hermite quadrature, plugging in each of the estimated densities. The expected gains from trade under the current mechanism can then be evaluate as in (15), replacing x with the estimate of x^D. Note that this estimate ignores any loss in surplus due to bargaining costs. Section 5.4 discusses bounds on this loss.

5.3 Putting It All Together: How Efficient Is Bargaining?

Results displayed in Figures 5–7 and Table 4–5.

5.4 Bounding surplus lost due to bargaining costs

The parameters δ, c_S, and c_B can be partially identified as follows.

6 Conclusion

This paper examined the efficiency of bargaining from a real-world setting with two-sided incomplete information. I developed a model and strategy for nonparametrically identifying and estimating the distributions of valuations on both sides of the market without relying on a particular structure or equilibrium for the bargaining game. I then mapped these distributions into the static, direct revelation mechanism framework which traces out the efficient frontier derived in Myerson and Satterthwaite (1983) and Williams (1987). I found that the deadweight loss due to incomplete information—the gap between the first-best trade line and the second-best frontier—is small in the wholesale used-car market. I also

[54]I use a Gaussian kernel with bandwidth set to the asymptotically optimal bandwidth for kernel density estimation of Ψ.

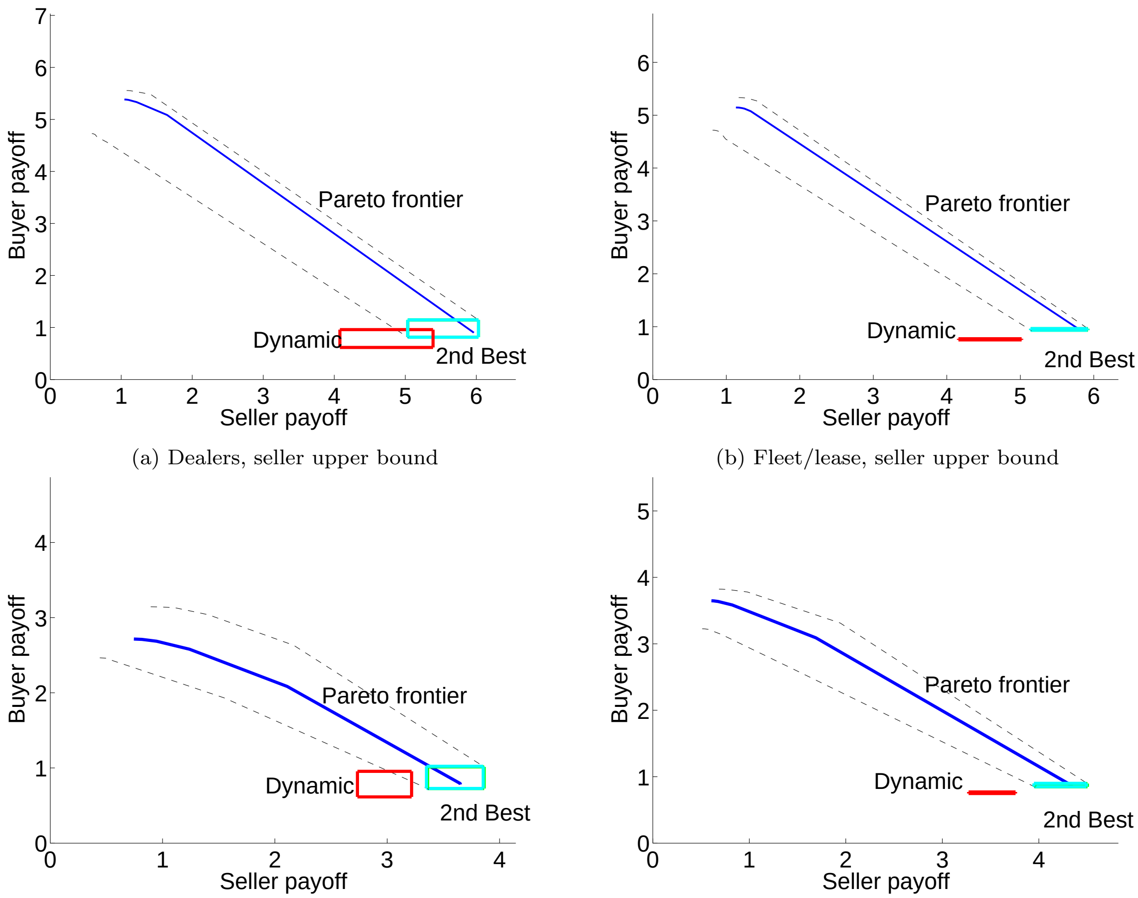

(a) Dealers, seller upper bound

(b) Fleet/lease, seller upper bound

(c) Dealers, seller lower bound

(d) Fleet/lease, seller lower bound

Figure 6: Expected gains from trade in dynamic mechanism and on Pareto frontier.

