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

Bradley Larsen*

Stanford University and eBay Research Labs

March 13, 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 300,000 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 88–96% of the surplus and 92–99% of the 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, Ariel Pakes, Paulo Somaini, and Juuso Toikka for extremely helpful suggestions. This paper also benefited from conversations with John Asker, Daniel Barron, Alessandro Bonatti, Maria Polyakova Breiter, Gabriel Carroll, Matias Cattaneo, Mingli Chen, Denis Chetverikov, Peter Cramton, Francesco Decarolis, Bernard Elyakime, Craig Farnsworth, Ian Gale, David Genesove, Kate Ho, Gaston Illanes, Dan Keniston, Kyoo-il Kim, Elena Krasnokutskaya, Katie Larsen, Greg Lewis, Thierry Magnac, Francesca Molinari, Sarah Moshary, Whitney Newey, Christopher Palmer, Christopher Parmeter, Parag Pathak, Jimmy Roberts, Nancy Rose, Mark Satterthwaite, Brad Shapiro, Matt Shum, Jean Tirole, Rob Townsend, Chris Walters, and Alex Wolitsky; as well as participants at the MIT Industrial Organization Seminar and field lunch, the MIT Econometrics and Theory field lunches, the 2012 International Industrial Organization Conference, and at seminars at Brigham Young University, HBS, NYU, University of Chicago, and Yale SOM. 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, the information-constrained efficient frontier, and to the indirect mechanisms studied in Satterthwaite and Williams (1989).

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 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 in Section 2, 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, independent, private values among bidders. 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 bidding strategy for buyers, and second, the seller's secret reserve price strategy is monotone. These two properties allow for nonparametric partial identification of the distributions of

[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 et al. (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."

buyer and seller types.

Section 4 describes the data, which comes from several independent auction houses from 2007 to 2010. The data contains 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, and cars sold by large institutions, such as rental companies, banks, and fleet companies. These two groups behave differently and tend to sell different types of cars, and hence I treat them separately in the analysis. I refer to the first group as used-car dealers and the latter group as fleet/lease sellers. The high bid from the auction exceeds the seller's secret reserve price only 16% of the time for cars sold by used-car dealers and only 37% of the time for fleet/lease sellers. Overall, the probability of trade is about three-fourths, including trade which occurs during and after the auction, motivating the question of whether or not this is an efficient level of trade.

Sections 5 through 7 constitute the heart of the paper. Section 5 presents the approach for estimating the three main objects necessary to determine the efficiency of bargaining: the distribution of buyer valuations, the distribution of seller valuations, and the costs of bargaining. After controlling for observable covariates and auction house fees, I estimate the distribution of buyer valuations using a standard order statistics inversion. I also exploit the size of and information in this dataset by obtaining multiple estimates of the distribution of bidder valuations using variation in the imputed number of bidders. I then construct an optimally weighted estimate of the buyer distribution using these different estimates.

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, implying that it can be inverted. I use moment inequalities defined by bargaining actions to bound the inverse of the seller's reserve price strategy. Without imposing that the estimates be monotonic, I find increasing upper and lower bounds on the inverse of the seller's strategy. 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. The parameters defining the costs of bargaining, which take the form of an additive and multiplicative cost, are also estimated using bounds.

Section 6 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 5 are crucial for solving for these mechanisms.

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

In Section 6 I also present the approach for solving for the equilibria of a k double auction, a simple mechanism where the buyer and seller simultaneously submit sealed bids and trade occurs if the buyer's bid is higher than the seller's, with the transaction price being a convex combination of the two bids with weight $k \in [0, 1]$. The solution approach comes from Satterthwaite and Williams (1989). Finally, I apply the Revelation Principle (Myerson 1979) to characterize a direct mechanism corresponding to the mechanism used at wholesale auto auctions. I solve for this direct mechanism using a quantile-matching approach and then use the direct mechanism to estimate buyer and seller surplus under dynamic, post-auction bargaining.

The main findings are presented in Section 7. First, I examine the 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, private 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 nearly 98% 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 limited commitment.[4] Findings indicate that the post-auction bargaining lies quite close to the efficient frontier, achieving 88–96% of the efficient level of surplus. This result is true in both the used-car dealers sample and the fleet/lease sellers sample. The deadweight loss due to limited commitment is therefore quite small in this market. However, while the bargaining does lie near the frontier, it does not lie at the point along the frontier which would maximize the total gains from trade or the probability of trade. The gains from trade and volume of trade could increase by about 14% and 8% if the bargaining were to move to a different point on the frontier.

Third, I compare the Pareto frontier and post-auction bargaining to what might be achieved under a simple $k = 1/2$ double auction. A stark difference between real-world bargaining and the theoretical

[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 et al. (1997) also referred to this as an issue of commitment.

mechanisms discussed in Myerson and Satterthwaite (1983) and Williams (1987) is that the theoretical mechanisms are direct mechanisms: the buyer and seller truthfully reveal their valuations to the mechanism designer who then follows a rule which determines whether they trade and at what price. While extremely useful as a theoretical tool, these mechanisms are unlikely to be adopted in practice as they would require the auction house to know the distributions of buyer and seller valuations. A k double auction, on the other hand, like alternating-offer bargaining, is simple to for both the auction house and the players to implement and understand. Section 7 shows various outcomes which could occur in a $k = 1/2$ double auction setting. Findings suggest that the post-auction bargaining mechanism achieves approximately the same level of surplus as the 75th percentile equilibrium of the $k = 1/2$ double auction (ranked by total surplus achieved).

Finally, I compare the region of the buyer and seller type space where trade occurs under theoretical mechanisms to where it occurs under the dynamic mechanism used at auto auctions. Plotting the region of trade is common in the theoretical bargaining literature to analyze a bargaining mechanism, but it is unknown what this region would look like in real-world bargaining. Here, as in the case of the efficient frontier, I put the theoretical bargaining and real-world bargaining on the same footing. I find that the region of trade is smaller for used-car dealers under the dynamic mechanism than under the theoretical second-best. For fleet/lease sellers the region of trade is similar to the theoretical second-best.

In Section 8 I discuss issues of model fit and robustness. I examine the fit of the estimated buyer valuation distribution using several tests from the empirical auctions literature. I also demonstrate that the estimates of the bounds on the seller's inverted secret reserve price are robust to various approaches to performing the procedure. Finally, I discuss and test several independence assumptions required for the estimation approach presented in this paper.

1.1 Related Literature

In addition to Myerson and Satterthwaite (1983) and Williams (1987), several papers examining the theoretical efficient frontier include Ausubel and Deneckere (1993), Ausubel et al. (2002), and Satterthwaite and Williams (1989). Ausubel and Deneckere (1993) and Ausubel et al. (2002) demonstrate 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) study the k double auction game and demonstrate that a continuum of equilibria exist. They also find that generically only the $k = 0$ or $k = 1$ double auctions reach the efficient frontier. Here I provide numerical evidence that equilibria of the $k = 1/2$ double auction can also come close to the frontier.

While the theoretical literature on incomplete information bargaining is large, the set of papers which structurally estimate these models is quite small.[5] These papers include Sieg (2000), Ambrus et al. (2011),

[5] 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.

and Keniston (2011).[6] The latter paper is the only one to estimate a model of two-sided uncertainty.[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 private information, can be estimated using 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 in the complete information case, there is no canonical model of alternating-offer bargaining under incomplete information with a continuum of types.[8]

Several other papers presented theoretical models of auctions followed by bargaining. Bulow and Klemperer (1996), Eklof and Lunander (2003), and Menezes and Ryan (2005) presented models of an auction followed by static bargaining where one party presents a single, take-it-or-leave-it offer to the other party, which is the optimal mechanism for the proposing party. Huh and Park (2010) and Elyakime et al. (1997) studied post-auction bargaining with alternating offers, but in both cases the authors make the assumption that bargaining takes place under complete information. This assumption allows for simpler solutions to the model, yielding the Rubinstein (1982) bargaining equilibrium, but also sacrifices an important element of reality, namely, incomplete information. Finally, Wang (2000) considered a theoretical model of post-auction bargaining at procurement auctions, where the bargaining takes place under one-sided incomplete information, adopting the equilibrium of Admati and Perry (1987). In addition to their theoretical models, Elyakime et al. (1997) and Eklof and Lunander (2003) structurally estimated the distributions of private information. The current paper is the first to present or estimate a model of an auction followed by dynamic bargaining with two-sided incomplete information.[9] Finally, Genesove (1991) 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.

[6]Note that the setting of the current paper benefits from the same features as Ambrus et al. (2011), in that players are professionals and regular participants, and the econometrician observes information, namely the secret reserve price, which is only known to one party.

[7]Merlo et al. (2008) also presented a situation two-sided incomplete information bargaining but adopted a reduced-form assumption about bargaining behavior rather than estimating these objects structurally.

[8]Most of the incomplete-information bargaining literature focuses on one-sided uncertainty, one-sided offers, or a two-type case. Two papers which model bargaining as an alternating-offer game and a continuum of types with two-sided incomplete information, where the incomplete information is about players' valuations, are Perry (1986), which predicted immediate agreement or disagreement, and Cramton (1992), which modeled the bargaining game as beginning with a war of attrition and consisting of players signaling their valuations through the length of delay between offers, as in Admati and Perry (1987). Neither of these models fits the type of bargaining observed at wholesale auto auctions. See Binmore et al. (1992), Ausubel et al. (2002), Roth (1995), and Kennan and Wilson (1993) for additional surveys of the theoretical and experimental bargaining literature.

[9]A somewhat overlapping strand of literature focuses on the use of secret reserve prices, including Vincent (1995), Horstmann and LaCasse (1997), Li and Tan (2000), Rosenkranz and Schmitz (2007), Ashenfelter (1989), Ashenfelter and Graddy (2003), Ji and Li (2008), Ji (2006), Bajari and Hortasu (2003), Katkar and Reiley (2007), Grant et al. (2006), Elyakime et al. (1997), Elyakime et al. (1994), and Eklof and Lunander (2003). One advantage of the data used in this paper is that the seller's secret reserve price is observable to the econometrician, unlike in many other studies involving secret reserves.

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 through 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. Buyers attending wholesale auto auctions are new and used car dealers.[10] Sellers may also be used or new car dealers selling off extra inventory, or they may be large institutions, such as banks, manufacturers, or rental companies selling repossessed, off-lease, or old fleet vehicles. Throughout the industry, the majority of auction house revenue comes from fees paid by the buyer and seller when trade occurs.

Sellers bring their cars to the auction house and report a secret reserve price. In the days preceding the auctioning of the car, potential buyers may view car details 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.[12] 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.[13] The entire bidding process usually takes 30–90 seconds.

If the high bid exceeds the secret reserve price, the car is awarded to the high bidder. If the high bid 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 high bid, end the negotiations, or propose a counteroffer.[14] If the seller counters, the auction house calls the buyer. Bargaining continues until one party accepts or terminates negotiations.

[10] 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.

[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] Tadelis and Zettelmeyer (2011) emphasized the fact that many lanes run simultaneously. The authors used a field experiment to demonstrate that online condition reports provide a quality disclosure mechanism which aids bidders in choosing which lane to participate in. Cho et al. (2010) presented an alternative information experiment at Korean auto auctions

[13] Lacetera et al. (2011) demonstrated that vehicle mileage, displayed on these monitors, is associated with large, discontinuous decreases in sale price as mileage readings cross 10,000-mile thresholds.

[14] If the seller a present during the auctioning of the car, the seller may choose to accept or reject the high bid immediately. This is discussed further in Section 4.

The typical time between calls is 2–3 hours.[15] 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.[16]

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.[17] 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.[18] This occurs less than three percent of the time in practice.

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 high bid exceeds the secret reserve price, the high bidder wins the item.
4. If the high bid 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.[19]

[15] 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.

[16] 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.

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

[18] 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.

[19] At wholesale auto auctions, some large institutional sellers are given the option to elect to eliminate step 4 above,

5. If the high bidder chooses to enter bargaining, the high bid becomes the first bargaining offer, and the high bidder and seller enter an alternating-offer bargaining game, mediated by the auction house.

Suppose there are N risk-neutral buyers participating in the ascending auction. Buyers have valuations, b_i, $i = 1, ..., N$, where each b_i is drawn from $F_b(\cdot)$ with atomless density $f_b(\cdot)$ and with support on $[\underline{b}, \overline{b}]$. Let the subscripts order values from highest to lowest (i.e., b_1 represents the max order statistic, b_2 represents the second order statistic, etc.). Suppose a risk-neutral seller has private value s, where s is drawn from $F_s(\cdot)$ with atomless density $f_s(\cdot)$ and with support on $[\underline{s}, \overline{s}]$. Throughout the paper I maintain the following assumptions:

Assumptions.

(A1) Buyers have symmetric, independent, private values (IPV), and the ascending auction follows a button auction model.

(A2) $F_s(\cdot)$ *and* $F_b(\cdot)$ *are independent.*

(A3) In bargaining players face a per-period disutility, $(c_B, c_S) > 0$*, as well as discount factor due to impatience,* $\delta_I \in [0, 1)$*, and a probability bargaining does not break down exogenously,* $\delta_E \in [0, 1)$*.*

A motivation for the IPV assumption is that buyer valuations arise primarily from their local demand and inventory needs.[20] Kim and Lee (2008) tested and failed to reject the symmetric IPV assumption at wholesale auto auctions.[21] 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.[22]

Like buyers, sellers who are used-car dealers have valuations arising from their local demand/inventory. Also, seller valuations can depend on the value at which the car was assessed as a trade-in. For a bank or

implying that when the high bid does not meet the secret reserve price, the high bidder is not allowed to immediately walk away from bargaining but must wait until the seller responds to the high bid. 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.

[20] 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).

[21] Roberts (2010) found contrasting evidence of unobserved auction-level heterogeneity at the same Korean auctions. The papers differ in approach: Kim and Lee (2008) presented an approach for testing the IPV assumption when the number of bidders is unknown, whereas Roberts (2010) estimated a model with and without accounting for unobserved heterogeneity. It is unclear why the results differ.

[22] Bid increments lie between \$25 and \$100. In contrast to the auctions studied in this paper, jump bidding does occur in the online auto auctions analyzed in Grether et al. (2012).

leasing company, valuations can arise from the size of the defaulted loan.[23] These sources of valuations serve as motivation for the assumption that buyer and seller types are independent. This does not preclude a common component among buyer and seller willingness to buy/sell, but I assume that this component is purely a function of observable characteristics of the car.

Bargaining costs of the kind mentioned in A3 are common in the literature.[24] Evidence that bargaining is costly lies in the facts that 1) bargaining does not continue indefinitely, having at most 10 counteroffers in the data; and 2) three percent of buyers, after learning the high bid did not meet the secret reserve, choose to walk away from bargaining before it begins. The parameter δ_E represents the probability that bargaining does not exogenously break down. In the data, exogenous breakdowns are observed when an auction house employee terminates the negotiations rather than the buyer or seller choosing to do so, occurring in 1–3% of bargaining interactions.[25]

In addition to the above assumptions, I model the seller as having some ex-ante uncertainty about the distribution of buyer valuations. This uncertainty is resolved once the auction takes place and the seller sees (or learns over the phone through an auction house employee) additional information, such as the level of buyer turnout/interest, which can be affected by weather, financial news, or other shocks. Such uncertainty is discussed by auto auction participants and is essential to rationalize why some sellers would set a reserve price and then accept a high bid below the reserve price, as discussed in Section 4. Let ζ be a finite vector parameterizing this uncertainty, where ζ is independent of buyer and seller valuations.[26] As only the secret reserve price will depend on this uncertainty, I use the convention that $F_b(\cdot)$ denotes the distribution of buyer valuations once ζ has been realized and is observed to all players, and I only use the notation $F_b(\cdot;\zeta)$ in Appendix A in discussing the seller's secret reserve price, which is set prior to knowing ζ.

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, $R(s)$, is strictly increasing in her type s. Note that throughout this section, auction house fees and auction heterogeneity are ignored. I discuss how these issues are handled in Section 5.1.

[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]See, for example, Cramton (1991). Note that the main reason for including discounting in this bargaining model is because Perry (1986) demonstrated that if bargaining costs take the form of an additive cost common to all buyers and an additive cost common to all sellers, as in this paper, and if there is no discounting/probability of breakdown, then the unique equilibrium is for bargaining to end immediately.

[25]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.

[26]Section 5.1 explains that a simple form this uncertainty might take would be an overall mean effect on prices on any given auction sale day, and the seller does not know the realization of this mean effect until the auction day arrives.

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 high bid at the auction, p^A. The seller then chooses between accepting (A), declining (D)—meaning terminating the negotiations—or making a counteroffer (C). Accepting ends the game, with trade taking place at the accepted price. Declining 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, declining, 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.

Suppose it is the seller's turn at time t. Let $h^t = \{p_\tau\}_{\tau=1}^{t-1}$ represent the history of previously rejected counteroffers by the players as well as the most recent offer on the table. 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 $d_t^S \in \{A, D, C\}$ represent the seller's decision in period t, and let $d_{t+1}^B \in \{A, D, C\}$ represent the buyer's decision in period $t+1$.

The seller's payoff at time t is given by the following. When the buyer's most recent offer is p_{t-1}^B, a seller of type s chooses to accept (A), decline (D), or counter (C), yielding the following payoffs:

$$
\begin{aligned}
&\text{A} : p_{t-1}^B \\
&\text{D} : s \\
&\text{C} : V_t^S\left(s|h^t\right) \\
&\quad = \delta_I\Bigg(\max_p p\delta_E \Pr\left(d_{t+1}^B = A|h^{t+1}\right) + s\left(\delta_E \Pr\left(d_{t+1}^B = D|h^{t+1}\right) + 1 - \delta_E\right) \\
&\quad + \delta_I\delta_E \Pr\left(d_{t+1}^B = C|h^{t+1}\right)\left(\delta_E E_{p_{t+1}^B}\left[\max\left\{p_{t+1}^B, s, V_{t+2}^S\left(s|h^{t+2}\right)\right\}\middle| h^{t+1}\right] + s(1-\delta_E)\right)\Bigg) - c_S
\end{aligned}
$$

where p is the counteroffer chosen by the seller. The per-period bargaining disutility ($c_S > 0$) is assumed to be common across sellers, and discounting due to impatience ($\delta_I < 1$) and the probability of not terminating exogenously ($\delta_E < 1$) are assumed to be common across sellers as well as buyers. The seller's counteroffer payoff takes into account that the buyer may either accept, decline, 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, decline, or counter. Exogenous breakdown may occur in any period, in which case the seller receives s as a payoff.