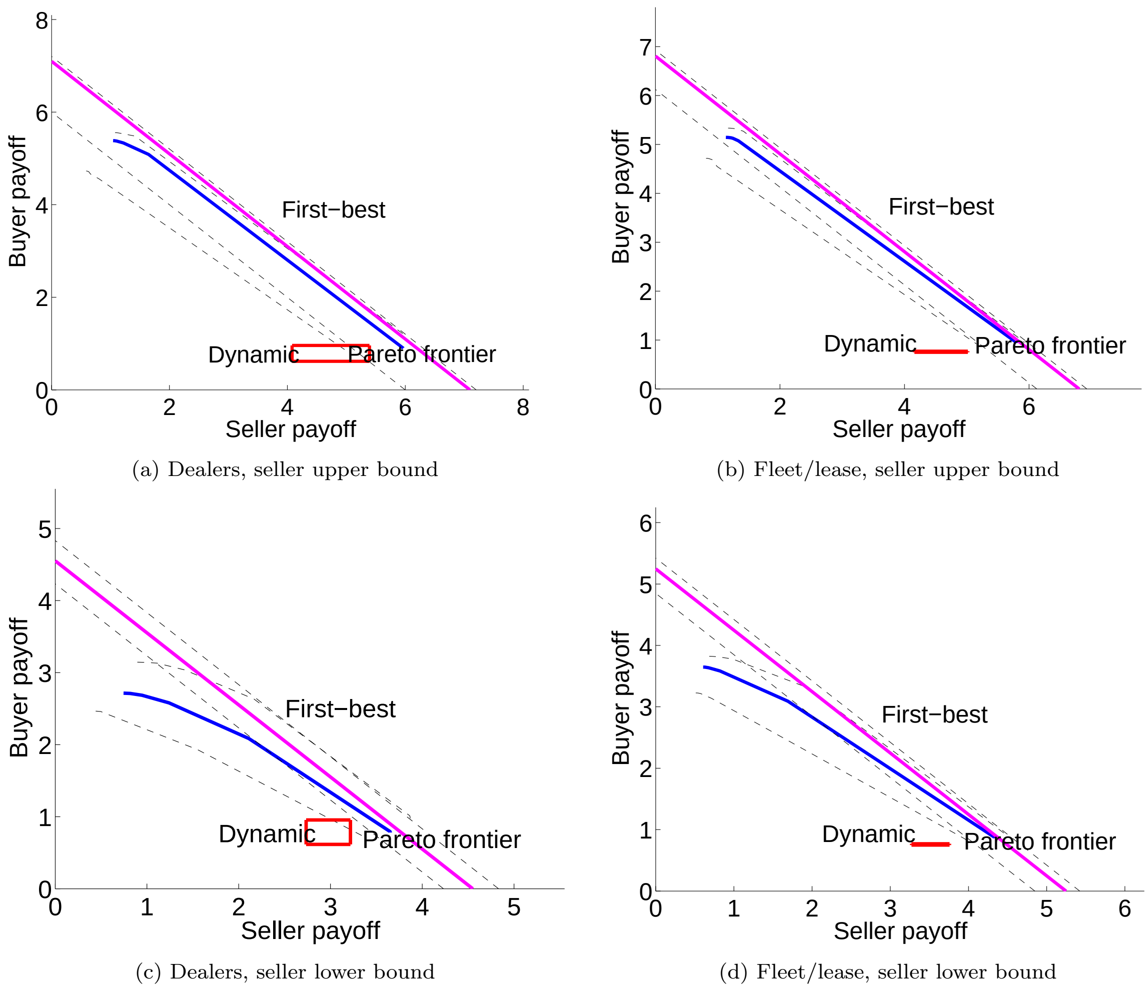

Figure 7: Expected gains from trade in dynamic mechanism, on Pareto frontier, and on first-best efficient frontier.

Table 4: Dealers sample: Expected gains from trade and prob of trade in counterfacual and current mechanisms

A. Counterfactuals under seller distribution upper bound (units=$1,000)					
	First-best	Second-best	Buyer-optimal	Seller-optimal	Dynamic mechanism
Expected gains from trade	7.097	6.866	6.428	6.862	6.527
	(0.262)	(0.236)	(0.260)	(0.238)	(0.339)
Buyer gains		0.903	5.384	0.898	0.714
		(0.067)	(0.199)	(0.066)	(0.068)
Seller gains		5.962	1.044	5.964	5.375
		(0.219)	(0.107)	(0.219)	(0.316)
Probability of trade	0.980	0.974	0.852	0.973	0.711
	(0.008)	(0.009)	(0.020)	(0.008)	--

A. Counterfactuals under seller distribution lower bound (units=$1,000)					
	First-best	Second-best	Buyer-optimal	Seller-optimal	Dynamic mechanism
Expected gains from trade	4.551	4.442	3.460	4.432	4.131
	(0.152)	(0.154)	(0.153)	(0.154)	(0.166)
Buyer gains		0.795	2.715	0.780	0.704
		(0.058)	(0.175)	(0.055)	(0.067)
Seller gains		3.647	0.745	3.652	2.989
		(0.134)	(0.111)	(0.134)	(0.130)
Probability of trade	0.916	0.868	0.541	0.862	0.711
	(0.008)	(0.010)	(0.025)	(0.010)	--

Notes: Dealers sample. Direct, static mechanisms compared to current dynamic mechanism based on estimated gains, payment, and probability of trade. Gains are in $1,000 units. Standard errors are from 200 bootstrap replications.

Table 5: Fleet/lease sample: Expected gains from trade and prob of trade in counterfacual and current mechanisms

A. Counterfactuals under seller distribution upper bound (units=$1,000)					
	First-best	Second-best	Buyer-optimal	Seller-optimal	Dynamic mechanism
Expected gains from trade	7.097	6.866	6.428	6.862	6.527
	(0.262)	(0.236)	(0.260)	(0.238)	(0.339)
Buyer gains		0.903	5.384	0.898	0.714
		(0.067)	(0.199)	(0.066)	(0.068)
Seller gains		5.962	1.044	5.964	5.375
		(0.219)	(0.107)	(0.219)	(0.316)
Probability of trade	0.980	0.974	0.852	0.973	0.711
	(0.008)	(0.009)	(0.020)	(0.008)	--

A. Counterfactuals under seller distribution lower bound (units=$1,000)					
	First-best	Second-best	Buyer-optimal	Seller-optimal	Dynamic mechanism
Expected gains from trade	4.551	4.442	3.460	4.432	4.131
	(0.152)	(0.154)	(0.153)	(0.154)	(0.166)
Buyer gains		0.795	2.715	0.780	0.704
		(0.058)	(0.175)	(0.055)	(0.067)
Seller gains		3.647	0.745	3.652	2.989
		(0.134)	(0.111)	(0.134)	(0.130)
Probability of trade	0.916	0.868	0.541	0.862	0.711
	(0.008)	(0.010)	(0.025)	(0.010)	--

Notes: Dealers sample. Direct, static mechanisms compared to current dynamic mechanism based on estimated gains, payment, and probability of trade. Gains are in $1,000 units. Standard errors are from 200 bootstrap replications.

found that the deadweight loss due to mechanism choice/limited commitment is quite small. This result is consistent with the hypothesis of Wilson (1986) and Ausubel and Deneckere (1993) who suggested that it may be that "[dynamic bargaining mechanisms] survive because they employ trading rules that are efficient for a wide class of environments."

The use of dynamic, post-auction bargaining may seem puzzling at first: why wouldn't used cars be sold with a standard auction format, such as an auction with no reserve price or an auction with a public reserve price? The findings of this paper shed some light on this question. Recall that, in this setting, it is the auction house, rather than the seller, who chooses the mechanism. An auction house is a platform in a two-sided market, required to attract both buyers and sellers, each with private information about his or her valuation for the good. A no-reserve auction could drive some high-value sellers out of the market. And while a public reserve auction is optimal for the seller, alternative mechanisms, including post-auction bargaining, may be preferred for the buyer or for the auction house, and may allow the market to achieve a more efficient allocation. Alternating-offer bargaining in particular is a natural mechanism which is easy for players to understand and for the auction house to implement, and which does not require the same level of commitment as static bargaining mechanisms, which, while more efficient, require players to sometimes walk away from negotiations even when it is discovered ex-post that gains from trade exist.[55]

As highlighted in Appendix C, the nature of wholesale used-car industry as a network of competing platforms likely affects auction houses' choice of mechanism—both the choice of whether or not to allow post-auction bargaining as well as the choice of fee structure. Studying the role of two-sided uncertainty and competition among auction houses in determining auction houses' choice of mechanism is infeasible in the data I use in this paper but would be a valuable goal for future research.

References

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

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

Ambrus, A., E. Chaney, and I. Salitsky (2011): "Pirates of the Mediterranean: An Empirical Investigation of Bargaining with Transaction Costs," Working Paper, Harvard.

Athey, S., and P. Haile (2007): "Nonparametric approaches to auctions," *Handbook of Econometrics*, 6, 3847–3965.

[55]An auction (or some form of bid solicitation) followed by bargaining occurs in numerous real-world settings, often in transactions which account for large shares of GDP. Examples include residential and commercial real estate, corporate takeovers (Burkart 1995), FCC spectrum auctions (Lenard, White, and Riso 2009), and many procurement auctions (Huh and Park 2010). Other examples include privatization auctions in post-Soviet Europe (Dooley, Friedman, and Melese-d'Hospital (1993)) and French timber auctions (Elyakime, Laffont, Loisel, and Vuong 1997).

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

Ausubel, L., and R. Deneckere (1989): "Reputation in bargaining and durable goods monopoly," *Econometrica*, pp. 511–531.

——— (1993): "Efficient sequential bargaining," *Review of Economic Studies*, 60(2), 435.

Bikhchandani, S. (1992): "A bargaining model with incomplete information," *Review of Economic Studies*, 59(1), 187.

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

Bulow, J., and P. Klemperer (1996): "Auctions Versus Negotiations," *American Economic Review*, 86(1), 180–194.

Burkart, M. (1995): "Initial Shareholdings and Overbidding in Takeover Contests," *The Journal of Finance*, 50(5), 1491–1515.

Chatterjee, K., and L. Samuelson (1988): "Bargaining under Two-Sided Incomplete Information: The Unrestricted Offers Case," *Operations Research*, 36(4), 605–618.

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

Chernozhukov, V., I. Fernandez-Val, and A. Galichon (2009): "Improving point and interval estimators of monotone functions by rearrangement," *Biometrika*, 96(3), 559–575.

Cho, I. (1990): "Uncertainty and delay in bargaining," *Review of Economic Studies*, 57(4), 575.

Coey, D., B. Larsen, and K. Sweeney (2014): "Identification in ascending and second price auctions with asymmetric bidders and common values," Working paper.

Compte, O., and P. Jehiel (2002): "On the Role of outside Options in Bargaining with Obstinate Parties," *Econometrica*, 70(4), 1477–1517.

Cramton, P. (1984): "Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty," *Review of Economic Studies*, 51(4), 579.

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

Cramton, P. C. (1991): "Dynamic Bargaining with Transaction Costs," *Management Science*, 37(10), 1221–1233.

Dooley, M., D. Friedman, and S. H. Melese-d'Hospital (1993): "Good and bad explanations for the failure of privatization," Paper for Conference: The Transition of Centrally Planned Economies in Pacific Asia.

Economides, N., and A. Siow (1988): "The Division of Markets Is Limited by the Extent of Liquidity (Spatial Competition with Externalities)," *American Economic Review*, 78(1), 108–121.

Edlin, A., and C. Shannon (1998): "Strict monotonicity in comparative statics," *Journal of Economic Theory*, 81(1), 201–219.

Eklof, M., and A. Lunander (2003): "Open outcry auctions with secret reserve prices: An empirical application to executive auctions of tenant owner's apartments in Sweden," *Journal of Econometrics*, 114(2), 243–260.

Elyakime, B., J.-J. Laffont, P. Loisel, and Q. Vuong (1997): "Auctioning and Bargaining: An Econometric Study of Timber Auctions with Secret Reservation Prices," *Journal of Business & Economic Statistics*, 15(2), 209–220.

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

Freyberger, J., and B. Larsen (2014a): "Identification in Auctions with Unobserved Heterogeneity and Unknown Number of Bidders," Working paper.

——— (2014b): "Partial Identification in Bargaining Games with Two-sided Incomplete Information," Working paper.

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 (1983): "Sequential bargaining with incomplete information," *Review of Economic Studies*, 50(2), 221.

——— (1991): *Game Theory*. Cambridge: MIT Press.

Genesove, D. (1991): "Coconuts, Lemons and Pies: Adverse Selection, Search and Bargaining in the Wholesale Used Car Market," Ph.D. thesis, Princeton University.

Genesove, D. (1993): "Adverse Selection in the Wholesale Used Car Market," *The Journal of Political Economy*, 101(4), 644–665.

——— (1995): "Search at Wholesale Auto Auctions," *The Quarterly Journal of Economics*, 110(1), 23–49.

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

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

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

Haile, P., H. Hong, and M. Shum (2003): "Nonparametric tests for common values at first-price sealed-bid auctions," Working Paper, Yale.

Haile, P. A., and E. Tamer (2003): "Inference with an Incomplete Model of English Auctions," *The Journal of Political Economy*, 111(1), 1–51.

Hammond, R., and T. Morrill (2012): "Strategic Exclusion of the Highest-Valued Bidders in Wholesale Automobile Auctions," Working Paper, North Carolina State University.

Henderson, D., J. List, D. Millimet, C. Parmeter, and M. Price (2011): "Empirical implementation of nonparametric first-price auction models," *Journal of Econometrics*.

Holmström, B., and R. Myerson (1983): "Efficient and durable decision rules with incomplete information," *Econometrica*, pp. 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.

Kennan, J., and R. Wilson (1993): "Bargaining with private information," *Journal of Economic Literature*, 31(1), 45–104.

Kim, K. I., and J. Lee (2008): "A Structural Analysis of Wholesale Used-Car Auctions: Nonparametric Estimation and Testing of Dealers' Valuations 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.

Kythe, P., and M. Schäferkotter (2005): *Handbook of Computational Methods for Integration*, vol. 1. Chapman & Hall/CRC.

Lacetera, N., B. Larsen, D. G. Pope, and J. R. Sydnor (2014): "Bid takers or market makers? The effect of auctioneers on auction outcomes," Working Paper.

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.

Lenard, T. M., L. J. White, and J. L. Riso (2009): "Increasing spetrum for broadband: What are the options?," Technology Policy Institute.

Matsuki, Y. (2013): "Essays on Analysis of Auctions," Ph.D. thesis, University of Wisconsin-Madison.

Menezes, F., and M. Ryan (2005): "Reserve price commitments in auctions," *Economics Letters*, 87(1), 35–39.

Myerson, R. (1979): "Incentive compatibility and the bargaining problem," *Econometrica*, pp. 61–73.

Myerson, R., and M. Satterthwaite (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*, pp. 313–321.

Roberts, J. (2010): "Unobserved Heterogeneity and Reserve Prices in Auctions," Working Paper, Duke.

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*, pp. 1151–1172.

——— (1985b): "Choice of conjectures in a bargaining game with incomplete information," in *Game Theoretic Models of Bargaining*, ed. by A. Roth, pp. 99–114. Cambridge Univ Pr.

Satterthwaite, M., and S. Williams (1989): "Bilateral trade with the sealed bid k-double auction: Existence and efficiency," *Journal of Economic Theory*, 48(1), 107–133.

Sieg, H. (2000): "Estimating a Bargaining Model with Asymmetric Information: Evidence from Medical Malpractice Disputes," *The Journal of Political Economy*, 108(5), 1006–1021.

Silveira, B. (2012): "Bargaining with Asymmetric Information: An Empirical Study of Plea Negotiations," Working Paper.

Sobel, J., and I. Takahashi (1983): "A multistage model of bargaining," *Review of Economic Studies*, 50(3), 411.

Tang, X., and A. Merlo (2010): "Identification of Stochastic Sequential Bargaining Models," Working Paper, University of Pennsylvania.