Similarly, the buyer's payoff at time $t+1$ is given by the following. When the seller's most recent offer

is p_t^S, a buyer of type b chooses to accept (A), decline (D), or counter (C), yielding the following payoffs:

$$\begin{aligned}
&\text{A} : b - p_t^S \\
&\text{D} : 0 \\
&\text{C} : V_{t+1}^B\left(b|h^{t+1}\right) \\
&\quad = \delta_I\delta_E\Bigg(\max_p (b-p)\Pr\left(d_{t+2}^S = A|h^{t+2}\right) \\
&\quad + \delta_I\delta_E \Pr\left(d_{t+2}^S = C|h^{t+2}\right) E_{p_{t+2}^S}\left[\max\left\{b - p_{t+2}^S, 0, V_{t+3}^B\left(b|h^{t+3}\right)\right\}\middle| h^{t+2}\right]\Bigg) - c_B
\end{aligned}$$

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, and hence δ_E enters the buyer's counteroffer payoff in the same fashion as δ_I.[27]

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, which represents his valuation for the car. A seller's strategy is a secret reserve price, R, which is a function of her type, s, the least she is willing to sell the car for. Let

$$\beta_j(b_j) = \max_{k \neq i} \beta_k(b_k)$$

That is, i will be the high bidder if and only if $\beta_i > \beta_j$. The expected payoff of bidder i from following bidding strategy $\beta_i(b_i)$ is given by

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

Buyer i would decide to enter bargaining if $\pi_B(\beta_j, b_i) > 0$, where $\pi^B(\beta_j, b_i)$ represents the buyer's expected payoff from entering bargaining, equivalent to the counteroffer payoff in the previous section but with the high bid 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 high bid, truth-telling is weakly dominant for bidders in the auction stage. That is,* $\beta_i(b_i) = b_i$.

[27] In reality, the outside option of a buyer is a complicated object that cannot be estimated in the scope of this data, as buyers who exit bargaining have the choice to obtain vehicles from a variety of sources, such as other sales at the same auction house, competing auction houses, online markets, and trade-ins.

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.[28] Intuitively, the assumption that the seller never accepts bargaining offers below the high bid 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 high bid. Moreover, bidders are not tempted to drop out before the bidding reaches their valuations because if the high bidder learns that the auction high bid did not meet the secret reserve, he can always opt out of bargaining. And because the high bid at the auction 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. Let p^A represent the high bid from the auction.

In choosing her secret reserve price, R, the seller wishes to maximize her ex-ante payoff.

$$E_\zeta \left[E_{p^A,b} \left[p^A * 1\left\{p^A > R\right\} + s * 1\left\{p^A < R, \pi^B(p^A, b) \leq 0\right\} + \pi^S\left(p^A, s\right) * 1\left\{p^A < R, \pi^B(p^A, b) > 0\right\} \middle| \zeta \right] \right] \tag{2}$$

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, $\pi^S\left(p^A, s\right) = \max\left\{p^A, s, V_2^S\left(s, p^A\right)\right\}$, which the seller receives when the high bid is below the reserve and bargaining occurs. The seller's bargaining payoff, the auction high bid, and the buyer types which would enter bargaining will all depend on the realization of ζ, which is unknown to the seller when choosing R. See Appendix A for further details.

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, $R^*(s)$, 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 $R(s)$ will be weakly increasing.

[28] 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.

4 Dataset

This section provides an overview of the data used in this paper, including summary statistics and a breakdown of when trade occurs in the wholesale auto auction mechanism. The data used in this paper come from 5-10 independent auction houses, each maintaining a large market share in the region in which it operates.[29] 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, odometer reading, accessories/options, condition report (prepared by the auction house), quality report (submitted by the seller), and vehicle identification number (VIN).[30] When bargaining occurs over the phone, the full sequence of buyer and seller actions (accept, decline, or counter) and offers/counteroffers are recorded. The data also records the seller's secret reserve price, the winning bid at the auction, a timestamp for the arrival of each bidder, and the identity of the seller and high bidder.

It is important to note that in this paper an observation refers to a run of a 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.[31]

I drop observations with no recorded auction house blue book estimate, cars less than one year or greater than 16 years old, observations with less than 100 miles or greater than 300,000 miles on the odometer, observations with incomplete bargaining sequences, and observations with normalized high bids or normalized reserve prices (where the normalization is described in Section 5.1) greater than the 0.99 quantiles of these variables..[32] I do not use observations where the high bid is missing.[33] I also drop observations for which the secret reserve price is missing, which usually occurs because the seller plans to be present at the auction sale, and hence would be able to accept or reject the high bid in person.[34] Observations are also dropped where the seller's reserve price was reportedly zero.[35] Finally, I

[29]These auction houses wish to remain anonymous, so the precise number of auction houses is not disclosed.

[30]The VIN contains information on vehicle characteristics. See Lacetera and Sydnor (2012) for an explanation of VINs and VIN decoding.

[31]Section 8.3 discusses this issue further.

[32]Together, observations falling into these categories make up 19% of all runs. Note that observations with no blue book estimate tend to be motor homes, boats, or other recreational vehicles.

[33]These account for 20% of all runs. These observations could be rationalized by a model of entry costs in the auction. In practice, 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 cancel the sale of the vehicle, leaving no high bid recorded. Given that bidders know they will only be awarded the car if they meet the secret reserve price or agree in bargaining, a small entry cost could lead to some auction sales where no bidders participate. Because the focus of this paper is on the post-auction bargaining, I do not model entry into the auction but instead work with the empirical distribution of auction outcomes observed in the data and hold this distribution fixed in counterfactual analysis.

[34]These account for 14% of all runs. In this scenario, the seller is effectively enforcing a secret reserve price even though none was recorded in the data. Therefore, I treat these as missing observations and omit them from the analysis.

[35]These account for less than one percent of all runs. It is unclear whether these observations were correctly recorded in

Table 1: Summary statistics: Number of bidders, prices, and car characteristics.

	Used-car dealers			Fleet/lease sellers		
	mean	s.d.	median	mean	s.d.	median
Number of bidders	60	39	51	90	83	64
Auction high bid	\$6,324	\$5,337	\$4,650	\$9,973	\$6,626	\$8,700
Secret reserve	\$7,253	\$5,605	\$5,500	\$10,173	\$6,561	\$9,000
Final price	\$6,346	\$5,347	\$4,668	\$9,987	\$6,623	\$8,800
Odometer reading	95,315	47,869	92,311	60,746	43,546	50,186
Age in years	6.45	3.45	6.00	3.49	2.92	3.00

Notes: Number of bidders is imputed by dividing the number of bidders present at the auction house by the number of active lanes at the moment the car was sold. Statics reported for the auction high bid, secret reserve, and final price are conditional on trade occurring.

drop observations for which fewer than ten vehicles were observed at a given make-model-year-trim-age combination.[36] In the end, I am left with 168,745 runs of cars sold by used-car dealers, and 182,843 sold by fleet/lease sellers.

Summary statistics are displayed in Table 1. The estimated number of bidders is constructed by dividing the number of bidders present at the auction house by the number of active auction lanes at the moment the car was sold, as described in Appendix B. The median number of bidders at fleet/lease sales is 64, while it is only 51 when cars are sold by used-car dealerships. Table 1 also displays the auction high bid, seller's secret reserve price, and the final transaction price, each calculated conditional on trade occurring. Prices are higher in the fleet/lease sample than in the used-car dealer sample. In each sample, the average secret reserve price exceeds the average high bid at the auction, and the final transaction price lies in between, close to the high bid. The auction high bid tends to be closer to sellers' reserve prices in the fleet/lease sample (\$200 below as opposed to \$900 below). The odometer and age rows demonstrate that used-car dealers tend to sell cars which are older and have more mileage.

Table 2 displays the proportion of observations in the used-car dealers sample and fleet/lease sellers sample which end in each major stage of the game, as well are the proportions which end in trade. The first column indicates that the auction high bid exceeds the reserve price only 15.77% of the time in the used-car dealer sample and 36.85% of the time in fleet/lease sample. Between the two samples (not shown in Table 2), the auction high bid fails to meet the reserve price about 73% of the time at these

the data, as in some cases a no-sale followed a reserve of zero. However, the main reason for dropping these observations is computational: the analysis below is simplified as there is no need to consider a possible atom in the seller's distribution at zero.

[36] These account for 12% of all runs.

Table 2: Probability of trade

	Used-car dealers (168,745 observations)		Fleet/lease sellers (182,843 observations)	
	% of Sample	% Trade	% of Sample	% Trade
High bid ≥ Reserve	15.77%	100.00%	36.85%	100.00%
High bid < Reserve cases				
A) Immediate agreement/ disagreement	61.45%	76.15%	44.16%	68.30%
B) Phone bargaining	22.78%	39.47%	18.99%	55.52%
Overall trade volume	71.55%		77.55%	

Notes: For each sample, the first column reports the proportion of the sample with the game ending at the auction, immediately after the auction, or in alternating-offer bargaining over the phone. The second column reports the percentage of time that trade occurs in each case.

auctions.

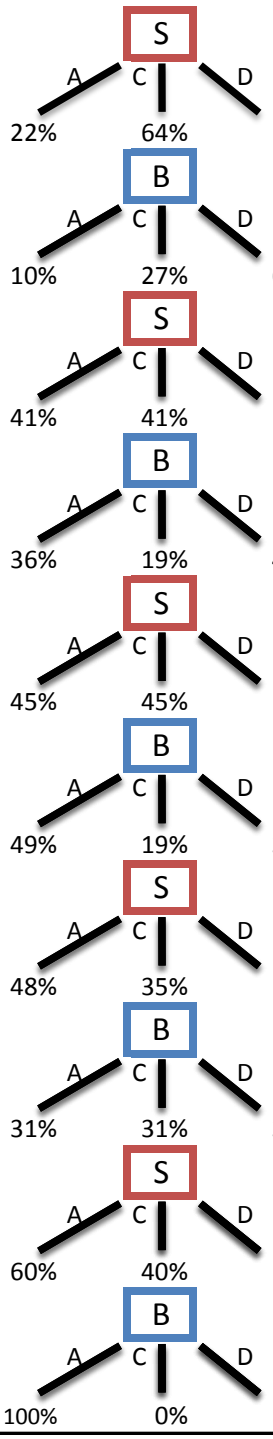
When the reserve is not met, the game ends in immediate agreement or disagreement 61.45% of the time for vehicles sold by used-car dealers and 44.16% of the time for cars sold by fleet/lease sellers. This immediate bargaining occurs when one of the following is true: 1) the seller is present at the auction house and can accept or reject the high bid; 2) the auctioneer rejects the high bid 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, which occurs about two percent of the time (not shown in Table 2) in both samples. As explained in Section 3, seller's setting secret reserve prices and then accepting offers below these reserve prices can be explained by demand uncertainty (buyer turnout/interest) which is resolved once the auction takes place.

Table 2 indicates that about 20% of the time, in both samples, the buyer and seller enter alternating-offer bargaining over the phone. In the used-car dealers sample, this bargaining results in trade 39.47% of the time and for the fleet/lease sellers sample, trade occurs 55.52% of the time. The overall probability of trade at these auction houses is slightly under three-fourths for cars sold by used-car dealers and slightly above three-fourths for fleet/lease sellers.

Figures 1 and 2 summarize the outcomes when bargaining occurs over the phone, beginning with the seller's decision to accept (A), decline (D), or counter (C) when faced with the auction high bid. The figures display the overall percentage of observations which reach each stage of bargaining, as well as the absolute number of observations. The figures also display the percentage of players who choose to accept, decline, or counter conditional on reaching period t of bargaining.

Observe in Figure 1 that when a used-car dealer is contacted over the phone at $t = 2$, 64% of the time she responds with a counteroffer, and 63% of those counteroffers are rejected by the buyer at $t = 3$. Buyers are more likely to decline and less likely to counter or accept than sellers. Only 7.29% of phone

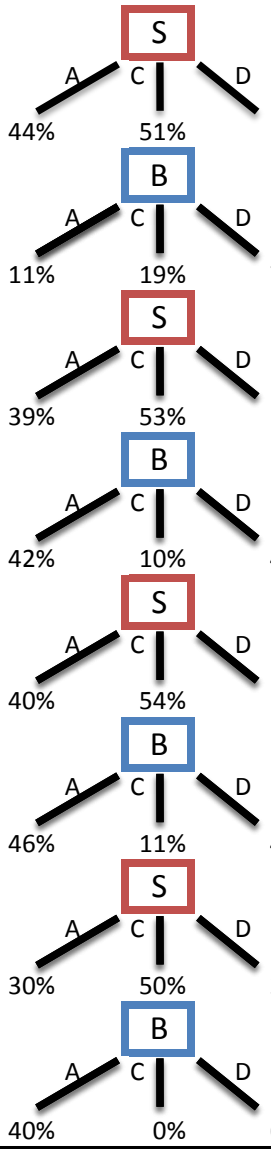
Figure 1: Summary of phone bargaining outcomes, used-car dealer sample



				# obs	% obs
t=2	22%	64%	14%	38,446	100.00%
t=3	10%	27%	63%	24,716	64.29%
t=4	41%	41%	17%	6,761	17.59%
t=5	36%	19%	44%	2,803	7.29%
t=6	45%	45%	10%	544	1.41%
t=7	49%	19%	32%	244	0.63%
t=8	48%	35%	17%	46	0.12%
t=9	31%	31%	38%	16	0.04%
t=10	60%	40%	0%	5	0.01%
t=11	100%	0%	0%	2	0.01%

Notes: Percentages at each node in period t describe the probability of the action accept, decline, or counter (A, D, C) conditional on arriving to period t. The final columns list the number and percentage of observations with phone bargaining which reach period t.

Figure 2: Summary of phone bargaining outcomes, fleet/lease sample



				# obs	% obs
t=2	44%	51%	5%	34,731	100.00%
t=3	11%	19%	70%	17,821	51.31%
t=4	39%	53%	8%	3,327	9.58%
t=5	42%	10%	49%	1,756	5.06%
t=6	40%	54%	7%	168	0.48%
t=7	46%	11%	43%	90	0.26%
t=8	30%	50%	20%	10	0.03%
t=9	40%	0%	60%	5	0.01%

Notes: Percentages at each node in period t describe the probability of the action accept, decline, or counter (A, D, C) conditional on arriving to period t. The final columns list the number and percentage of observations with phone bargaining which reach period t.

bargaining observations reached the fifth period of bargaining or beyond. However, this still leaves 2,803 records reaching this period. Two pairs of players endured 11 periods of bargaining, finally accepting in the end.

Figure 2 displays similar patterns for the fleet/lease sample. 51% of the time, the seller responds to the high bid by proposing a counteroffer, which the buyer declines 70% of time. 5.06% of observations reach period $t = 5$, and five pairs endured nine rounds of bargaining, with two of these pairs trading in the end.

5 Estimation of Buyer and Seller Distributions and Bargaining Costs

In this section, I exploit the model properties derived above in order to estimate the distribution of buyer and seller valuations as well as the bargaining cost parameters. I first perform two preliminary steps, controlling for observable heterogeneity and adjusting for auction house fees. I then estimate the distribution of buyer valuations using an order statistics approach. To estimate seller valuations, I present a moment inequalities approach using observed bargaining actions to bound the inverse of the seller's secret reserve price. Bargaining costs are also estimated using a bounds approach.

5.1 Preliminary Steps: Heterogeneity and Auction House Fees

This section describes how vehicle heterogeneity and auction house fees are handled. The model described in Section 3 is for the sale of a homogeneous good but can easily be extended to allow for heterogeneity by controlling for auction-level covariates in a first-stage regression, as is now common in the empirical auctions literature (see Haile et al. 2003). To do so, I first state the following lemma, similar to results found in Haile et al. (2003) but generalized to this specific mechanism. Let $\Gamma(X_j)$ be a scalar summarizing the observable value of vehicle j, where X_j contains auction-level covariates.

Lemma (1)**.** *Suppose the following:*

(a) When the observable value of the car is given by $\Gamma(X)$, the seller and buyer values are $\Gamma(X)s$ and $\Gamma(X)b$, and the additive costs of bargaining are $\Gamma(X)c_S$ and $\Gamma(X)c_B$.

(b) When $\Gamma(X) = 1$, the optimal secret reserve is R, the optimal bargaining offer in period t is p_t^S, and the optimal offer in period $t+1$ is p_{t+1}^B.

Then, for any positive $\Gamma(X)$, the optimal secret reserve will be $\Gamma(X)R$ and the optimal offers will be $\Gamma(X)p_t^S$ and $\Gamma(X)p_{t+1}^B$.