Townsend, R. (1983): "Theories of intermediated structures," in *Carnegie-Rochester Conference Series on Public Policy*, vol. 18:1, pp. 221–272. Elsevier.

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.

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.

Watson, J. (1998): "Alternating-Offer Bargaining with Two-Sided Incomplete Information," *Review of Economic Studies*, 65(3), 573–594.

Williams, S. (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.

A Proofs and Generalizations

A.1 Proofs

Proof of Proposition 1

Proof. Note that $\pi^B(\beta, b_i)$ is given by

$$\pi^B(\beta, b_i) = \delta\Bigg((b_i - \beta) \Pr\left(D_2^S = A|\beta\right) \\ + \delta \Pr\left(D_2^S = C|\beta\right) E_{P_2^S}\left[\max\left\{b_i - P_2^S, 0, V_3^B\left(b_i|\{\beta, P_2^S\}\right)\right\}\Bigg|\beta\right]\Bigg) - c_B$$

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

Because the high bidder, after learning that the high bid did not meet the secret reserve price, has the option to immediately walk away without entering bargaining, the payoff $M(\beta, b_i)$ cannot be negative. To see that truth-telling is a dominant strategy, suppose first that bidder i drops out at some $\beta_i < b_i$.

1. If $b_i \leq \beta$, then $\beta_i < b_i \leq \beta$, so bidder i is not the high bidder, and would not have been even if he had bid $\beta_i = b_i$.

2. If $b_i > \beta$, then the following is true:

   (a) If $\beta < \beta_i < b_i$, then bidder i is the high bidder and gets an expected payoff of $M(\beta, b_i)$.

   (b) If $\beta_i < \beta < b_i$, then bidder i loses, but i would have been the high bidder if he had bid b_i, and would have again made $M(\beta, b_i)$.

Thus, dropping out at a price less than b_i will never raise bidder i's payoff, and in some cases may decrease it.

Now, suppose that bidder i drops out at some $\beta_i > b_i$

1. If $\beta \leq b_i$, then $\beta \leq b_i < \beta_i$, then bidder i is the high bidder and gets payoff $M(\beta, b_i)$, but would have received this same payoff dropping out at b_i. Also, as noted above, because it is the auctioneer, rather than the bidders, who calls out bids, a player cannot actually outbid himself in an attempt to win the object while avoiding costly bargaining.

2. If $\beta > b_i$, then the following is true:

   (a) If $b_i < \beta_i < \beta$, then bidder i loses, and would not have been the high bidder even if he had bid $\beta_i = b_i$.

   (b) If $b_i < \beta < \beta_i$, then bidder i is the high bidder, but would not choose to enter bargaining because the condition that the seller never accepts offers below the auction high bid rules out the possibility that bidder i could receive a positive payoff by bargaining.

□

Proof of Proposition 2

Proof. In order to prove this result, the following lemma is useful

Lemma (A). *For any finite T and history h_t, $V_t^S\left(s|h^t\right)$ is weakly increasing in s and $V_{t+1}^B\left(b|h^{t+1}\right)$ is weakly increasing in b for all $t \leq T$.*

Proof of Lemma A

The proof proceeds by induction on the number of periods remaining. 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. Let h^{T-1} represent the history at the beginning of period $T-1$ and h^T the history in the final period. The seller's payoff from countering at a price of p is then

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

Let $p^*(s|h^{T-1}) = \arg\max_p W_T^S(s,p|h^{T-1})$. That is, $V_{T-1}^S(s|h^{T-1}) = W_{T-1}^S(s, p^*(s|h^{T-1})|h^{T-1})$. Now let $V_{T-1}(s,s'|h^{T-1})$ represent the payoff to the seller of type s who mimics type $s' < s$. 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, let h^T represent the history in period T when the seller of type s has mimicked type s' in period $T-1$. That is, $h^T = \{h^{T-1}, p^*(s'|h^{T-1})\}$. Observe that

$$V_{T-1}(s,s'|h^{T-1}) = \big(p^*(s'|h^{T-1})\delta \Pr(D_T^B = A|h^T) + s(\delta(1-\Pr(D_T^B = A|h^T)) + 1 - \delta)\big) - c_S,$$

and

$$V_{T-1}(s',s'|h^{T-1}) = \big(p^*(s'|h^{T-1})\delta \Pr(D_T^B = A|h^T) + s'(\delta(1-\Pr(D_T^B = A|h^T)) + 1 - \delta)\big) - c_S$$

Thus,

$$\begin{aligned} V_{T-1}(s,s'|h^{T-1}) - V_{T-1}(s',s'|h^{T-1}) &= (s-s')(\delta(1-\Pr(D_T^B = A|h^T)) + 1 - \delta) \\ &\geq 0 \end{aligned}$$

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 $V^S_{T-(t-1)}(s|h^{T-(t-1)})$ denote the seller's counteroffer payoff with $t-1$ periods remaining, and suppose $V^S_{T-(t-1)}(s|h^{T-(t-1)})$ is weakly increasing s. Note that when there are t periods remaining, $V_{T-t}(s,s|h^{T-t}) \geq V_{T-1}(s,s'|h^{T-t})$ by the same argument as above. It remains to be shown that $V_{T-t}(s,s'|h^{T-t}) \geq V_{T-t}(s',s'|h^{T-t})$. Let

$$\begin{aligned} h^{T-(t-1)} &= \{h^{T-t}, p^*(s'|h^{T-t})\} \\ H^{T-(t-2)} &= \{h^{T-t}, p^*(s'|h^{T-t}), P^B_{T-(t-1)}\} \end{aligned}$$

Note that

$$\begin{aligned} &V_{T-t}(s,s'|h^{T-t}) - V_{T-1}(s',s'|h^{T-t}) \\ &= (s-s')\left(\delta \Pr\left(D^B_{T-(t-1)} = Q|h^{T-(t-1)}\right) + 1 - \delta\right) \\ &+ \ \delta \Pr\left(D^B_{T-(t-1)} = C|h^{T-(t-1)}\right) \\ &\times \left(\delta E_{P^B_{T-(t-1)}}\left[\max\left\{P^B_{T-(t-1)}, s, V^S_{T-(t-1)}\left(s,s'|H^{T-(t-2)}\right)\right\}\right.\right. \\ &- \left.\left.\max\left\{P^B_{T-(t-1)}, s', V^S_{T-(t-1)}\left(s',s'|H^{T-(t-2)}\right)\right\} \middle| h^{T-(t-1)}\right] + (s-s')(1+\delta)\right) \\ &\geq 0 \end{aligned}$$

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, $V^B_{t+1}\left(b|h^{t+1}\right)$, is increasing follows by the same steps.

□

(Continuation of Proof of Proposition 2)

Let $\chi(b)$ be defined by $0 = \pi^B(\chi, b)$, where π^B is defined in the proof of Proposition 1. Intuitively, χ is the high bid at the auction which would make a high bidder of type b indifferent between bargaining and not bargaining. Note that, for $b' > b$, $\pi^B(\chi(b), b') > 0$, because $V_3^B(\cdot)$ is increasing in b by Lemma A. Thus, $\chi(b') > \chi(b)$, and hence χ^{-1}, the inverse, exists and is also strictly increasing. To make notation clear, if $y = \chi(b)$, then this inverse function gives $b = \chi^{-1}(y)$, which defines the lowest buyer type who would enter bargaining when the high bid is y. Also, note that $\chi(b) < b$ because $\pi^B(b, b) < 0$ due to $c_B > 0$.

The seller's payoff can then be re-written as

$$\int_{\rho}^{\overline{B}} b^{(2)} f_{B^{(2)}}(b^{(2)}) db^{(2)} + \int_{\underline{B}}^{\rho} \left[\int_{b^{(2)}}^{\chi^{-1}(b^{(2)})} s f_B(b) db + \int_{\chi^{-1}(b^{(2)})}^{\overline{B}} \pi^S\left(b^{(2)}, s\right) f_B(b) db \right] \frac{f_{B^{(2)}}(b^{(2)})}{1 - F_B(b^{(2)})} db^{(2)}$$

$$= \int_{\rho}^{\overline{B}} b^{(2)} f_{B^{(2)}}(b^{(2)}) db^{(2)} + \int_{\underline{B}}^{\rho} \left[s\left(F_B(\chi^{-1}(b^{(2)})) - F_B(b^{(2)})\right) + \pi^S\left(b^{(2)}, s\right)\left(1 - F_B(\chi^{-1}(b^{(2)}))\right) \right] \frac{f_{B^{(2)}}(b^{(2)})}{1 - F_B(b^{(2)})} db^{(2)}$$

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

$$\frac{\partial}{\partial \rho} = -\rho + s \frac{F_B(\chi^{-1}(\rho)) - F_B(\rho)}{1 - F_B(\rho)} + \pi^S(\rho, s) \frac{1 - F_B(\chi^{-1}(\rho))}{1 - F_B(\rho)}$$

Lemma A implies that $\pi^S\left(b^{(2)}, s\right)$ is weakly increasing in s, and thus $\frac{\partial}{\partial \rho}$ will be strictly increasing in s because $F_B(\chi^{-1}(\rho)) > F_b(\rho)$ given that $\chi^{-1}(\cdot)$ is strictly increasing and $f_B(\cdot)$ is atomless. Given that $\frac{\partial}{\partial \rho}$ is strictly increasing in s, the Edlin and Shannon (1998) Theorem implies that, as long as the optimal $\rho^*(s)$ lies on the interior of the support of ρ, $\rho^*(s)$ will be strictly increasing in s. The support of ρ is the real line, thus completing the proof. Note that without costly bargaining a weak monotonicity can still be obtained following Topkis Theorem. □

Proof of Lemma 1

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

That the auction high bid will be additively separable in W is obvious, given that the bidding function is the identity function. To demonstrate that bargaining offers are also additively separable, the proof proceeds by induction on the number of periods remaining. 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.

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

$$\begin{aligned}
\tilde{p}^* &= \arg\max_{\tilde{p}} \{\delta\tilde{p}(1-F_{\tilde{B}}(\tilde{p})) + \tilde{s}\,(\delta F_{\tilde{B}}(\tilde{p})+1-\delta) - c_S\} \\
&= w + \arg\max_{p} \{\delta p(1-F_B(p)) + s\,(\delta F_B(p)+1-\delta) - c_S + w(1-F_B(p)) + w\,(\delta F_B(p)+1-\delta)\} \\
&= w + \arg\max_{p} \{[\delta p(1-F_B(p)) + s\,(\delta F_B(p)+1-\delta) - c_S] + w\}
\end{aligned}$$

Therefore, the penultimate bargaining offer in the heterogeneous setting is 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. 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, declining, or countering in period $T-t$ can be written as follows:

$$\begin{aligned}
&\text{A} : \tilde{p}^B_{T-(t+1)} = w + p^B_{T-(t+1)} \\
&\text{D} : \tilde{s} = w + s \\
&\text{C} : \tilde{V}^S_{T-t}\left(\tilde{s}|\tilde{h}^{T-t}\right) \\
&\quad = \max_{\tilde{p}} \tilde{p}\delta \Pr\left(D^B_{T-(t-1)} = A|\tilde{h}^{T-(t-1)}\right) + \tilde{s}\left(\delta\Pr\left(D^B_{T-(t-1)} = Q|\tilde{h}^{T-(t-1)}\right) + 1-\delta\right) \\
&\quad + \delta\Pr\left(D^B_{T-(t-1)} = C|\tilde{h}^{T-(t-1)}\right) \\
&\quad \times \left(\delta E_{\tilde{P}^B_{T-(t-1)}}\left[\max\left\{\tilde{P}^B_{T-(t-1)}, \tilde{s}, \tilde{V}^S_{T-(t-2)}\left(\tilde{s}|\tilde{H}^{T-(t-2)}\right)\right\}\middle| h^{T-(t-1)}\right] + \tilde{s}(1-\delta)\right) - c_S \\
&\quad = w + \max_{p} p\delta\Pr\left(D^B_{T-(t-1)} = A|h^{T-(t-1)}\right) + s\left(\delta\Pr\left(D^B_{T-(t-1)} = Q|h^{T-(t-1)}\right) + 1-\delta\right) \\
&\quad + \delta\Pr\left(D^B_{T-(t-1)} = C|h^{T-(t-1)}\right) \\
&\quad \times \left(\delta E_{P^B_{T-(t-1)}}\left[\max\left\{P^B_{T-(t-1)}, s, V^S_{T-(t-2)}\left(s|H^{T-(t-2)}\right)\right\}\middle| h^{T-(t-1)}\right] + s(1-\delta)\right) - c_S
\end{aligned}$$

The last line follows by removing w from each expression and from the following claim: the probability of the buyer accepting, declining, or countering in period $T-(t-1)$ will be the same in the heterogeneous good model as in the homogeneous good model. To prove this claim note that the buyer's payoffs for each action are given by:

$\text{A} : \tilde{b} - \tilde{p}^S_{T-t} = b - p^S_{T-t}$

$\text{D} : 0$

$$
\begin{aligned}
\text{C} &: \tilde{V}^B_{T-(t-1)}\left(\tilde{b}|\tilde{h}^{T-(t-1)}\right)\\
&= \delta\Bigg(\max_{\tilde{p}}(\tilde{b} - \tilde{p}) \Pr\left(D^S_{T-(t-2)} = A|\tilde{h}^{T-(t-2)}\right)\\
&+ \delta \Pr\left(D^S_{T-(t-2)} = C|\tilde{h}^{T-(t-2)}\right) E_{\tilde{P}^S_{T-(t-2)}}\left[\max\left\{\tilde{b} - \tilde{P}^S_{T-(t-2)}, 0, \tilde{V}^B_{T-(t-3)}\left(\tilde{b}|\tilde{H}^{T-(t-3)}\right)\right\}\middle|\tilde{h}^{T-(t-2)}\right]\Bigg) - c_B\\
&= V^B_{T-(t-1)}\left(b|h^{T-(t-1)}\right)
\end{aligned}
$$