The same result follows immediately for the auction high bid. Using this lemma, the observed reserve prices, bids, and counteroffers in auction/bargaining sequence j can be homogenized by dividing by an

estimate of $\Gamma(X_j)$. As in common in the auction literature, I specify $\Gamma(X_j) = \exp^{X_j\theta}$ and estimate θ by the following regression

$$\ln \tilde{R}_j = X_j\theta + \ln R_j \tag{3}$$

where $\tilde{R}_j$ is the secret reserve price observed by the econometrician and R_j is the underlying, homogenized secret reserve.[37] The vector X_j contains the auction houses' blue-book estimate, the odometer reading, the number of pictures displayed online, and a dummy for used-car dealer vs. fleet/lease seller, as well as dummies for each make-model-year-trim-age combination (where age refers to the age of the vehicle in years), odometer bin, condition report grade (ranging from 1-5), auction house location, and year-month (of the sale).[38] I also include dummies for 32 different vehicle damage categories recorded by the auction house and for each fleet/lease seller who sells over 500 cars in the sample. The R^2 from this first-stage regression is 0.89, implying that most of the variation in reserve prices is explained by observables.[39]

Lemma 1 is extremely useful from a computational standpoint, but part (a) does embed two unattractive features. The first is the assumption that additive costs of bargaining scale with the value of the car. This does not end up being restrictive, as I find these costs to be very small. The second is that the function of observables $\Gamma(X)$ entering the buyer's overall value is the same as the function entering the seller's overall value. While analagous assumptions are common in the empirical auctions literature, the assumption implies that the first-stage regression approach cannot be used to directly control for differences in buyer and seller overall valuations which are due to the seller's demand uncertainty mentioned in Section 3.[40] Instead, this uncertainty can be controlled for by explicitly estimating each piece of the model in different subsamples of the data in which ζ should differ, as discussed in Section 8.3.[41]

Auction house fees are a function of the transaction price, and are approximately linear. Let

$$w^B(p) = \alpha_B + \beta_B p$$
$$w^S(p) = \alpha_S + \beta_S p$$

[37] As explained below, θ could also be estimated using the auction high bid instead of the secret reserve price.

[38] Hortaçsu et al. (2010) demonstrated that market prices fluctuated weekly during a portion of the sample period I use. I find that the estimates from (3) are unchanged when I use year-week instead of year-month effects. I also find similar results when I use the interaction of location dummies with year-month effects.

[39] I estimate regression (3) pooling together the fleet/lease sellers and used-car dealers samples. All other estimation in the paper treats the two samples separately.

[40] For example, if this uncertainty is due to the seller not knowing the value of the year-month or year-month times location effect entering $\Gamma(X)$, then Lemma 1 does not allow for one to control for this uncertainty simply through the first-state regression.

[41] I find that the correlation is 0.98 between $\hat{\Gamma}_R(X)$, the estimate obtained from the regression in (3), and $\hat{\Gamma}_{p^A}(X)$, the estimate obtained from a regression using high bids instead of secret reserve prices. Also, the correlation between secret reserve prices normalized by $\hat{\Gamma}_R(X)$ and secret reserve prices normalized by $\hat{\Gamma}_{p^A}(X)$ is 0.71. Similarly, the correlation between high bids normalized by $\hat{\Gamma}_R(X)$ and high bids normalized by $\hat{\Gamma}_{p^A}(X)$ is 0.74. These correlations are suggestive that restricting $\Gamma(X)$ to be the same across buyers and sellers may not be greatly affecting the estimated distributions of buyer and seller valuations. Note also that the proof of Lemma 1 implies that if $\Gamma(X)$ is not common to buyers and sellers, one could condition estimation of the distributions of buyer and seller valuations on $\hat{\Gamma}_R(X)$ and $\hat{\Gamma}_{p^A}(X)$. This is discussed further in Section 8.3.

denote the fee paid by the buyer and seller, respectively, when the transaction price is p. Note that these fees can vary by auction house and can change from year to year, but the fees are observed in the data when transactions occur and thus the fee structure can be estimated with a simple linear regression, estimating the regression separately for each auction house and year.

After estimating $w^B(p)$, $w^S(p)$, and $\Gamma(X)$, I adjust observed bids and bargaining offers as follows. Let $\bar{p}^A$ represent the unadjusted, original high bid. Let $\bar{p}$ denote an original, unadjusted bargaining offer. These offers are only used in the estimation of the seller type, and hence are adjusted by the seller fee. The fee-adjusted, homogenized high bids and offers are then given by

$$p^A \equiv \frac{\overline{p}^A + w^B(\overline{p}^A)}{\Gamma(X)} \qquad \text{and} \qquad p \equiv \frac{\overline{p} - w^S(\overline{p})}{\Gamma(X)}$$

Note that the adjustment for auction house fees is just an approximation, because in reality the fixed components of fees, α_S and α_B, do not vary with $\Gamma(X)$, but I treat them as though they do in order to be able to apply the multiplicative structure implied by Lemma 1. This approximation should have a negligible affect overall given that in most cases the fixed component of fees is small (about \$130) relative to the transaction price of a car (\$8,000 on average).

After performing this adjustment, one additional rescaling is useful in order to put all prices and private valuations in terms of a market value of the vehicle. I estimate the overall average of the homogenized final transaction price, $\mu = 1.03$, and divide all prices by μ. Throughout the remainder of this paper, prices and private values refer to the fee-adjusted, homogenized prices/values, scaled by $1/\mu$ in order that the final price will have unit mean. I refer to prices and values as being fractions of the market value or observable value of the car.

5.2 Estimation of Buyers' Distribution

In this section I present the approach for estimating the distribution of buyer valuations and then display the resulting estimates. Proposition 1 above implies that the high bid from the auction will be the second order statistic of bidder valuations, as is the case in a standard ascending auction. Therefore, $F_b(\cdot)$ is nonparametrically identified from the distribution of observed high bids from auction sales in which there were N bidders present. As discussed in Athey and Haile (2007), the following relationship holds between the distribution of the second order statistic, $F_{p^A}(\cdot)$, and the underlying distribution, $F_b(\cdot)$:

$$F_{p^A}(b) = NF_b(b)^{N-1}(1 - F_b(b)) + F_b(b)^N \tag{4}$$

One can solve numerically for $F_b(b)$ over a grid of values for b. Similarly, the density of buyer valuations, $f_b(b)$, is given by

$$f_b(b) = \frac{f_{p^A}(b)}{N(N-1)F_b(b)^{N-2}(1 - F_b(b))}$$

where $f_{p^A}(b)$ is the density of high bids. Different estimates of $F_{p^A}(b)$ are available using different values of N. Therefore, I use these estimates from different N to obtain an optimally weighted estimate of $F_b(b)$,

constraining the function to be monotone. This procedure is mentioned by Athey and Haile (2007), but as it is has not been implemented in practice (to my knowledge), I provide computational details in Appendix C.1. In practice, I use auction observations in which N was among the ten most common values for N observed in the sample. Depending on the sample, these ten values range from about 45 to 60.

To estimate $f_{p^A}(b)$ and $F_{p^A}(b)$, I employ kernel density estimation using a Gaussian kernel and the integral of a Gaussian kernel, respectively, with the asymptotically optimal bandwidth.[42] The bounds of the buyer support, $[\underline{b}, \overline{b}]$, used below in counterfactual integration, are chosen to be the lower 0.001 and upper 0.999 quantiles of the observed auction high bid.[43]

Figures 3 and 4 display, respectively the original distribution/density of high bids and the estimated, underlying distribution/density.[44] Panels (c) and (d) of Figure 3 demonstrate that the distribution of high bids is approximately centered about one in both samples, although it is more tight for cars sold by fleet/lease sellers. Panels (a) and (b) demonstrate that in both samples less than ten percent of cars yield auction high bids lower than 0.5 or greater than 1.5 times the observable value of the vehicle.

Figure 4 demonstrates, in panels (c) and (d), that the estimated, underlying distribution of buyer valuations has more mass at lower values of the buyer support than does the density of high bids, as should be the case given that the high bid is the second order statistic among bidders drawn from this underlying distribution. Panels (a) and (b) demonstrate that 70% of the time buyers have valuations lower than one, indicating that most buyers are not willing to pay the amount for which the car typically sells for at auction. This makes intuitive sense given that buyers have differing inventory needs. Cars sold by used-car dealers tend to display more mass at very low values than do fleet/lease sales. However, the density appears quite similar in both samples. In both samples, buyer valuations exceed 1.5 times the observable value of the car less than two percent of the time.

5.3 Estimation of Bounds on Sellers' Types and Distribution

In this section I present the moment inequality approach for estimating bounds on the seller's secret reserve strategy and then display the resulting estimates. Proposition 2 showed that the seller's secret reserve price strategy, $R(s)$, is strictly increasing in s. Therefore, the reserve price can be inverted to

[42] The asymptotically optimal bandwidth is given by $h = 1.06\hat{\sigma}n^{-\frac{1}{5}}$, where n is the number of observations in the sample used in estimating the density/distribution, and $\hat{\sigma}$ is the sample standard deviation. Instead of using kernel estimation to estimate $F_{p^A}(b)$, one could use the empirical CDF of high bids. I experimented with this approach as well and found similar estimates overall, but using a kernel estimate for the CDF aided in obtaining an estimated distribution and density which satisfied $F_b(b) = \int_{\underline{b}}^{b} f_b(v)dv$. This is important in the analysis of counterfactual mechanisms below, where I estimate a variety of integrals over the estimated densities and distributions. Henderson et al. (2011) also advocated using a kernel for both the estimation of a density and the corresponding distribution in an auction setting.

[43] The counterfactual analysis relies on choosing nodes for numerical integration, as explained in Appendix C.2. For a fixed number of nodes, the integration is more accurate over a smaller support, so I choose the bounds as described. Overall this choice has little or no effect on the counterfactuals other than increasing accuracy.

[44] Note that Figure 3 pools all auctions together for illustrative purposes. The actual distribution of high bid varies depending on the number of bidders present.

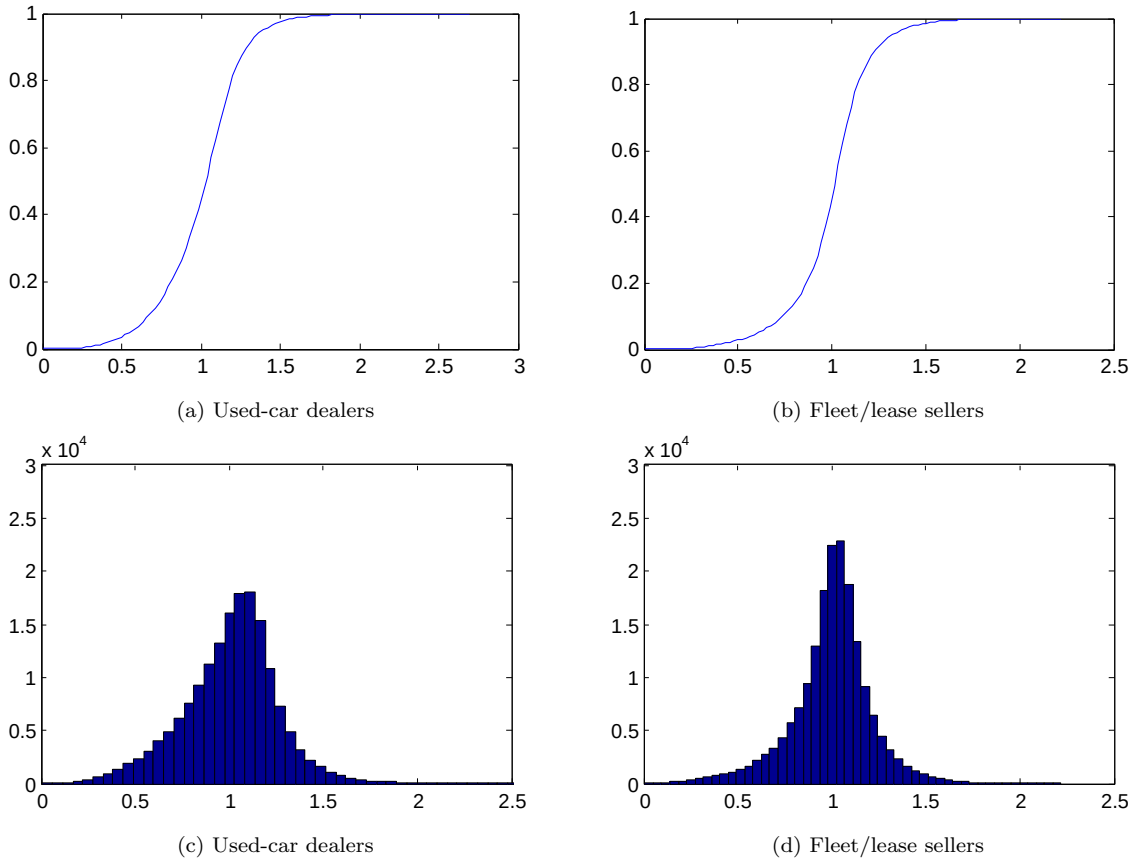


Figure 3: Empirical CDF and histogram of normalized high bid in used-car dealers sample and fleet/lease sellers sample. Units are in terms of the observable value of the car.

obtain the seller's type as a function of the reported reserve, $s(R)$. Post-auction bargaining actions yield bounds on $s(R)$:

$s(R) \leq p^B$ if p^B is a buyer offer which the seller accepted

$s(R) \geq p^B$ if p^B is a buyer offer at which the seller declined (exited)

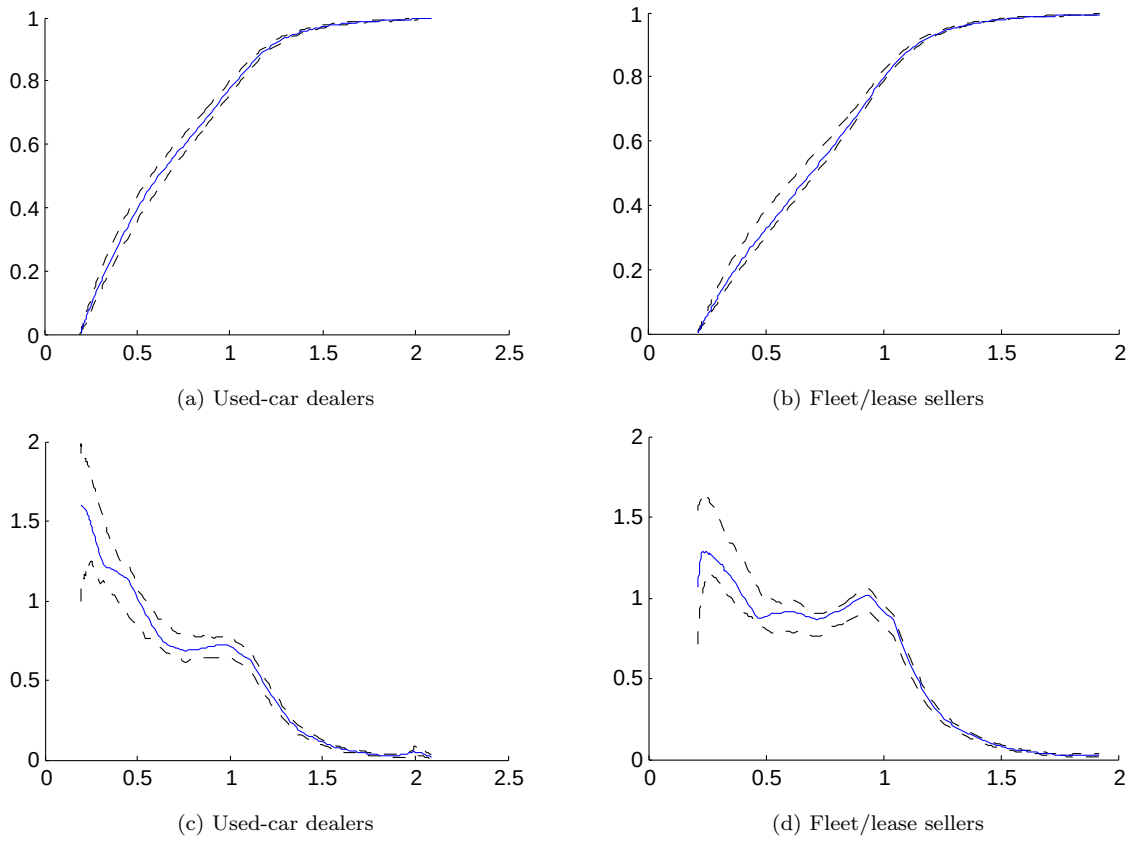
$s(R) \leq p^S$ if p^S is a price at which the seller countered

In order to estimate these bounds, I assume that all observations were generated under a single equilibrium of the post-auction bargaining game. This assumption can be relaxed by estimating the model separately on different subsamples of data, which I do in Section 8.3.[45]

One explanation for violations of these bounds is measurement error in the recorded bargaining offers, and the estimation approach I present here is consistent with this explanation.[46] In practice, violations

[45] Model estimates may differ in different subsamples for other reasons as well, as discussed in Section 8.3.

[46] Note that these offers are recorded by an auction house sales representative who is mediating negotiations for multiple parties at once after the auction takes place. The reserve price, on the other hand, is chosen several days before the sale



(a) Used-car dealers

(b) Fleet/lease sellers

(c) Used-car dealers

(d) Fleet/lease sellers

Figure 4: Estimated CDF and density of buyer valuations in used-car dealers sample and fleet/lease sellers sample. Pointwise 95% confidence bands shown by dashed lines. Units are in terms of the observable value of the car.

at the individual observation level may also occur due to heterogeneity in individual sellers' secret reserve price strategies which is not captured in this model, and in this case the measurement error model described here is simply an approximation. See Section 8.2 for further discussion.

First, note that not all observations in the data contain a lower bound; an informative lower bound (other than $\underline{s}$) is only observed if bargaining occurred and the seller declined a buyer offer sometime during bargaining, or if immediate disagreement occurred. Similarly, an informative upper bound (other than R itself or $\overline{s}$) is only observed if the seller did not immediately decline the auction high bid. Let bargaining sequences with an informative lower bound be represented by $L_j = L_j^* + \nu_j$, where L_j^* is the

and is observed by sales representatives, the auctioneer, and a recording clerk who aides the auctioneer. Thus, human measurement error is more likely to occur in the recording of bargaining offers than in the reserve price. In practice, the first offer in bargaining, the auction high bid, can also be recorded with error as all bargaining offers, including the high bid, are recorded in a separate portion of the database from auction outcomes. In the estimation of, however, I use the same recorded value of the auction high bid in estimating the buyer's distribution and in bounding seller types.

correctly measured lower bound and ν is a mean-zero error term.[47] Similarly, let $U_j = U_j^* + \nu_j$ represent a bargaining sequence with an informative upper bound. Let $\mathcal{U}$ and $\mathcal{L}$ represent the set of observations (j) for which an informative upper or lower bound is observed, respectively.