Finally, consider the seller's secret reserve price in the auction with auction-level heterogeneity w. Let $\tilde{\chi}$ satisfy $0 = \pi^B(\tilde{\chi}, \tilde{b})$. Note that $\pi^B(\tilde{\chi}, \tilde{b}) = \pi^B(\chi, b)$ by the above arguments for the buyer's bargaining payoff. The first order condition for the seller's secret reserve, $\tilde{r} = \rho(\tilde{s})$, from the proof of Proposition 2, will be given by

$$
\begin{aligned}
\frac{\partial}{\partial \tilde{r}} &= -\tilde{r} + \tilde{s}\frac{F_{\tilde{B}}(\tilde{\chi}^{-1}(\tilde{r})) - F_{\tilde{B}}(\tilde{r})}{1 - F_{\tilde{B}}(\tilde{r})} + \pi^S\left(\tilde{r}, \tilde{r}\right)\frac{1 - F_{\tilde{B}}(\tilde{\chi}^{-1}(\tilde{r}))}{1 - F_{\tilde{B}}(\tilde{r})}\\
&= -\tilde{r} + w + s\frac{F_B(\chi^{-1}(\tilde{r} - w)) - F_B(\tilde{r} - w)}{1 - F_B(\tilde{r} - w)} + \pi^S\left(\tilde{r} - w, s\right)\frac{1 - F_B(\chi^{-1}(\tilde{r} - w))}{1 - F_B(\tilde{r} - w)}
\end{aligned}
$$

Therefore, the optimal secret reserve price in the heterogeneous setting will be w above the optimal reserve in the homogenous setting, completing the proof.

□

Proof of Corollary 1

Proof. This proof follows similar steps to those in the proof of Theorem 2 of Myerson and Satterthwaite (1983) and relies on results from Theorem 1 of Williams (1987). The problem is to find an allocation rule $x : [b^{(2)}, \overline{B}] \times [\underline{S}, \overline{S}] \to [0, 1]$ to maximize

$$\int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} x(s, b) f_S(s) f_B(b) ds db$$

subject to the players' participation constraint, which is

$$0 \leq \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} \left(\phi_B(b, 1) - \phi_S(s, 1)\right) x(s, b) f_S(s) f_B(b) ds db$$

See Myerson and Satterthwaite (1983) for more details. Letting λ denote the Lagrange multiplier, the unconstrained problem is to maximize

$$\int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} \left(1 + \lambda \left(\phi_B(b, 1) - \phi_S(s, 1)\right)\right) x(s, b) f_S(s) f_B(b) ds db$$

For any $\lambda \geq 0$, the Lagrangian is maximized when $x(s,b) = 1$ if and only $(1 + \lambda(\phi_B(b,1) - \phi_S(s,1))) \geq 0$. To achieve this result, let

$$\frac{1}{\lambda} = \frac{2\kappa}{1-\kappa}$$

$\kappa \in [0,1)$ may then be solved for to equate the participation constraint to zero. That is, let

$$\tilde{G}(\kappa) = \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\overline{S}} (\phi_B(b,1) - \phi_S(s,1))\, x^{\kappa}(s,b) f_S(s) f_B(b) ds db$$

where

$$x^{\kappa}(s,b) = 1\left\{\phi_S(s,1) - \frac{2\kappa}{1-\kappa} \leq \phi_B(b,1)\right\}$$

Observe that $x^{\kappa}(s,b)$ is decreasing in κ. Therefore, for some $\alpha < \kappa$, $\tilde{G}(\alpha)$ will differ from $\tilde{G}(\kappa)$ only because $0 = x^{\alpha}(s,b) < x^{\kappa}(s,b) = 1$ for some (s,b) where $\phi_B(b,1) < \phi_S(s,1) - \frac{2\alpha}{1-\alpha}$, implying that at that same (s,b), it must be the case that $\phi_B(b,1) < \phi_S(s,1)$. Thus, as κ increases, $x^{\kappa}(s,b)$ yields trade at regions of the type space at which $(\phi_B(b,1) - \phi_S(s,1))$ is negative. Therefore, $\tilde{G}(\kappa)$ is decreasing in κ.

To prove the $\tilde{G}(\kappa)$ is continuous, note that if $\phi_S(s,1)$ and $\phi_B(b,1)$ are both strictly increasing, then given any b and κ, the equation $\phi_B(b,1) = \phi_S(s,1) - \frac{2\alpha}{1-\alpha}$ has at most one solution in s, so $\tilde{G}(\kappa)$ can be written as

$$\tilde{G}(\kappa) = \int_{b^{(2)}}^{\overline{B}} \int_{\underline{S}}^{\tilde{g}(b,\kappa)} (\phi_B(b,1) - \phi_S(s,1))\, f_S(s) f_B(b) ds db$$

where $\tilde{h}(b,\kappa)$ is continuous in b and κ, so $\tilde{G}(\kappa)$ is continuous. Note also that $\tilde{G}(0) >= 0$, and $\lim_{\kappa \to 1} \tilde{G}(\kappa) = -\infty$. Therefore, there exists a unique $\kappa \in [0,1)$ such that $\tilde{G}(\kappa) = 0$. By Theorem 1 of Williams (1987), the transfer function of this mechanism will by given by (20) and (21). □

A.2 Generalization to Demand-learning Case

Let ...

B Additional Computational Details

B.1 Gaussian Quadrature

The counterfactual analysis in this paper requires the evaluation of a significant number of integrals, such as (14). In order to achieve accuracy and limit the computational burden, I employ Gauss-Chebyshev integration, as advocated by Judd (1998), with a large number of nodes. Specifically, let $z_k, k = 1, ..., K$ be the Chebyshev nodes, given by $z_k = \cos(\pi(2k-1)/(2K))$. Let $g(v)$ be the function to be integrated. Then

$$\int_{\underline{v}}^{\overline{v}} g(v) dv \approx \frac{\pi(\overline{v} - \underline{v})}{2K} \sum_{k=1}^{K} g(x_k) w_k$$

where $x_k = (1/2)(z_k + 1)(\overline{v} - \underline{v}) + \underline{v}$ and $w_k = (1 - z_k^2)^{1/2}$. For integration in multiple dimensions, I use a tensor product:

$$\int_{\underline{u}}^{\overline{u}} \int_{\underline{v}}^{\overline{v}} g(v,u)dvdu \approx \frac{\pi^2(\overline{v} - \underline{v})(\overline{u} - \underline{u})}{(2K)^2} \sum_{j=1}^{K} \sum_{k=1}^{K} g(x_k, x_j) w_k w_j$$

See Kythe and Schäferkotter (2005) or Judd (1998) for additional details. In the estimation of integrals in this paper, I use 200 nodes when integrating in the dimension of the seller's or high bidder's type. Accuracy in these two dimensions is essential, as this is the level at which I solve for counterfactual mechanisms (conditional on the high bid). The integration over the high bid, on the other hand, is not involved in solving for mechanisms, so I use 25 nodes in this dimensions. Increasing the number of nodes beyond 25 does not change results.