Bounds on the function $s(R)$ are then given by

$$E[L^*|R] = E[L^* + \nu|R] \leq s(R) \leq E[U^* + \nu|R] = E[U^*|R] \tag{5}$$

The expression in (5) states that the seller's type, s, which is a one-to-one function of R, lies between the lower and upper bounds implied by bargaining actions when the reserve price was R. Bounds estimation can then be implemented by dividing the support of R into K bins with end points $\{r_0, r_1, ..., r_K\}$. For $k = 1, ..., K$, let $\mathcal{R}_k = \{R_j : r_{k-1} \leq R_j < r_k\}$ be the set of observed reserve prices lying in bin k. Further, let $\mathcal{R}_k^U = \{j : R_j \in \mathcal{R}_k, j \in \mathcal{U}\}$ and $\mathcal{R}_k^L = \{j : R_j \in \mathcal{R}_k, j \in \mathcal{L}\}$ be the set of observed reserve prices in bin k which have an informative lower and upper bound, respectively. Let $\bar{R}_k$ represent the median of the bin k. Then, averaging over the lower and upper bounds in a given bin k yields

$$\frac{1}{\#(\mathcal{R}_k^L)} \sum_{j \in \mathcal{R}_k^L} L_j \leq s(\bar{R}_k) \leq \frac{1}{\#(\mathcal{R}_k^U)} \sum_{j \in \mathcal{R}_k^U} U_j$$

where $\#(\cdot)$ returns the number of elements for a set. All values of $s(\bar{R}_k)$ satisfying these bounds constitute the identified set for that grid point. For the counterfactual analysis below, I select either the midpoint of the estimated bounds or the lower or upper 95% confidence band as an estimate of $\hat{s}(R)$. Linear interpolation is used to evaluate $\hat{s}(R)$ at non-grid points. To estimate the density and distribution of $\hat{s}$, I use kernel estimation with a Gaussian kernel and the intergral of Gaussian kernel, as in Section 5.2.[48] The bounds of the seller support, $[\underline{s}, \overline{s}]$, are chosen to be the lower 0.001 and upper 0.999 quantiles of $\hat{s}$.

Monotonicity of the seller's strategy is not ex-ante imposed in estimation. However, it can be used ex-post to obtain tighter bounds. This can be performed by simply replacing the lower bound at R with the greatest lower bound for $R' < R$, and replacing the upper bound at R with the least upper bound for $R' > R$. An alternative approach, which yields bounds which are slightly more conservative but smoother, is to use the rearrangement method advocated by Chernozhukov et al. (2009). In this setting, rearrangement is performed by simply sorting the estimated lower bounds and reassigning them to the original grid points, and similarly for the upper bounds. Chernozhukov et al. (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.[49]

Figures 5, 6, and 7 show the observed distribution of secret reserve prices, the estimated bounds, and the estimated distribution of seller types in the used-car dealers and fleet/lease sellers samples. Figure

[47] The measurement error here is written at the j level, rather than the (j,t) (period of the bargaining game) level. The resulting estimates are identical in either setting, but the former aids in identifying bargaining costs below and is consistent with the fact that no violations are observed *within* the same bargaining sequence j.

[48] As above, the bandwidth is given by $h = 1.06\hat{\sigma}n^{-\frac{1}{5}}$.

[49] After rearranging, I obtain conservative confidence bands from the outer limits of bootstrapped confidence intervals for the rearranged bounds.

5 demonstrates, in panels (c) and (d), that reserve prices are tightly centered around one for fleet/lease sellers. Used-car dealers, on the other hand, choose a wider variety of reserve prices which tend to be centered around 1.2, exceeding the average observable value of the car. Panels (a) and (b) demonstrate that only about 20% of used-car dealers report reserves lower than one, while about half of fleet/lease sellers do.

The estimated bounds on seller types are displayed in panels (a) and (b) of Figure 6 using $K = 50$ grid points. The estimates displayed are before applying rearrangement. For used-car dealers, in panel (a), each estimated lower bound lies below the estimated upper bound except at the far left tail, where few observations exist, as was demonstrated in Figure 5. In the fleet/lease sample, shown in panel (b) of Figure 6, the estimated lower bound lies above the estimated upper bound at 12 out of 50 grid points. However, as panel (b) shows, the violations are very small from an economic standpoint. Also, not shown in Figure 6 is the fact that nowhere does the 95% confidence interval for the lower bound lie completely above the confidence interval for the upper bound. Panels (c) and (d) display the 95% pointwise confidence bands after rearrangement.

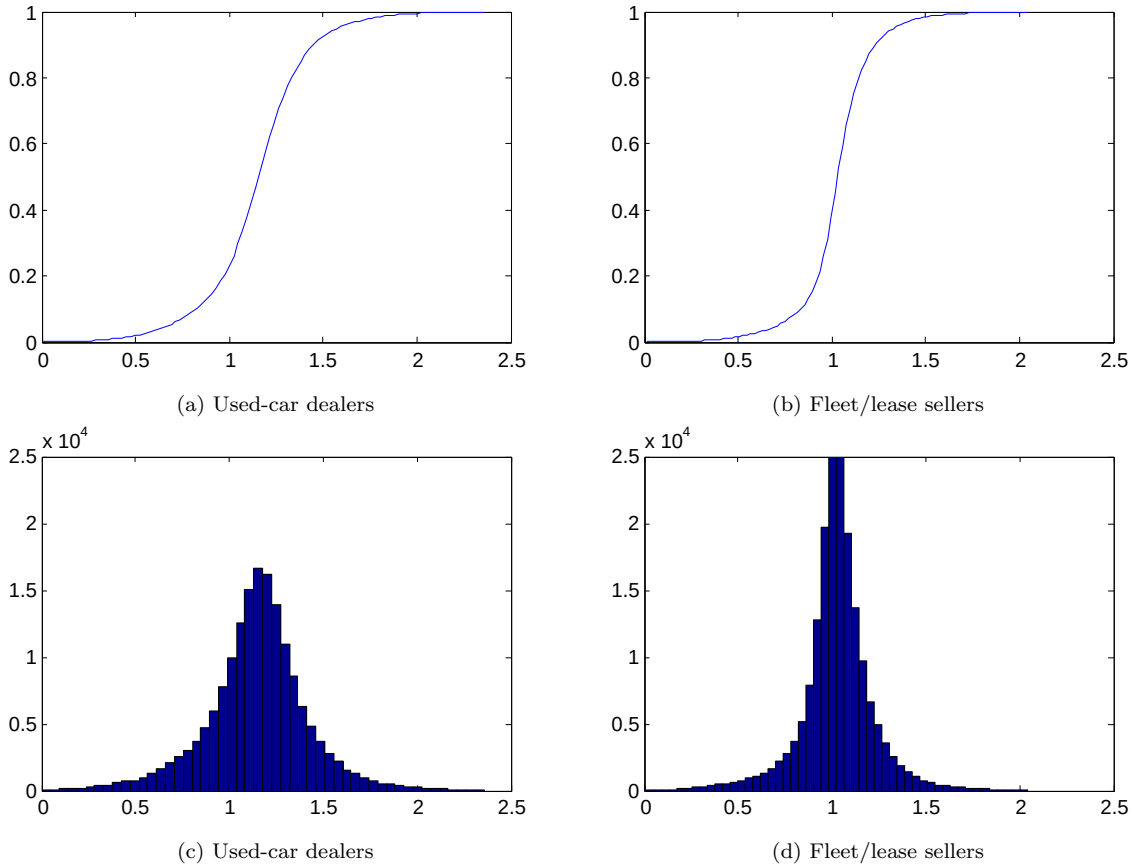
Figure 6 can be seen as evidence in favor of the model from Section 3. For example, in both samples the estimated bounds lie below the 45 degree line ($s(R) < R$), indicating that sellers never report secret reserve prices which are lower than the least they would be willing to sell the car for. Also, note that nowhere in the estimation of panels (a) and (b) of Figure 6 is the restriction imposed that $s(R)$ be strictly increasing, yet the estimates appear to bound a strictly increasing function.

One implication of panel (a) of Figure 6 is that when a used-car dealer sets a secret reserve price of 1.2 times the observable value of the car, she is really willing to let the car go at a price equal to the one. Similarly, in panel (b), when a fleet/lease seller sets a reserve price of one, she is actually willing to sell the car at a price of 0.9.

The estimated distribution and density of seller valuations, $\hat{s}$, are shown in Figure 7, using the midpoint of the bounds as the estimate of $\hat{s}$. Panels (c) and (d) demonstrate that much of the mass lies below one, which is to be expected: the existence of this market implies that sellers value these goods less than buyers (i.e., less than the high bidders). Panels (a) and (b) indicate that, for used-car dealers, about 60% of the time their valuation lies below one, while fleet/lease sellers have valuations less than one nearly 100% of the time.

5.4 Estimation of Bargaining Costs

This section presents an approach for estimating bounds on bargaining costs c_S, c_B, and $\delta = \delta_E \delta_I$. Note that δ_I and δ_E are not separately identified by the approach discussed here. These bounds may be estimated using observations in which one party chooses to counter. A necessary condition for a party to choose to counter is that the payoff from the player's opponent accepting with probability one must



(a) Used-car dealers

(b) Fleet/lease sellers

(c) Used-car dealers

(d) Fleet/lease sellers

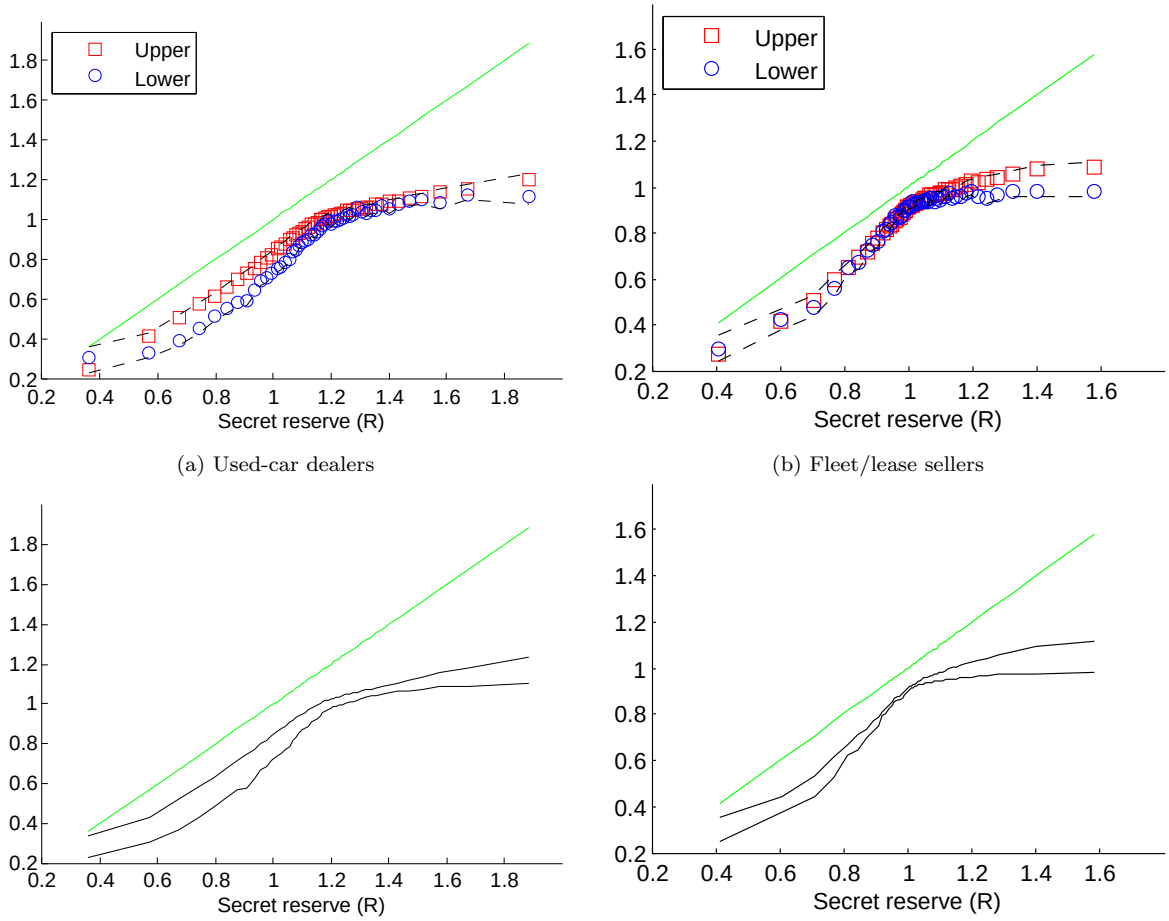
Figure 5: Empirical CDF and histogram of normalized secret reserve prices in the sample of cars sold by used-car dealers and the sample of cars sold by fleet/lease sellers. Units are in terms of the observable value of the car.

exceed the player's payoff from accepting the current offer on the table. That is,

$$\delta p_t^S - c_S \geq p_{t-1}^B$$
$$\delta(b - p_{t+1}^B) - c_B \geq b - p_t^S$$

Noting that $\delta < 1$ and $(c_S, c_B) > 0$, yields

$$\delta p_t^S \geq p_{t-1}^B$$
$$p_t^S - p_{t-1}^B \geq c_S$$
$$p_t^S - p_{t+1}^B \geq c_B$$



(a) Used-car dealers

(b) Fleet/lease sellers

(c) Used-car dealers

(d) Fleet/lease sellers

Figure 6: Bounds on sellers type. Panels (a) and (b) display estimated bounds on seller types in used-car dealers sample and fleet/lease sellers sample. Pointwise 95% confidence bands shown by dashed lines. Panels (c) and (d) display 95% confidence bands alone after applying rearrangement. Units are in terms of the observable value of the car.

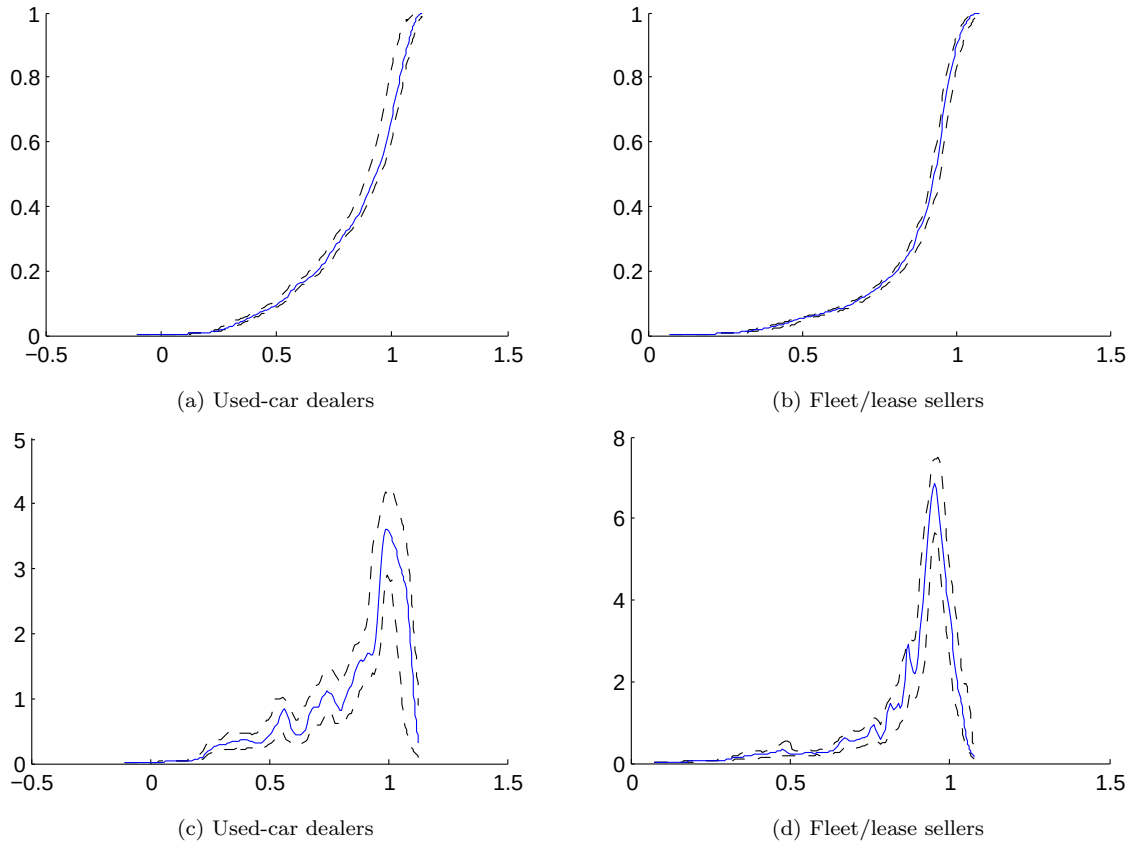
Bounds for additive bargaining costs can then be estimated by

$$0 \leq c_S \leq \min_{\{(j,t):d^S_{j,t}=C\}} \left(p^S_{j,t} - p^B_{j,t-1}\right)$$

$$0 \leq c_B \leq \min_{\{(j,t):d^B_{j,t}=C\}} \left(p^S_{j,t} - p^B_{j,t+1}\right)$$

As above, recorded bargaining offers are considered to be measured with error ν_j, which drops out of the expression for bounds on additive bargaining costs, but not the expression of δ bounds.[50] Therefore,

[50] As noted in Section 5.3, the estimated bounds on $s(R)$ are identical if measurement error is considered to vary with t. The estimated upper bounds on c_S and c_B are high and uninformative in this case, however. An empirical finding



(a) Used-car dealers

(b) Fleet/lease sellers

(c) Used-car dealers

(d) Fleet/lease sellers

Figure 7: Estimated CDF and density of seller valuations in used-car dealers sample and fleet/lease sellers sample. Pointwise 95% confidence bands shown by dashed lines. Units are in terms of the observable value of the car.

bounds on δ are given by taking expectations, and these bounds are made tighter by taking expectations in different samples of data, for example, by conditioning on R:

$$1 \geq \delta \geq \sup_{R} \frac{E[p^B_{t-1}|R]}{E[p^S_t|R]}$$

This can be implemented as follows. Divide the support of R into K bins, as in Section 5.3. Let $\mathcal{R}^S_k = \{(j,t) : j \in \mathcal{R}_k, d^S_{j,t} = C\}$ represent the set of observed bargaining periods, (j,t) in which the seller chose to counter and in which the reserve price was in bin k. Then bounds on δ are given by

$$1 \geq \delta \geq \max_{k} \frac{\frac{1}{\#(\mathcal{R}^S_k)} \sum_{(j,t)\in\mathcal{R}^S_k} p^S_{j,t}}{\frac{1}{\#(\mathcal{R}^S_k)} \sum_{(j,t)\in\mathcal{R}^S_k} p^B_{j,t-1}}$$

supporting the assumption that measurement error lies at the j level is that no violations of the bounds in Section 5.3 are observed within a given auction/bargaining sequence j.