In evaluating the SNP likelihood function in Section 4.3, I employ Gauss-Hermite quadrature as follows. Let ...

B.2 Imposing Monotonicity and Solving for Implied Density/Distribution

I impose that $\phi_S(s)$ and $\phi_B(b)$ be weakly increasing following the rearrangement approach of Chernozhukov, Fernandez-Val, and Galichon (2009). In practice, this operation can be performed as follows. Let a grid of values on $[\underline{S}, \overline{S}]$ be given by $z^S = [z_1^S, ..., z_K^S]$ and on $[\underline{B}, \overline{B}$ be given by $z^B = [z_1^B, ..., z_K^B]'$. Let $\hat{\phi}_S(z^S)$ and $\hat{\phi}_B(z^B)$ be the estimates of ϕ_S and ϕ_B obtained by plugging in the estimated distributions and densities from Sections 4.4-4.5 evaluated at the elements of z^B and z^S. Rearrangement is performed by simply sorting the vector $\hat{\phi}_S(z^S)$ and reassigning the sorted values to the original z^S vector, and similarly for $\hat{\phi}_B(z^B)$. Let $\hat{\phi}_S^*(z^S)$ and $\hat{\phi}_B^*(z^B)$ denote the rearranged estimates.

The implied densities and distributions corresponding to the rearranged estimates can then solved for by noting that $d \ln F_S(s)/ds = f_S(s)/F_S(s)$, which implies

$$\int_{\underline{S}}^{s} \frac{1}{\phi_S(u) - u} du = \ln F_S(s) - \ln F_S(\underline{S}) \qquad \Rightarrow F_S(s) = e^{\left(\int_{\underline{S}}^{s} \frac{1}{\phi_S(u)-u} du + \ln F_S(\underline{S})\right)}$$

and similarly for F_B. Thus,

$$\hat{F}_S^*(z_k^S) = e^{\left(\int_{\underline{S}}^{z_k^S} \frac{1}{\hat{\phi}_S^*(u)-u} du + \ln \hat{F}_S(z_1^S)\right)} \qquad \text{and } \hat{f}_S^*(z_k^S) = \frac{\hat{F}_S^*(z_k^S)}{\hat{\phi}_S^*(z_k^S) - z_k^S}$$

$$\hat{F}_B^*(z_k^B) = 1 - e^{\left(\int_{\underline{B}}^{z_k^B} \frac{1}{\hat{\phi}_B^*(u)-u} du + \ln(1-\hat{F}_B(z_1^S))\right)} \qquad \text{and } \hat{f}_B^*(z_k^B) = \frac{1 - \hat{F}_B^*(z_k^B)}{z_k^B - \hat{\phi}_B^*(z_k^B)}$$

Figure 8 displays the results of this monotization on the estimated densities.

B.3 Solving for the Pareto Frontier and the Transfer Function

The Pareto frontier which can be achieved by static, incentive compatible, individually rational, bilateral trade mechanisms can be solved for using Theorem 3 of Williams (1987), which states the following. Recall that each mechanism is summarized by two objects, (x, q), defined in Section 5.