I use a grid of $K = 1,000$ points. I then obtain a bias-corrected, one-sided confidence interval for each of these estimated bounds using recent methods for moment inequality inference developed in Chernozhukov et al. (2011), which accounts for the fact that the estimated bound for δ will be biased upward given that it is derived from taking a maximum, and vice-verse for c_S and c_B.[51]

Table 3 displays the bargaining cost bound estimates, with one-sided 95% confidence estimates reported in parentheses. The estimated lower bound on δ is 0.9365 in the sample of cars sold by used-car dealers and 0.9413 in the sample of cars sold by fleet/lease sellers. These values are close to one, as expected. However, because the empirical probability of exogenous breakdown in bargaining is 1–3%, naively using these bounds as estimates of δ would imply a larger discount factor due to impatience than is realistic, given that the time between bargaining rounds is only 2–3 hours. Therefore, in the analysis below I choose the estimate of δ to be the midpoint between one and the estimated bound on δ.

The upper bounds on the per-period costs of bargaining are relatively close to zero. The 95% confidence bound for c_S is 0.0043 in the used-car sample and 0.0051 in the fleet/lease car sample. The bound on c_B is slightly tighter than on c_S in the used-car sample. Recalling that used-car dealers are both buyers and sellers in this sample, this result suggests that buyers may dislike bargainging less when they are on the buy side than on the sell side.

To make the comparison more concrete, recall from Table 1 that the average final price, in nominal terms, is \$6,346 for cars sold by used-car dealers and \$9,987 for fleet/lease sellers. Therefore, bounds of 0.0043 and 0.0051 for for c_S corresponds to an upper bound of about \$28 for a used-car seller's disutility of making bargaining offers, and \$51 for a fleet/lease seller. For buyers, the corresponding nominal upper bounds for the disutility of bargaining would be \$15 and \$52. This suggests that both parties are more willing to bargain over cars sold by used-car dealers than by fleet/lease sellers, consistent with the reduced-form statistics of Table 2, which demonstrated that bargaining occurs more often in the used-car dealers sample. Also, as part of the cost embedded in c_S and c_B consists of the time spent on the phone with the auction house and off the phone determining a new counteroffer, a value of \$15, for example, seems reasonable relative to what might be the hourly wage of a used-car dealer. In the estimation of expected gains in the dynamic mechanism below, I plug in the upper 95% bounds as my estimates of additive costs, but as the costs are quite small the results are not sensitive to using other values within the estimated range.

6 The Pareto Frontier, Counterfactual Mechanisms, and the Dynamic Mechanism

This section describes how the Pareto frontier and each counterfactual mechanism can be solved for once the distributions of seller and buyer valuations are known. This section also details the quantile-matching approach I use to solve for the direct mechanism corresponding to the dynamic mechanism. The

[51] Chernozhukov et al. (2011) refer to estimators corrected in this fashion as being half-median unbiased.

Table 3: Estimated Bargaining Costs

	Used-car dealers	Fleet/lease sellers
$\delta \geq$	0.9365	0.9413
	(0.8852)	(0.8973)
$c_S \leq$	0.0037	0.0029
	(0.0043)	(0.0051)
$c_B \leq$	0.0005	0.0030
	(0.0023)	(0.0052)

Notes: Estimated bounds on the costs of bargaining. The parameter δ represents discounting due to impatience as well as the probability that bargaining does not exogenously breakdown. The parameters c_S and c_B represent the seller's and buyer's per-period disutility of bargaining in terms of the observable value of the car. One-sided 95% confidence bounds are reported in parentheses.

counterfactuals hold the distribution of buyer and seller types fixed. In reality, changing the mechanism could change the distribution of types.[52]

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 identifies the high bidder, and the lower bound for the buyer's support in the post-auction bargaining game becomes p^A, the high bid at the auction.[53] The post-auction bargaining game between the high bidder and seller can then be mapped into the direct revelation mechanism framework provided by Myerson and Satterthwaite (1983) and Williams (1987), or an indirect mechanism framework, that of a k double auction, provided by Chatterjee and Samuelson (1983) and Satterthwaite and Williams (1989).

6.1 Direct Bargaining Mechanisms

In this section I discuss how I solve for the Pareto frontier and other direct, efficient mechanisms. 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

[52]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.

[53]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.

the seller.[54] 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 $\overline{s}$.[55] The ex-ante expected utility of the buyer and seller in a mechanism (x, q) is given by

$$\overline{U}_s(x,q) = q + \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} x(s,b) F_s(s) \frac{f_b(b)}{1 - F_b(p^A)} ds db \tag{6}$$

$$\overline{U}_b(x,q) = G(x) - q + \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} x(s,b) \frac{(1 - F_b(b))}{1 - F_b(p^A)} f_s(s) ds db \tag{7}$$

where

$$G(x) = \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(\phi_b(b) - \phi_s(s)\right) x(s,b) f_s(s) \frac{f_b(b)}{1 - F_b(p^A)} ds db \tag{8}$$

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

$$\gamma \overline{U}_s + (1 - \gamma) \overline{U}_b \tag{9}$$

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

$$x^{\alpha_1(\gamma),\alpha_2(\gamma)}(s,b) = 1\left\{\phi_b(b,\alpha_2(\gamma)) \geq \phi_s(s,\alpha_1(\gamma))\right\}$$

The parameters $(\alpha_1(\gamma), \alpha_2(\gamma))$ can be solved for at each γ using an approach developed in Williams (1987) and described in Appendix C.4. Intuitively, the approach maximizes (9) subject to $G(x^{\alpha_1(\gamma),\alpha_2(\gamma)}) \geq 0$, where $G(x)$ is defined in (8). 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 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 et al. (2009).[56]

[54] 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.

[55] The transfer function, p, is not essential for the results here, but can be found in Appendix C.4.

[56] One might be tempted to impose monotonicity of $\phi_s(s)$ and $\phi_b(b)$ simply by adjusting the kernel bandwidths used in calculating the corresponding densities and distributions. Henderson et al. (2011) demonstrated that this approach is not valid, and proposed a constraint-weighted bootstrapping approach to imposing monotonicity. I find that their approach yields similar estimates to the rearrangement approach, but rearrangement is more efficient computationally given the size of my dataset.

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 C.3. Overall, the rearrangement has no effect on the densities and distributions other than smoothing out small deviations from monotonicity.

Several mechanisms of interest fit into this framework, such as

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): $\gamma = 1/2$, $\alpha_1 = \alpha_2 = \alpha^*$, where α^* solves $G(x^{\alpha^*,\alpha^*}) = 0$.
3. Seller-optimal: $\gamma = 1$, $\alpha_1 = 0$, $\alpha_2 = 1$.
4. Buyer-optimal: $\gamma = 0$, $\alpha_1 = 1$, $\alpha_2 = 0$.
5. Pareto frontier: mechanisms maximizing (9) subject to $G(x^{\alpha_1(\gamma),\alpha_2(\gamma)}) \geq 0$ for $\gamma \in [0,1]$.

Note that an auction followed by the seller-optimal mechanism is equivalent to a public reserve auction.[57] 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 D.

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).[58] 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 probability of trade in any of these direct mechanisms is given by

$$\int_{\underline{b}}^{\overline{b}} \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} x(s,b;p^A)\hat{f}_s(s)\frac{\hat{f}_b(b)}{1-\hat{F}_b(p^A)}\hat{f}_{p^A}(p^A)dsdbdp^A \tag{10}$$

To perform this integration, I use Gauss-Chebyshev quadrature, as described in Appendix C.2, with 200 nodes in the s and b dimensions, and 25 nodes in the p^A dimension.[59]

[57] See Menezes and Ryan (2005).

[58] Note that the expected transfer functions for the mechanism in Corollary 1 are given by (17) and (18) in Appendix C.4

[59] Increasing the number of nodes in either dimension did not change the results. A greater degree of accuracy in the b and s dimensions than in the p^A dimension is useful, however, as each mechanism is solved conditional on p^A.

6.2 Indirect Mechanisms: The k Double Auction

This section discusses the estimation of equilibria in the k double auction, an alternative to dynamic bargaining which would be simple for the auction house and players to understand and implement. A k double auction, as in Chatterjee and Samuelson (1983) and Satterthwaite and Williams (1989), is an indirect mechanism in which both the seller and buyer submit sealed offers, p^S and p^B, to the auctioneer. If $p^B \geq p^S$, trade occurs at price

$$p = kp^S + (1-k)p^B$$

A k double auction with $k = 1$ corresponds to the seller-optimal direct mechanism (or the public reserve auction) and a k double auction with $k = 0$ corresponds to the buyer-optimal direct mechanism. Satterthwaite and Williams (1989) demonstrated that, for $k = (0, 1)$, a continuum of equilibria exist satisfying the following linked differential equations:

$$p^{B(-1)}(p^S(s)) = p^S(s) + kp^{S'}(s)\phi_s(s,1) \tag{11}$$

$$p^{S(-1)}(p^B(b)) = p^B(b) + (1-k)p^{B'}(b)\phi_b(b,1) \tag{12}$$

where $p^{B(-1)}(\cdot)$ and $p^{S(-1)}(\cdot)$ are the inverses of the buyer's and seller's strategies.

Solving the k double auction is somewhat more involved than estimating the Pareto frontier. Satterthwaite and Williams (1989) provided an approach for solving for equilibria numerically. A point (s, b, p) is chosen in the set $\mathcal{P} = \{(s, b, p) : \underline{s} \leq s \leq p \leq b \leq \overline{b}, s \leq \overline{s}, b \geq \underline{b}\}$, and then a one-dimensional manifold passing through this point is traced out using differential equations defined by (11) and (12). This path traces out an equilibrium. See Appendix C.5 for details. As with the solution method for the Pareto frontier, this approach relies on $\phi_b(b)$ and $\phi_s(s)$ being weakly increasing.

Given that a continuum of equilibria exist, and given that the k double auction is not a supermodular game, summarizing possible outcomes in this game is not simple.[60] However, the purpose of examining the $k = 1/2$ double auction in the context of this paper is to compare the currently used, *dynamic* mechanism, which has easy rules for players to understand and for the auction house to implement, to a *static* mechanism which could also be easily understood and implemented by the auction house. Therefore, in order to obtain some idea of possible outcomes under the $k = 1/2$ double auction given the estimated market conditions of buyer and seller valuations, I choose a grid of approximately 100 uniformly spaced points in $\mathcal{P}$ and solve for the equilibria passing through these points. I then rank these

[60] To see that the k double auction is not a supermodular game, let $\pi_b^{DA}(p^B, p^S)$ represent the buyer's payoff, which is 0 if $p^S < p^B$ and $b - kp^B - (1-k)p^S$ otherwise. Let Δ be a small, positive constant. Then write the following second difference

$$\begin{aligned}&\pi_b^{DA}(p+\Delta, p-2\Delta) - \pi_b^{DA}(p-\Delta, p-2\Delta) - \left[\pi_b^{DA}(p+\Delta, p) - \pi_b^{DA}(p-\Delta, p)\right] \\ &= -b + p - k\Delta\end{aligned}$$

The final term can be positive or negative depending on b and p. Even if one requires $p \leq b$, other examples exist showing ambiguity of the sign of the second difference. Therefore, the game is not supermodular or submodular.

equilibria based on the total gains from trade achieved, and report the equilibria corresponding to the minimum, maximum, and quartiles of this ranking.[61] An example of a solution path crossing through a point in $\mathcal{P}$ is shown in Figure 8.

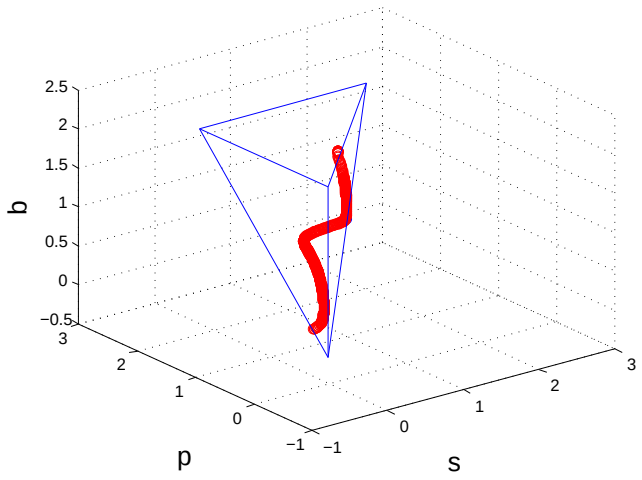


Figure 8: A solution to the k=1/2 double auction, lying with the tetrahedron $\mathcal{P} = \{(s, b, p) : \underline{s} \leq s \leq p \leq b \leq \overline{b}, s \leq \overline{s}, b \geq \underline{b}\}$. Units are in terms of the observable value of the car.

The expected gains from trade for the buyer in the $k = 1/2$ double auction can be evaluated by

$$\int_{\underline{b}}^{\overline{b}} \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(b - \frac{kp^B(b) + (1-k)p^S(s)}{2} \right) 1\{p^S(s) \leq p^B(b)\} \hat{f}_s^*(s) \frac{\hat{f}_b^*(b)}{1 - \hat{F}_b^*(p^A)} \hat{f}_{p^A}(p^A) ds db dp^A \tag{13}$$

where $p^B(b)$ and $p^S(s)$ are equilibrium strategies solved for using the approach in Appendix C.5. The seller's expected gains from trade and the probability of trade are calculated similarly.

6.3 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, the direct mechanism equivalent to the dynamic bargaining mechanism currently used at wholesale auto auctions can be characterized by functions x and p determining whether or not trade will occur and at what price. However, in the dynamic mechanism, these functions will vary by the period of the bargaining game. As with the above mechanisms, x and p will also in general depend on the high bid at the auction, p^A, as this is the lower bound of the support of buyer types when bargaining takes place.

[61] Due to numerical approximation error in the differential equations techniques described in Appendix C.5 and the approach in solving for the Pareto frontier in Appendix C.4, equilibria in the k double auction can be estimated to lie outside the Pareto frontier. Therefore, I eliminate these equilibria and rank the equilibria which lie within the Pareto frontier.

It is important to note that the estimation procedures in Section 5 yields an estimate of bounds on $\hat{s}_j$, the seller's value for car j. I then select a point estimate by using the lower bound, upper bound, midpoint of the 95% confidence band. Section 5 does not, however, yield an estimate of the buyer's value. If the buyer's value could be estimated, the direct mechanism (x, p) would be trivial to estimate. However, given that only the *distribution* of buyer values is identified, a more sophisticated approach is required.

Let the allocation function in the dynamic mechanism be written $x_\tau^D(s, b; p^A)$, representing the probability that trade occurs in bargaining period τ given that the game has lasted up until the τ^{th} period and the players' types are s and b. I assume this function is an indicator function, as with the theoretical mechanisms in Section 6.1.[62] The function $x_\tau^D(s, b; p^A)$, can be recovered nonparametrically using the result that $s(R)$ is one-to-one. Let

$$m_\tau(R, p^A) = \Pr(x_\tau^D(s(R), b; p^A) = 1 | R, p^A, \tau)$$

The object $m_\tau(R, p^A)$ can be estimated by the empirical probability of trade in period τ when the reported secret reserve was R and the high bid was p^A and the game had lasted at least to the beginning of the period τ. I estimate this function using a linear regression of a dummy for trade on a tensor product of cubic b-splines with 10 uniformly spaced knots in each dimension.[63] Then, because $x_\tau^D(s(R), b; p^A)$ should be weakly increasing in b given R and p^A, $x_\tau^D(s(R), b; p^A)$ can be expressed as

$$x_\tau^D(s(R), b; p^A) = 1\{b \geq g_\tau(R, p^A)\}$$

where $g_\tau(R, p^A)$ can be solved for numerically as the $(1 - m_\tau(R, p^A))$ quantile of $\frac{F_b(b) - F_b(p^A)}{1 - F_b(p^A)}$, the distribution of the high valuation buyer conditional on the auction high bid.[64] , the fact that bargaining occurs Because s(R) is one-to-one, $g_\tau(R, p^A)$ may also be written $g_\tau(s, p^A)$, and, for each p^A, delineates the boundary of trade in (s, b) space.

With the objects $m_\tau(R, p^A)$ and $g_\tau(s, p^A)$ in hand, the ex-ante expected surplus for the buyer and seller in the dynamic mechanism, $\overline{U}^{D,B}$ and $\overline{U}^{D,S}$, can be calculated as follows. Let $\mathcal{J}_\tau$ be the set of all auction/bargaining sequences which ended in period τ. Let $p_j^{D,B}$ and $p_j^{D,S}$ be the final price paid by the buyer and received by the seller in observation j. These prices are both zero if no trade occurs, and they

[62] More generally, the function $x_\tau^D(s, b; p^A)$ could take on values on the interior of $(0, 1)$, but in this case the function would not be identified using the approach described here.

[63] B-splines are finite-dimensional, piece-wise polynomials. See Chen (2007) for an explanation of b-splines and other sieve spaces. Note that this linear regression can predict values lying outside of $[0, 1]$. In these cases, I replace the estimated $m_\tau(R, p^A)$ with 0 or 1. Overall, this truncation is only necessary in a few cases and has little effect on the estimated $x_\tau^D(s(R), b; p^A)$.

[64] Note that the lower bound of the buyer's support in the currently used dynamic mechanism will not necessarily be exactly p^A, because the seller also learns from the fact that the buyer must have chosen not to opt out of bargaining. The actual lower bound will be $\chi^{-1}(p^A)$, defined in the proof of Proposition 2. However, $\chi^{-1}(p^A)$ is not identified without imposing and solving for the equilibrium of the bargaining game. Therefore, I use p^A in place of $\chi^{-1}(p^A)$.

are not equivalent to each other because both buyer and seller pay auction house fees. Then

$$\overline{U}^{D,B} = \sum_{\tau=1}^{T} \delta^{\tau-1} \left(\left[\frac{1}{J} \sum_{j \in \mathcal{J}_t, t \geq \tau} \int_{g_\tau(s_j, p_j^A)}^{\overline{b}} b f_b(b) db \right] - \left[\frac{1}{J} \sum_{j \in \mathcal{J}_\tau} p_j^{D,B} \right] - \left[c_B \frac{\tau}{2} 1\{\tau \text{ even}\} \frac{\#(\mathcal{J}_\tau)}{J} \right] \right) \quad (14)$$

$$\overline{U}^{D,S} = \sum_{\tau=1}^{T} \delta^{\tau-1} \left(\left[\frac{1}{J} \sum_{j \in \mathcal{J}_\tau} p_j^{D,S} \right] - \left[\frac{1}{J} \sum_{j \in \mathcal{J}_t, t \geq \tau} s_j m_\tau(R_j, p_j^A) \right] - \left[c_S \frac{\tau - 1}{2} 1\{\tau \text{ odd}, \tau > 1\} \frac{\#(\mathcal{J}_\tau)}{J} \right] \right) \quad (15)$$

In (14), the expression multiplying δ inside the first summation consists of three pieces. The first term is the expected gross payoff (ignoring price) to the buyer from trade occurring in period τ. The second term is the expected price paid by the buyer for trade occurring in period τ. The final term is the expected additive bargaining cost paid by the buyer in a game which ended in period τ. Recall that the buyer pays this cost whenever he makes a counteroffer, so if the game ends at an even period τ (the seller's turn), the buyer will have paid this cost $\tau/2$ times.

Similarly, in (15), the expression multiplying δ inside the first summation is divided into three parts. First is the expected price received by the seller for trade occurring in period τ. Second is the forgone outside option for the seller for trade occurring in period τ. Third is the expected additive bargaining cost paid by the seller for a game ending in period τ. The seller pays this cost whenever she makes a counteroffer, so if the game ends at an odd period $\tau > 1$ (the buyer's turn is at odd periods), then the seller will have paid this cost $(\tau - 1)/2$ times.

7 Putting It All Together: How Efficient Is Bargaining?

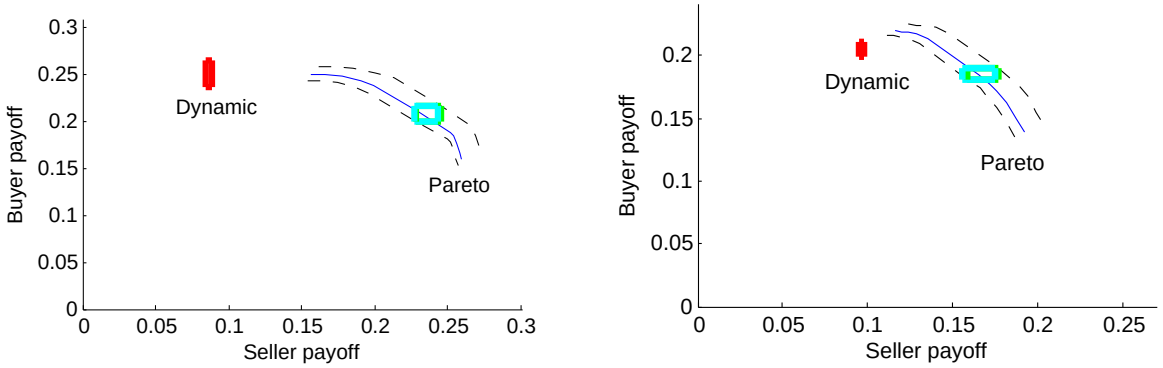
This section puts together all of the pieces from estimation (buyer and seller distributions and bargaining costs) and the mechanism design objects solved for (the Pareto frontier, other counterfactual mechanisms, and surplus in the current mechanism) to determine how efficient the currently used dynamic mechanism is. I first study the comparison to the direct mechanisms described above, followed by a comparison to the indirect, $k = 1/2$ double auction. Finally, I present an analysis of the region of the type space where trade occurs. For the fleet/lease sample, the analysis below uses the midpoint of the 95% confidence bounds for the seller's secret reserve price as the estimate of $\hat{s}(R)$. For the used-car dealers sample, I instead use the lower confidence bound, as this resulted in a better model fit, as discussed in Section 8.

7.1 Efficiency of Bargaining Relative to Pareto Frontier and Direct Mechanisms

In this section I present a qualitative and then quantitative comparison of dynamic bargaining to the Pareto frontier and other direct mechanisms. The performance of the currently used dynamic mechanism relative to the Pareto frontier is displayed in Figure 9. Each panel of Figure 9 also displays the location of the second-best and trade-maximizing mechanisms with respect to the frontier, and for each mechanism

the 95% confidence interval is displayed. In both samples, the second-best and trade-maximizing mechanisms are indistinguishable from one another on the frontier. Panel (a) demonstrates that, in sales of cars owned by used-car dealers, the auction followed by alternating-offer bargaining results in a surplus level which lies close to the Pareto frontier, and for fleet/lease sellers the mechanism lies even closer, indicating that bargaining is quite efficient.

Any gap between the frontier and the dynamic mechanism would represent a deadweight loss—neither bargaining party captures this lost surplus. I refer to this gap as the deadweight loss due to limited commitment, as this gap could theoretically be eliminated by committing to a static, direct mechanism which lies on the frontier. Figure 9 suggests that this deadweight loss is small, but still positive. In particular, the seller payoff would be higher along the Pareto frontier than it is in the currently used mechanism.



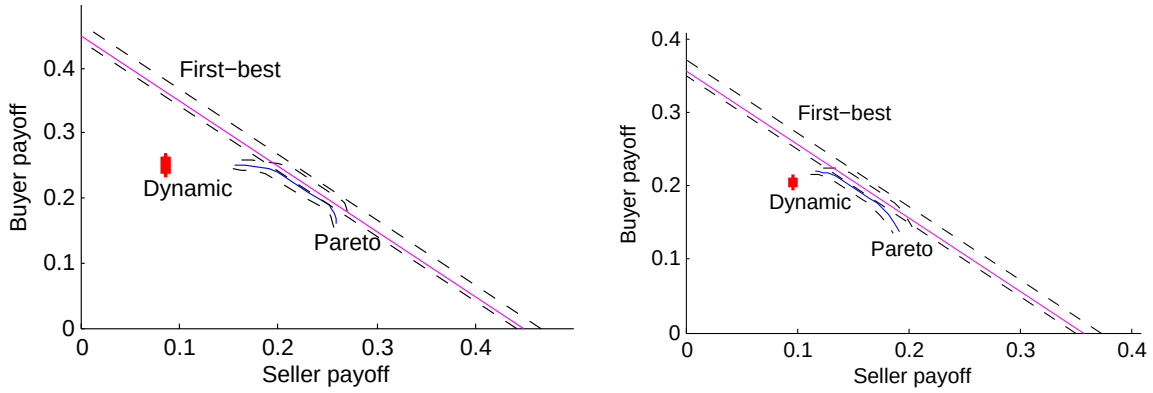
(a) Used-car dealers

(b) Fleet/lease sellers

Figure 9: Estimated performance of dynamic mechanism relative to Pareto frontier and other direct mechanisms. The dynamic mechanism is marked in red. The second-best (marked in light blue) and trade-maximizing (marked in light green) lie close to one another on the frontier and are indistinguishable in the figure. For each mechanism, the pointwise 95% confidence region is displayed. The Pareto frontier is marked in blue, with pointwise 95% confidence bands marked by dashed lines. Units are in terms of the observable value of the car.

An additional source of deadweight loss is incomplete information. Myerson and Satterthwaite (1983) demonstrated that this deadweight loss is inevitable in bilateral bargaining with two-sided uncertainty. Intuitively, both parties trade off the conflicting incentives of increasing the probability of trade and increasing the rent extracted from the other party, leading buyers to shade their offering price downward and sellers to shade their asking price upward, leading some trades to fail to be consummated even when the buyer values the good more than the seller. This deadweight loss is depicted graphically in Figure 10 as the gap between the Pareto frontier and the first-best (infeasible) surplus line. In both the used-car dealer sample and the fleet/lease sample, the gap is small. In fact, the estimated Pareto frontier lies within the 95% confidence band of the first-best line, indicating that incomplete information per se is not

the primary reason for inefficiency in this market.



(a) Used-car dealers

(b) Fleet/lease sellers

Figure 10: Estimated performance of dynamic mechanism relative to Pareto frontier and first-best (infeasible) surplus line. The dynamic mechanism is marked in red. The first-best line is marked in purple and the Pareto frontier is marked in blue, each with pointwise 95% confidence bands marked by dashed lines. Units are in terms of the observable value of the car.

Table 4 analyzes these mechanisms in more detail, displaying the total expected gains from trade, the gains to each player, and the probability of sale under several static direct mechanisms. The current dynamic mechanism is displayed in the final column. Because buyer and seller gains take into account fees paid to the auction house, the total expected gains from trade for the dynamic mechanism in the first row is given by the sum of buyer and seller gains and fees paid to the auction house. All outcomes are in terms of the wholesale observable value, or market value, of the vehicle.

The final column of Table 4 shows that, for cars sold by used-car dealers (panel A), the expected gains from trade in the current mechanism is 38.8% of the market value of the vehicle, with 25.6% going to the buyer, 8.7% going to the seller, and the remainder going to the auction house through fees. For buyers, 25.6% can be interpreted as an expected retail markup, as a buyer's valuation should be the amount he expects to sell the car for at his own lot. For a seller, 8.7% is the expected gain from selling the car at the auction today as opposed to selling it through her next-best option. This next-best option includes the possibility of taking the car back to the seller's own lot or leaving it at the auction house to run again the following week. Therefore, the low expected gains from trade for the seller imply that cars are being sold for close to what the seller would be able to get for the vehicle elsewhere.

In the fleet/lease sample, shown in panel B of Table 4, the expected gains from trade is 33.4% of the market value of the car, with 20% going to the buyer and 9.6% going to the seller. From conversations with used car dealers and sellers, as well as auctioneers and auction house employees, these markups for buyers and sellers in both the used-car dealer sample and the fleet/lease sample are in line with reality: A buyer expects to obtain a retail markup ranging from about 10-25%, and and sellers expects to make a smaller gain relative to her next best option, ranging from about 0-10%. The markups are therefore

Table 4: Performance of bargaining relative to direct mechanisms

A. Used-car dealers sample

	First-best	Second-best	Buyer-optimal	Seller-optimal	Trade-maximizing	Dynamic mechanism
Expected gains from trade	0.449 (0.006)	0.440 (0.006)	0.408 (0.006)	0.419 (0.006)	0.436 (0.006)	0.388 (0.007)
Buyer gains		0.210 (0.005)	0.251 (0.004)	0.159 (0.003)	0.209 (0.005)	0.256 (0.007)
Seller gains		0.231 (0.005)	0.157 (0.002)	0.260 (0.004)	0.230 (0.005)	0.087 (0.001)
Probability of trade	0.893 (0.003)	0.773 (0.005)	0.620 (0.007)	0.662 (0.006)	0.777 (0.005)	0.715 (0.002)

B. Fleet/lease sellers sample

	First-best	Second-best	Buyer-optimal	Seller-optimal	Trade-maximizing	Dynamic mechanism
Expected gains from trade	0.358 (0.005)	0.349 (0.005)	0.335 (0.005)	0.330 (0.006)	0.345 (0.005)	0.334 (0.003)
Buyer gains		0.184 (0.002)	0.219 (0.002)	0.139 (0.002)	0.185 (0.002)	0.203 (0.003)
Seller gains		0.165 (0.005)	0.116 (0.003)	0.191 (0.004)	0.162 (0.005)	0.096 (0.001)
Probability of trade	0.907 (0.003)	0.773 (0.006)	0.705 (0.006)	0.654 (0.007)	0.780 (0.006)	0.774 (0.002)

Notes: Direct, static mechanisms compared to current dynamic mechanism based on estimated gains and probability of trade. Standard errors are from 200 bootstrap replications. Units (other than for probability of trade) are in terms of the observable value of the car.

slightly high for buyers, but still within a reasonable range.

As was shown graphically in Figure 9, Table 4 demonstrates quantitatively that the second-best mechanism achieves nearly the first-best level of surplus. In the used-car dealers sample, the second-best surplus is 0.440 of the market value relative to 0.449 of the market value in the first-best, indicating that the inefficiency due to incomplete information in this market is less than one percent of the market value

of the car. Alternatively, the second-best achieves 98% of the first-best surplus level. In the fleet/lease sample, the second-best also achieves about 98% of the first-best surplus.

Comparing the buyer-optimal to the dynamic mechanism gives an idea of what could occur if the auction house were to commit to an auction followed by a take-it-or-leave-it offer from buyer to seller. Buyer gains from trade in this case would be approximately the same, about 25% of the market value of the car. Seller gains, on the other hand, increase from 8.7% to 15.7% of the market value in the user-car dealers sample and increase from 9.6% to 11.6% of the market value in the fleet/lease sample.

The seller-optimal mechanism is equivalent to a public reserve auction in this setting. Panel A of Table 4 implies that seller gains could be much higher under a public reserve auction than in the currently used dynamic mechanism, moving from 0.087 to 0.260 of the market value of the car. Buyer gains, on the other hand, decrease to 0.159 of the market value. This may be one reason for the use of this post-auction bargaining: while a public reserve auction would be optimal for the seller, post-auction bargaining in this market appears to improve the buyer's payoff relative to a public reserve auction. Panel A demonstrates that the total gains from trade in the used-car dealers sample would be slightly higher in a public reserve auction than in the current dynamic mechanism, while the opposite is shown in panel B for fleet/lease sellers. It is important to note that each of these counterfactual comparisons holds fixed the distribution of buyers and sellers who choose to participate. In reality, if the auction house were to implement a public reserve auction, some buyers would likely choose not to participate.

The deadweight loss due to limited commitment can be evaluated by comparing the expected gains from trade in final column of Table 4 to the second-best column. In the used-car dealers sample, the dynamic mechanism yields gains of 0.388 of the market value relative to 0.440 in the second-best, meaning the dynamic mechanism captures 88% of the second-best surplus. Therefore, the deadweight loss due to limited commitment is small, but still larger than the deadweight loss due to incomplete information. Comparing the dynamic mechanism to the first-best demonstrates that the dynamic mechanism captures 86% of the first-best surplus in the used-car dealers sample.

In the fleet/lease sample, the total deadweight loss due to limited commitment in the dynamic mechanism is also small: the dynamic mechanism yields total gains of 0.334 of vehicle market value and the second-best mechanism would yield gains of 0.349 of market value. Therefore, the dynamic mechanism captures 96% of the second-best gains from trade. A comparison with the first-best demonstrates that the dynamic mechanism captures 93% of the first-best surplus.

The costs of bargaining account for only a small amount of lost surplus in this setting. To measure this, I evaluate (14) and (15) with no costs of bargaining ($\delta = 1$, $c_S = 0$, $c_B = 0$). I find that buyer and seller surplus increase by at most 0.005 of the market value of the car in either sample.

The final row of each panel in Table 4 compares the probability of trade in each mechanism. The probability of trade at wholesale auto auctions is used by industry analysts and individual auction houses as a performance measure, and improving the probability of trade is a primary goal of auction houses.[65] Table 4 provides insight into how this number might change if players were to commit to a static mech-

[65]See Sawyers (2012) for a recent industry discussion of the probability of trade, also referred to as the conversion ratio.

anism. In sales by used-car dealers, trade occurs in the currently used mechanism 71.5% of the time. Under the trade-maximizing mechanism, this sales rate could increase to 77.7%. Thus, the dynamic mechanism captures 92% of the trade volume achievable under the trade-maximizing mechanism. With the buyer-optimal or seller-optimal mechanisms, the volume of trade would decrease to 62-66%.