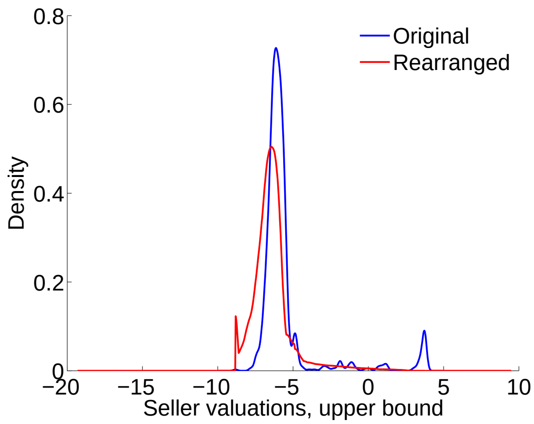

(a) Dealers, seller upper bound

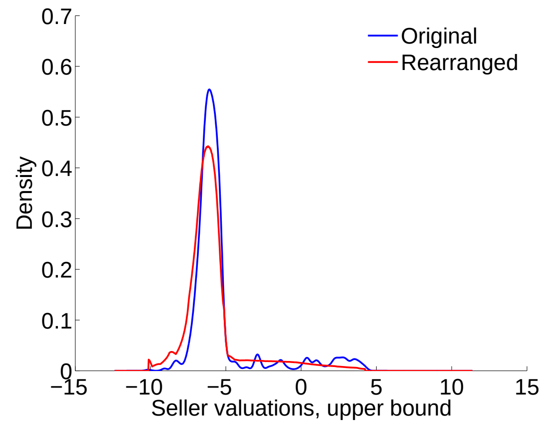

(b) Fleet/lease, seller upper bound

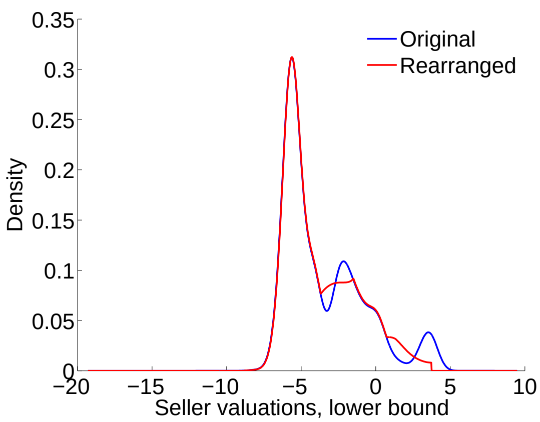

(c) Dealers, seller lower bound

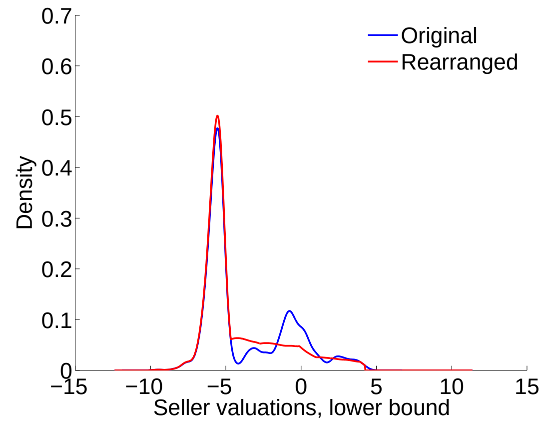

(d) Fleet/lease, seller lower bound

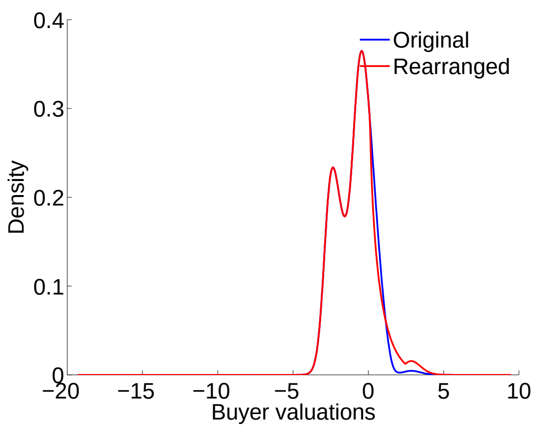

(e) Dealers, buyer

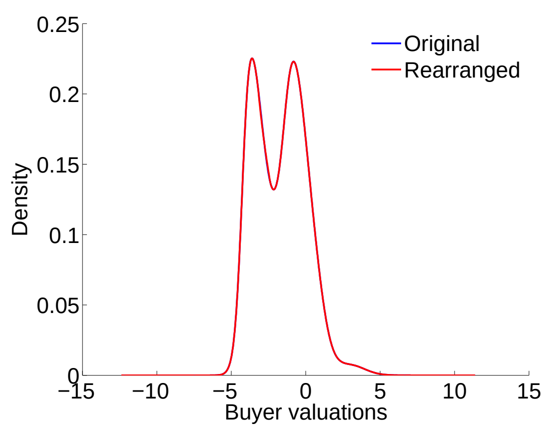

(f) Fleet/lease, buyer

Figure 8: Densities of seller valuations (upper and lower bound on seller valuations) and buyer valuations before and after imposing monotonicity through rearrangement.

Theorem (from Williams 1987)**.**

Suppose $\phi_s(s)$ *and* $\phi_b(b)$ *are weakly increasing. Then*

1. *For* $0 \leq \eta < 1/2$, *if* $G(x^{\bar{\alpha}_1,0}) \geq 0$ *for* $\bar{\alpha}_1 = 1 - \eta/(1-\eta)$, *then* $(x^{\bar{\alpha}_1,0}, 0)$ *is the unique solution maximizing (13) for this* η*; if* $G(x^{\bar{\alpha}_1,0}) < 0$, *then there exists a unique* (α_1^*, α_2^*) *that satisfies the equations* $G(x^{\alpha_1,\alpha_2}) = 0$ *and* $(\alpha_2 - 1) = (\alpha_1 - 1)(1-\eta)/\eta$, *and* $(x^{\alpha_1^*,\alpha_2^*}, 0)$ *is the unique solution maximizing (13) for this* η.

2. *For* $1/2 < \eta \leq 1$, *if* $G(x^{0,\bar{\alpha}_2}) \geq 0$ *for* $\bar{\alpha}_2 = 1 + (\eta - 1)/\eta)$, *then* $(x^{0,\bar{\alpha}_2}, G(x^{0,\bar{\alpha}_2}))$ *is the unique solution maximizing (13) for this* η*; if* $G(x^{0,\bar{\alpha}_2}) < 0$, *then there exists a unique* (α_1^*, α_2^*) *that satisfies the equations* $G(x^{\alpha_1,\alpha_2}) = 0$ *and* $(\alpha_2 - 1) = (\alpha_1 - 1)(1-\eta)/\eta$, *and* $(x^{\alpha_1^*,\alpha_2^*}, 0)$ *is the unique solution maximizing (13) for this* η.

For these direct mechanisms defining the Pareto frontier, Theorem 1 of Williams (1987) implies that, given (x, q), the expected transfer for a seller of type s or for a buyer of type b, which I denote $p_S(s)$ and $p_B(b)$, respectively, are given by

$$p_S(s) = q + s \int_{b^{(2)}}^{\overline{B}} x(s,b) \frac{f_B(b)}{1 - F_B(b^{(2)})} db + \int_s^{\overline{S}} \int_{b^{(2)}}^{\overline{B}} x(u,b) \frac{f_B(b)}{1 - F_B(b^{(2)})} f_S(u) db du \tag{20}$$

$$p_B(b) = G(x) - q + b \int_{\underline{S}}^{\overline{S}} x(s,b) f_S(s) ds + \int_{b^{(2)}}^{b} \int_{\underline{S}}^{\overline{S}} x(s,u) \frac{f_B(u)}{1 - F_B(b^{(2)})} f_S(s) ds du \tag{21}$$

C Auction House Revenues and the Broker Optimal Mechanism

Myerson and Satterthwaite (1983) demonstrated that the mechanism which would maximize revenue for a broker with market power is given by allocation function $x^{1,1}$, with transfers given by $p^B(s,b)$, the amount paid by the buyer to the auction house, and $p^S(s,b)$, the amount which the auction house then passes on to the seller. The difference constitutes auction house revenue. These transfer functions can be defined in many ways. One such way is given by Myerson and Satterthwaite (1983) as

$$p^B(s,b) = x^{1,1}(s,b) * \min\{u | u \geq \underline{B}, \phi_B(u) \geq s\}$$
$$p^S(s,b) = x^{1,1}(s,b) * \max\{v | v \leq \overline{S}, \phi_S(v) \leq b\}$$

Revenue is given by $G(x^{1,1})$, where $G(\cdot)$ is defined in (12). This expression is the participation constraint which must be satisfied in any individually-rational, incentive-compatible mechanism. In the mechanisms which maximize the gains from trade or the probability of trade, this expression is equal to zero. In the mechanism maximizing the auction house revenue, however, the auction house wishes to leave some slack in the participation constraint in order to extract surplus from participants.

Table 6: Performance of broker-revenue maximizing mechanism

The performance of this mechanism relative to the dynamic mechanism is shown in Table 6. The expected revenue for the auction house in the dynamic mechanism was estimated using data on fees when trade occurred. Table 6 demonstrates that the broker-optimal mechanism would result in auction house revenues of ...

The expected auction house revenue can also be seen as the difference between the payment from the buyer to the auction house and the amount which the auction house passes on to the seller after removing fees. Note that in the broker-optimal mechanism, the payments themselves are much smaller but the gap between payments of buyers and seller is larger. The probability of trade and total expected gains from trade are both lower under the broker-optimal mechanism, as this mechanism introduces an additional deadweight loss due to the broker's rent-extraction behavior.

It is difficult to interpret these results given that a shift to this mechanism would likely drive buyers and sellers away from the auction house and toward competing sourcing venues, and this competition is not expressed in the model. Therefore, while auction house revenue is clearly of primary interest to the auction house, competition among auction houses may impede an individual auction house from achieving the payoff of the broker optimal mechanism. Townsend (1983) demonstrated, in a general equilibrium framework, that competition among auction houses, or even the threat of competition, leads to the Walrasian equilibrium as the number of buyers and sellers grows large. Thus, auction houses may appear to behave as though they were maximizing surplus rather than achieving the optimal revenue for a solo auction house. However, Economides and Siow (1988) showed, in a competition-on-a-line framework, that liquidity provides positive externalities for buyers and sellers which are not fully internalized by the auction house, and this may prevent efficient entry of auction houses and hence prevent the market from achieving the surplus-maximizing allocation. It is theoretically ambiguous how close auction houses would come to achieving the revenue-maximizing outcomes in a setting with two-sided uncertainty. For these reasons, and due to the fact that I have no data on competing auction houses, I do not focus on this mechanism.