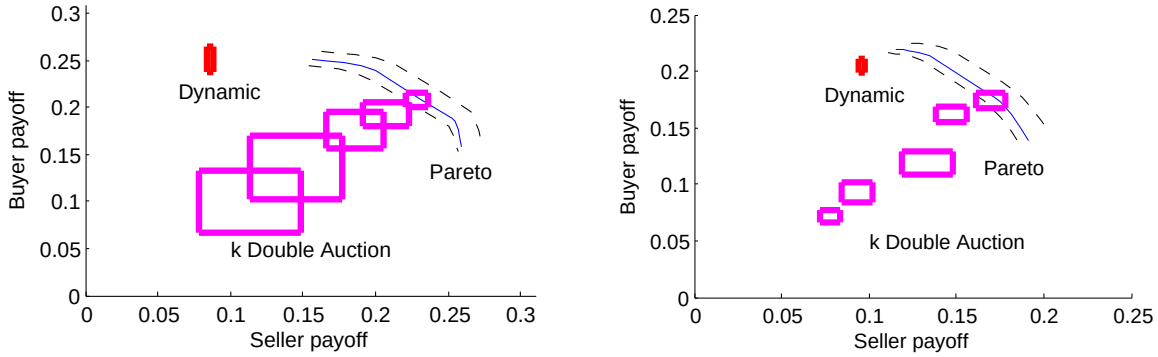
In the fleet/lease sample, trade occurs 77.4% of the time in the dynamic mechanism and 78% of the time in the trade-maximizing mechanism, implying that the dynamic mechanism achieves 99% of the consummated sales which are possible under two-sided uncertainty. The level of trade is higher in the dynamic mechanism than in the seller-optimal/public reserve auction mechanism (65.4%). This again speaks to one reason why post-auction bargaining may be used: it can potentially lead to a higher volume of trade than a public reserve auction, which can benefit the auction house.

Recall that the first-best would have trade occur whenever the buyer's value is higher than the seller's. Therefore, the first-best probability of trade will not be one, but will depend on the distributions of buyer and seller valuations. The first-best column demonstrates that in the absence of incomplete information, 89.3% of buyer and seller matches should result in trade in the used-car dealers sample and 90.7% in the fleet/lease sample. The second-best volume of trade is about 12-13 percentage points lower. Given that the second-best surplus level is quite close to the first-best, this result implies that the trades which fail to occur due to incomplete information are trades which would have generated little surplus—that is, they are cases where the buyer only values the good slightly more than the seller. The trade-maximizing mechanism yields only a slightly higher volume of trade than the second-best mechanism, increasing by less than one percentage point in both samples.

7.2 Efficiency of Bargaining Relative to a $k = 1/2$ Double Auction

As mentioned in the introduction, one reason alternating-offer bargaining is such a common form of trade may be its simplicity. It is easy for players to understand and easy for the auction house to implement, as it does not require the auction house to know the distributions of buyer and seller valuations. On the other hand, the second-best direct mechanism is unlikely to actually be implemented in practice, as it would require, first, that players understand the mechanism and recognize that truth-telling is optimal, and second, that the auction house have knowledge of buyer and seller distributions in order to implement the mechanism (i.e., the auction house would need to be able to evaluate the allocation and transfer functions, (x, p), described above). It is therefore useful to compare the currently used mechanism to an alternative mechanism which is similarly simple to understand and implement: a $k = 1/2$ double auction. However, unlike alternating-offer bargaining, a k double auction requires commitment, just as with the direct mechanisms examined above. Therefore, comparing outcomes under the dynamic mechanism yields an estimate of what might be achieved if the auction house and players could commit to a alternative, simple trading mechanism.

The results of this comparison are displayed in Figure 11 under different possible equilibria for the $k = 1/2$ double auction. I rank double auction outcomes by the total gains from trade and select equilibria corresponding to the quartiles of these outcomes as well as the minimum and maximum.



(a) Used-car dealers

(b) Fleet/lease sellers

Figure 11: Estimated performance of dynamic mechanism relative to $k = 1/2$ double auction and relative to the Pareto frontier. 95% confidence bands for the Pareto frontier are shown by dashed lines, and 95% confidence bands for the minimum, maximum and quartile-ranked equilibria of the $k = 1/2$ double auction are shown in purple. The confidence region for the dynamic mechanism is shown in red. Units are in terms of the observable value of the car.

Figure 11 demonstrates that the $k = 1/2$ double auction can perform quite well, in particular for the seller, relative to the dynamic mechanism. The equilibria do not appear to perform better for the buyer than the dynamic mechanism, as the dynamic mechanism already yields close to the maximum buyer payoff possible under two-sided uncertainty.

Figure 11 also demonstrates new numerical evidence about the efficiency of the $k = 1/2$ double auction. Chatterjee and Samuelson (1983) demonstrated that when both buyer and seller types are distributed uniformly on $[0, 1]$, the $k = 1/2$ double auction reaches the Pareto frontier. However, Satterthwaite and Williams (1989) showed the symmetric uniform setting is a knife-edge case, and that generically only the $k = 1$ and $k = 0$ double auctions can attain the Pareto frontier. Intuitively, the $k = 1$ and $k = 0$ double auctions correspond exactly to the seller-optimal and buyer-optimal mechanisms, lying on the frontier, whereas for $k = 1/2$ or for any $k \in (0, 1)$, it is not obvious how close outcomes will lie relative to the efficient frontier.[66] However, Figure 11 illustrates that the $k = 1/2$ double auction has equilibria which can closely approximate the frontier.

Table 5 displays the quantitative results corresponding to Figure 11.[67] The second-best direct mechanism is also reported for comparison. In the used-car dealers sample, shown in panel A, the dynamic mechanism achieves a similar total gains from trade as does the 0.75 quantile equilibrium of the $k = 1/2$ double auction: both yield gains of approximately 0.39 of the market value of the car. The seller's gains are lower in the dynamic mechanism than in all of the k double auction equilibria in Table 5. For example, if the 0.25 quantile equilibrium of the $k = 1/2$ double auction were to occur, the seller's payoff would

[66] Evaluating (11) and (12) at $k = 1$ and $k = 0$ yields the seller-optimal and buyer-optimal mechanisms discussed in Section 6.1.

[67] The estimated minimum equilibrium is omitted from Table 5 for the sake of space.

Table 5: Performance of bargaining relative to $k = 1/2$ double auction

A. Used-car dealers sample

	Second-best	0.25 quantile	0.5 quantile	0.75 quantile	Max	Dynamic mechanism
Expected gains from trade	0.440 (0.006)	0.256 (0.034)	0.348 (0.019)	0.390 (0.012)	0.431 (0.007)	0.388 (0.007)
Buyer gains	0.210 (0.005)	0.136 (0.018)	0.164 (0.011)	0.195 (0.006)	0.211 (0.004)	0.256 (0.007)
Seller gains	0.231 (0.005)	0.120 (0.020)	0.185 (0.012)	0.195 (0.010)	0.221 (0.004)	0.087 (0.001)
Probability of trade	0.773 (0.005)	0.332 (0.081)	0.513 (0.059)	0.605 (0.042)	0.732 (0.012)	0.715 (0.002)

B. Fleet/lease sellers sample

	Second-best	0.25 quantile	0.5 quantile	0.75 quantile	Max	Dynamic mechanism
Expected gains from trade	0.349 (0.005)	0.183 (0.006)	0.249 (0.007)	0.305 (0.005)	0.341 (0.006)	0.334 (0.003)
Buyer gains	0.184 (0.002)	0.088 (0.005)	0.109 (0.006)	0.162 (0.004)	0.173 (0.004)	0.203 (0.003)
Seller gains	0.165 (0.005)	0.096 (0.005)	0.140 (0.008)	0.144 (0.005)	0.168 (0.004)	0.096 (0.001)
Probability of trade	0.773 (0.006)	0.307 (0.013)	0.443 (0.015)	0.642 (0.015)	0.743 (0.011)	0.774 (0.002)

Notes: Indirect (k=1/2 double auctoin) mechanism compared to current dynamic mechanism based on estimated gains and probability of trade. First column reports the second-best direct mechanism. Columns 2 through 5 report the outcomes of various equilibria in the k double auction, ranked by estimated gains from trade. Standard errors are from 200 bootstrap replications. Units (other than for probability of trade) are in terms of the observable value of the car.

increase from 0.087 of the market value of the car to 0.120 of the market value. The probability of trade, on the other hand, is higher in the dynamic mechanism in all but the maximum estimated equilibrium.

In the fleet/lease sample (panel B), seller gains are again estimated to be higher in each of the equilibria

of the $k = 1/2$ double auction. Specifically, seller gains would increase from 0.096 to 0.168 of the market value of the car if the maximum equilibrium were played. By other measures, on the other hand, it is the dynamic mechanism which outperforms the $k = 1/2$ double auction. Only the maximum equilibrium achieves a higher total gains from trade than the dynamic mechanism, increasing from 0.334 to 0.341 of the observable value of the car. The probability of trade (77.4%) and buyer gains (0.203 of the market value) are higher in the dynamic mechanism than in all of the double auction equilibria estimated. These findings are suggestive that a $k = 1/2$ double auction would do little to improve outcomes in fleet/lease bargaining, but could potentially result in slightly more efficiency for bargaining over vehicles sold by used-car dealers.

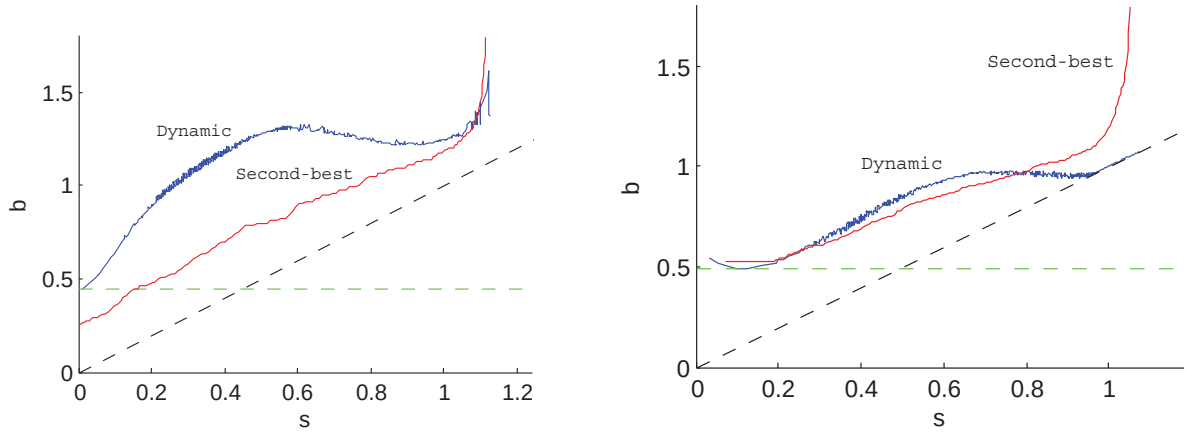
7.3 The Region of Trade

This section takes the analysis of trade one step further by examining the region of buyer and seller types where trade occurs. This type of plot is common in the theoretical bargaining literature; here I merge the plot of the theoretical region with an estimate of this region under real-world bargaining.

Figure 12 plots this estimated region under the dynamic mechanism (marked in blue) and under the second-best mechanism maximizing the total gains from trade (marked in red). For the dynamic mechanism, this region of trade is determined by the allocation function x_τ^D estimated using the quantile-matching procedure described in Section 6.3, where I use the region of trade for $\tau = 2$, or bargaining that ended in the second period of the game. The region of trade for the second-best mechanism is determined by the allocation function, x, from Myerson and Satterthwaite (1983) and Williams (1987), described in Section 6.1. Both mechanisms depend on the auction high bid, so it is only possible to plot the region of trade after selecting a particular value of the auction high bid. Figure 12 displays the region of trade when the high bid is approximately 0.5 of the market value, represented by the horizontal dashed line. The diagonal dashed line marks the 45 degree line.

The first-best outcome would be for trade to occur whenever the buyer values the good more than the seller, which is the region northwest of the dashed lines. The gap between the dashed lines and the second-best mechanism represents the region of buyer and seller types where trade will fail to occur because of incomplete information. This gap appears to be of similar magnitude in both the used-car and fleet/lease sellers samples, and the failed trades occur primarily in cases close to the 45 degree line, where the buyer and seller have similar valuations for the car. For example, panel (a) indicates that if a seller values a car at 0.6 of the market value, trade will occur under the second-best mechanism if and only if the buyer's valuation exceeds 0.8 of the market value.

Figure 12, panel (a) indicates that at a high bid value of 0.5 the region of trade is smaller under the dynamic bargaining mechanism than under the second-best for cars sold by used-car dealers. The figure implies that if a seller values a car at 0.4 of the market value, trade will only occur in the dynamic mechanism when the buyer's valuation exceeds the market value (that is, a normalized valuation of one). In panel (b), for cars sold by fleet/lease sellers, on the other hand, the region of trade appears quite similar under the second-best and the dynamic mechanism. When the seller has a valuation of 0.4 of the



(a) Used-car dealers

(b) Fleet/lease sellers

Figure 12: Estimated boundary of trade for the dynamic mechanism (marked in blue) and the boundary of trade in the second-best mechanism maximizing the total gains from trade (marked in red). The horizontal dashed line represents the auction high bid (approximately 0.5) at which the region of trade was evaluated for this figure (marked in green) and the diagonal dashed line represents the 45 degree line (marked in black). The seller's type is displayed on the horizontal axis, and the buyer's type on the vertical axis. Units are in terms of the observable value of the car.

market value, trade occurs under either mechanism when the buyer's valuation is approximately 0.7 of the market value.

8 Model Fit and Robustness

Various issues of model fit and robustness have already been discussed above in the presentation of results. Here I discuss several additional tests looking at the fit of the buyer distribution, the moment inequalities in the estimation of seller types, and the independence of the model's various moving pieces.

8.1 Tests of Buyer Distribution Fit

The estimation approach in this paper relies on the assumption that the auction takes place in a symmetric, independent, private values environment.[68] One test of this assumption is a linear regression of

[68] As discussed in Athey and Haile (2007), identification results in ascending auctions outside of the IPV framework are quite limited. Roberts (2010) provided one approach to account for unobserved auction heterogeneity, but the approach requires the assumption that reserve prices contain no private information of sellers. This assumption is natural in the setting of Roberts (2010), where it is the auction house and not sellers who choose the reserve price. The measurement error approach to allowing for unobserved heterogeneity, developed in Li and Vuong (1998) and applied in Krasnokutskaya (2011) and Asker (2010), has not been applied to ascending auction bids because of dependencies among order statistics, as pointed out in Aradillas-Lopez et al. (2012) and Athey and Haile (2002). It may be possible to apply the measurement error

normalized high bids on the number of bidders present, treating the number of bidders as exogenous. Under the null hypothesis that the auction environment is a symmetric IPV setting, bids should be increasing in the number of bidders.[69] In both the used-car and fleet/lease sellers samples, this regression yields a slope coefficient which is positive and significant at the 95% level. A more powerful test is to regress quantiles of bids on the number of bidders present, as suggested in Chetverikov et al. (2012).[70] Estimating this regression for the 0.05, 0.1, ..., 0.95 quantiles yields positive and significant slope coefficients at every quantile except the 0.9 and 0.95 quantiles of the fleet/lease sample, where the slope coefficients are insignificant. The test indicates that the distribution of the second order statistic increases as the number of competitors increases, consistent with a symmetric IPV environment.

Two additional tests are discussed in Athey and Haile (2007). The first test relies on the property that, under the null hypothesis of symmetric IPV, the estimate of the distribution of buyer valuations obtained using observations in which N bidders were present should be the same as the estimate obtained using observations in which $N' \neq N$ bidders were present. To implement this test, I use a two-sample Kolmogorov-Smirnov test to compare the estimated distribution from auctions in which 48 bidders were present (the most frequently occurring N) to the estimated distributions from each of the nine other most frequently occurring values of N, which range from 45 to 57. For the used-car dealers sample, four of the nine tests failed to reject the null of symmetric IPV. For the fleet/lease sellers sample, all nine tests failed to reject symmetric IPV. Therefore, by this standard, evidence in favor of the symmetric IPV environment is weaker in the used-car dealers sample than in the fleet/lease sample.[71]

A second test suggested by Athey and Haile (2007) is to compare estimates of the buyer distribution obtained by using different order statistics; under the null hypothesis of symmetric IPV, these estimated distributions should coincide. In order to apply this approach, I use a small subset of 6,362 auction sales by fleet/lease sellers for which both the second and third order statistics are observed. The details of the construction of this subset are found in Appendix B. Again using the ten most frequently observed values of N in the data, I find that, at nine out of these ten values, a Kolmogorov-Smirnov test fails to reject that the estimated distributions are the same when using the second vs. third order static, again

approach using the auction high bid and secret reserve price as two different measures of the unobserved heterogeneity, as Decarolis (2009) does in a first price auction setting, but the approach would not yield an estimate of seller types needed for the estimation of the dynamic mechanism described in Section 6.3. Finally, Aradillas-Lopez et al. (2012) examined ascending auctions with correlated values and presented an approach for estimating bounds on buyer and seller surplus but not for estimating the underlying distribution of buyer valuations, and hence the approach cannot be used to compare alternative mechanisms.

[69] As pointed out in Athey and Haile (2007), this condition is necessary but not sufficient for a setting to be symmetric IPV; prices can also increase with the number of bidders in a common values environment.

[70] I perform this regression for a given quantile u by selecting, for each value of N (the number of bidders present) with at least 1,000 auction observations, the u^{th} quantile of high bids from auctions with N bidders. I then regress these quantiles on N. This regression could also be performed using standard quantile regression. For a comparison of standard quantile regression to this type of grouped quantile regression estimation, see Chetverikov et al. (2012).

[71] Note that in both the used-car sellers sample and the fleet/lease sellers sample, when comparing the estimates of the buyer type distribution from auctions with N far below the most frequently observed range of 45 to 57 to the estimates from auctions with N lying within the most frequently observed range, a Kolmogorov-Smirnov test does reject that the estimated distributions are equivalent.

consistent with the symmetric IPV assumption.[72]

Each of the tests in this section and the approach for obtaining the buyer distribution in Section 5.2 uses estimated values of N, the number of bidders present at each auction sale. An alternative approach which does not rely on knowing the number of bidders was proposed by Song (2005) and applied to Korean wholesale auto auctions by Kim and Lee (2008). The key insight of Song (2005) is that the density of the k^{th} order statistic conditional on the j^{th} order statistic (where $k < j$) does not depend on N.[73] Therefore, as long as two or more order statistics are observed, the underlying distribution can be estimated using semi-nonparametric maximum likelihood, as in Gallant and Nychka (1987) and Chen (2007). I perform this exercise using the sample of 6,362 auction sales for which the second and third order statistics are observed. I then compare the estimated distribution using the Song (2005) approach to the main estimates obtained in Section 5.2. A Kolmogorov-Smirnov test rejects that the distributions are the same at the 95% level, indicating that, maintaining the symmetric IPV assumption, the estimated values of N are at best an imperfect measure of the true number of bidders participating. See Appendix C.6 for further details on this estimation procedure.

The procedure applied in this paper for imputing the number of bidders assumes that all bidders present at the auction house are indeed participating. Therefore, those who are not actively participating are treated as though they had very low valuation draws. An alternative to this approach would be to assume that only some fraction of bidders present at any given time actually receives a draw from the distribution of buyer valuations and that the remaining bidders are not participating. To implement this approach, I estimate the model assuming that only half of the bidders present at any point in time were actually partipating in the auction. The resulting estimates yield a worse fit in terms of the theory, yielding estimates of buyer surplus which lie slightly beyond the Pareto Frontier.[74]

8.2 Robustness of Moment Inequalities Estimation of Seller's Type

This section provides evidence of the robustness of the moment inequalities approach for estimating the seller's inverted secret reserve price, $s(R)$. The main results of the paper use a grid over R with $K = 50$ grid points. This approach can result in an estimated lower bound exceeding the estimated upper bound at certain grid points. When $K = 50$, bounds cross at one grid point in the used-car dealers sample and at twelve grid points in the fleet/lease sample. With $K = 100$, these numbers increase to four and 29 grid points with crossings, and at $K = 25$, there are zero and three crossings, respectively.

[72]The order statics relationship used to estimate the distribution of underlying valuations from the observed third order statistic, b_3, is given by

$$F_{b_3}(b) = \frac{1}{2}N(N-1)F_b(b)^{N-2}(1-F_b(b))^2 + NF_b(b)^{N-1}(1-F_b(b)) + F_b(b)^N \tag{16}$$

[73]This conditional density is given by

$$f^{k|j}(y|x) = \frac{(j-1)!(F(y)-F(x))^{j-k-1}(1-F(y))^{k-1}f(y)}{(j-k-1)!(k-1)!(1-F(x))^{j-1}}$$

[74]The resulting estimates are similar to those in panels (a) and (c) of Figure 13.

Pakes et al. (2011), Bontemps et al. (2011), and others discuss the case when bounds cross, leading to a point estimate rather than a set. One test of model specification in this setting is given by Andrews and Soares (2010), and is analogous to an over-identification J test in a traditional GMM setting. Under the null hypothesis of correct specification, the confidence set calculated from a generalized moment selection procedure should be nonempty.[75] I employ this approach by modeling $s(R)$ as a b-spline function with 15 uniformly spaced knots, and estimate the confidence set using the approach described in Andrews and Soares (2010). I find that the test fails to reject the null hypothesis of correct model specification.

Moreover, several arguments suggest that these crossing bounds are not problematic. First, from Figure 6 above, it is clear that when the lower bound does lie above the upper bound, it is not by an economically relevant amount. Also, not immediately evident from Figure 6 is the fact that nowhere does the 95% confidence interval for the lower bound lie strictly above the 95% confidence interval for the upper bound. Finally, as K increases from 25 to 100, the pointwise confidence interval for the identified set does not change by an observable amount.

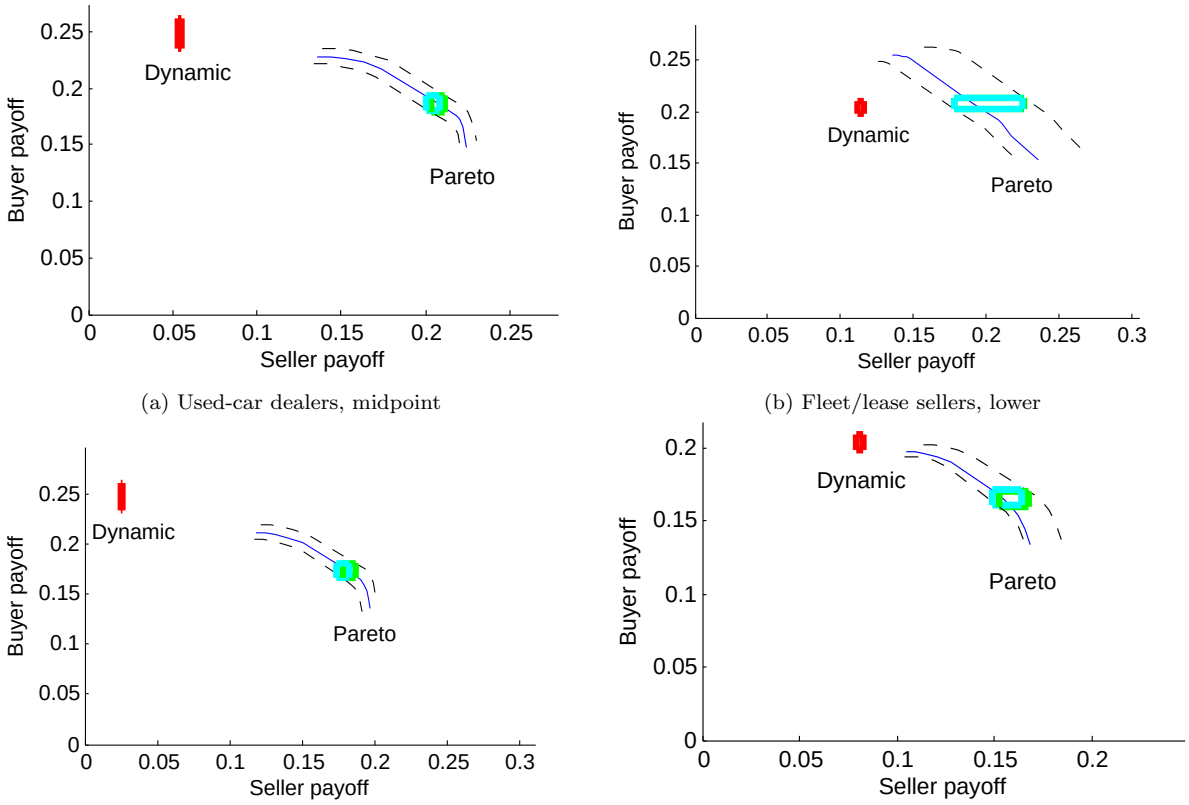
While the number of crossings of the estimated upper and lower bounds does not appear large, the number of individual observations violating these bounds is large. Recall that the overall upper and lower bounds for seller types are estimated by taking means of individual-level upper and lower bounds. I find that the individual-level observations with informative lower bounds (cases in which the seller walked away from bargaining) lie above the overall estimated upper bound 55% of the time in the fleet/lease sellers sample and 47% of the time in the used-car dealers sample. Similarly, individual-level observations with informative upper bounds lie below the overall estimated lower bound 29% of the time in the fleet/lease sellers sample and 27% of the time in the used-car dealers sample. Under the measurement error interpretation of these violations, these results would imply that one-fourth to one-half of the bargaining records with informative bounds are recorded with error (note that this does not imply that one-fourth to one-half of all observations are recorded with error, as not all records have informative bounds, as described in Section 5.3).[76] In practice, these violations may be due to other factors such as heterogeneity in sellers' bargaining costs or different equilibria being played in the bargaining game, either of which would imply that different sellers can have differing secret reserve price functions. Unfortunately, accounting for these alternative factors structurally would be infeasible, as the outcomes of the dynamic bargaining mechanism would depend on these factors and surplus in the dynamic bargaining would no longer be identified.[77] Therefore, the additive, mean-zero measurement

[75] Andrews and Barwick (2011) provided a refinement of this generalized moment selection procedure.

[76] When considering all records, 8% of observations display an informative lower bound less than the estimated lower bound in both the used-car dealers sample and the fleet/lease sellers sample. 14% display an informative upper bound greater than the estimated lower bound in the fleet/lease sample, while 18% fall into this category in the used-car dealers sample.

[77] For example, consider the case where the per-offer bargaining cost, c_S, varies by seller. Then the secret reserve price strategy, $R(s, c_S)$, would depend on s as well as c_S. It can be shown that $R(s, c_S)$ should be strictly decreasing in c_S for any s, and hence $s(R, c_S)$ should be strictly increasing in c_S for any R. Therefore, it is true that at each R the $\tau-$quantile of c_S, $Q^{\tau}_{c_S}$, satisfies $Q^{\tau}_{L|R} \leq s(R, Q^{\tau}_{c_S}) \leq Q^{\tau}_{U|R}$, where $Q^{\tau}_{L|R}$ and $Q^{\tau}_{U|R}$ represent the τ-quantile of the individual-level lower and upper bounds observed from bargaining when the secret reserve price was R. Using this approach, one can bound a different

error model may be an approximation to a more complex structure of heterogeneity among sellers.[78]



(a) Used-car dealers, midpoint

(b) Fleet/lease sellers, lower

(c) Used-car dealers, upper

(d) Fleet/lease sellers, upper

Figure 13: Estimated performance of dynamic mechanism relative to Pareto frontier using the midpoint and upper confidence bounds of seller types for used-car dealers, and using the lower and upper confidence bound of seller types for fleet/lease sellers. Second-best and trade-maximizing mechanisms are shown on the Pareto frontier. 95% confidence bands are shown around the frontier and each mechanism. Units are in terms of the observable value of the car.

In order to analyze the sensitivity of the welfare estimates to the choice of using the midpoint, upper bound, or lower bound of the pointwise confidence interval for $s(R)$, I estimate the model under each specification. The baseline estimates from Figure 9 use the lower confidence bound for the used-car

inverted secret reserve function at each quantile of c_S. The main challenge with carrying this approach through the rest of the steps of the paper is that the direct revelation mechanism corresponding to the dynamic bargaining mechanism would depend not only on the valuations of the buyer and seller, but also on the bargaining costs of the buyer and seller, and as these costs are unobserved the approach presented in Section 6.3 to estimate the direct mechanism (and, consequently, the surplus under the dynamic bargaining) would no longer be feasible.

[78] For example, if sellers have heterogeneous per-offer costs, the approach of taking expectations of individual-level upper and lower bounds would be capturing bounds on the conditional expectation of s given R rather on the true $s(R)$. That is, $E[L|R] \leq E[s(R, c_S)|R] \leq E[U|R]$.

dealers sample and the midpoint for the fleet/lease sample. The other specifications are displayed in Figure 13. Note that the choice of the midpoint, lower bound, or upper bound does not affect the estimate of buyer surplus in the dynamic mechanism, as this is solved for using the quantile-matching approach from Section 6.3. The choice does affect the estimated seller surplus and the estimated Pareto frontier.

As seen in panels (a) and (c) of Figure 13, for the used-car dealers sample, estimation using the upper bound or midpoint yields an outcome for dynamic bargaining which lies slightly beyond the outcomes which the estimated frontier predicts are possible. The fit using the lower bound, shown in Figure 9, is much better, yielding a Pareto frontier which is consistent with estimated surplus because the 95% band about the Pareto frontier overlaps the 95% confidence band about the dynamic bargaining surplus. In the fleet/lease sellers sample, estimation with the lower bound results in more seller gains, and estimation with the upper bound results in less seller gains. Intuitively, using a lower estimate of the sellers types s implies that sellers must have been gaining more from trades which occurred. In both panels (b) and (d), the 95% confidence region for the dynamic bargaining outcome lies below or overlaps the region of surplus which is predicted to be achievable by the Pareto frontier, yielding a better fit that in panels (a) and (c).[79]

8.3 Additional Tests of Independence Assumptions

The model and estimation approach in this paper assume independence of various moving pieces. For example, the approach controlling for observables assumes that covariates are independent of players' private valuations. A simple test of this assumption is to examine whether the variance of normalized high bids depends on $\Gamma(X)$, the observable component of buyers and sellers types estimated in Section 5.1. I find that the variance is quite different for cars with $\Gamma(X) <$\$5,000 than for cars with $\Gamma(X) \geq$\$5,000, implying there is heteroskedasticity in the first-stage regression. I then re-estimate the distributions and counterfactuals using only observations with $\Gamma(X) <$\$5,000 . The results are displayed in row 2 of Table 6. The full sample estimates are displayed for comparison (using the midpoint of the confidence bounds for the seller types). For both used-car dealers and fleet/lease sellers, the estimated gains from trade lie within 0.49 to 0.60 of the market value of the car, whereas in the baseline case gains were in the range of 0.33 to 0.40 of market value. This finding is indicative that gains from trade (relative to the market value of the car) are higher on vehicles with a lower market value.

Estimating the model on low vs. high-value cars can be viewed as a check on whether heteroskedasticity matters in the first-stage regression estimated in (3). Alternatively, the exercise can be viewed as a check on whether differences in seller demand uncertainty, ζ, affect estimation, as I find that high-value

[79] While the exercise illustrated in Figure 13 treats the outcome in panels (a) and (c) as an issue of model fit, Bazerman et al. (1998) argue that real-world bargaining can potentially yield *more* efficient outcomes than the theoretical bounds of Myerson and Satterthwaite (1983) and Williams (1987) would predict possible. Bazerman et al. (1998) review evidence from the behavioral and experimental literature suggesting that more efficient outcomes can occur due to non-traditional utility functions (where one player's utility nests the other's), limits on players' abilities to mimic other types, and other features of bounded rationality.

cars exhibit smaller differences in estimates of $\Gamma(X)$ from (3) using secret reserve prices and estimates from an analogous regression using high bids. While the gains from trade clearly differ in low vs. high-value cars, the qualititative findings are the same. First, the deadweight-loss due to incomplete information is small, as the first-best gains are close to the second best gains. Second, the deadweight loss due to limited commitment is also small, as the dynamic mechanism lies close to the second-best in the used-car dealers sample and even more so in the fleet/lease sample.

A similar check on this assumption uses only cars which are less than four years old, shown in row 3 of Table 6. For both used-car dealers and fleet/lease sellers, the estimated gains from trade decrease to a range of 0.17 to 0.23 of the market value, indicating that less gains from trade exist on newer cars, a finding consistent with the estimated expected gains from trade from websites such as Edmunds.com. As above, the main qualitative results are unchanged.

Table 6: Estimated expected gains from trade in selected subsamples

	Used-car dealers sample			*Fleet/lease sellers sample*		
	First-best	Second-best	Dynamic mechanism	First-best	Second-best	Dynamic mechanism
Full Sample, midpoint of seller bounds	0.404 (0.005)	0.393 (0.004)	0.353 (0.007)	0.358 (0.005)	0.349 (0.005)	0.334 (0.003)
Observable value < $5,000	0.538 (0.005)	0.525 (0.005)	0.493 (0.010)	0.598 (0.006)	0.585 (0.006)	0.579 (0.009)
Less than 4 years old	0.2111 0.0405	0.2051 0.0364	0.1727 0.0291	0.2344 0.0038	0.232 0.0038	0.2187 0.0061
First run of car	0.417 (0.005)	0.408 (0.005)	0.387 (0.010)	0.397 (0.010)	0.390 (0.010)	0.373 (0.006)
First car sold by seller	0.422 (0.007)	0.411 (0.007)	0.359 (0.011)	0.409 (0.012)	0.401 (0.012)	0.406 (0.015)

Notes: Estimates of expected gains from trade using various sample restrictions. First row displays full sample estimtes using the midpoint of seller bounds for s(R). Second row uses only cars with an estimated observable value, Γ(X), less than $5,000. Third row uses only cars less than four years old. Fourth row uses only the first attempt to sell a given vehicle. Fifth row uses only the first car offered by a seller on a given auction day. Standard errors are from 200 bootstrap replications. Units (other than for probability of trade) are in terms of the observable value of the car.

A potential concern with the estimation approach in this paper is that players' values for cars in different auction observations may not be independent. I examine here two different ways in which sellers types may be correlated across auctions.[80] First, I limit the sample to the first time a given car went up

[80]Correlation may also exist in buyer types across different auctions but testing is simpler with seller types as each buyer is only observed if he is the high bidder.

for sale (recall that cars can go up for sale multiple times if they do not sell on the first run), as a seller's value for the same car in different runs is likely to be correlated. The results are displayed in row 4 of Table 6, and yield estimates similar to the full sample estimates. Second, I re-estimate the model using only the first car which a seller sold on a given auction day, as any correlation between auction sales due to inventory behavior is likely to be the strongest for cars sold on the same day. The results are displayed in row 5 of Table 6. The estimated expected gains from trade are again quite similar to the full sample estimates, indicating that this type of correlation may not be a concern.[81]

One limitation of this paper is the maintained assumption that buyer and seller types are independent. The correlation between buyer and seller types cannot be tested directly because buyer types are not identified (only the distribution of buyer types is identified). However, the raw correlation between normalized secret reserve prices and normalized auction high bids is available. Before outliers above the 99^{th} percentile are trimmed (as explained in the description of which observations were dropped in Section 4), this correlation is 0.05 in both samples. However, after eliminating outliers, the correlation is 0.36 in the fleet/lease sample and 0.57 in the used-car sellers sample. Extending the above analysis to take this correlation into account would pose significant challenges. In particular, the estimation of the allocation function corresponding to the dynamic mechanism would not be straightforward.[82]

One test of the implications of ignoring correlation between seller and buyer values is to simulate from the observed marginal distributions of the high bid and seller's secret reserve price to measure how often the high bid should exceed the reserve. In the used-car dealers sample, this simulation implies that trade should occur through the auction 30% of the time, whereas it only occurs 16% of the time in practice (from Table 2). For fleet/lease sellers the fit is much better: in the data, the high bid exceeds the reserve price 37% of the time, and simulating from the marginal distributions implies this should occur 41% of the time under independence of the two variables. Therefore, correlation between buyer and seller types appears to be less problematic in the fleet/lease sellers sample than in the used-car dealers sample, and estimating the model separately for the two samples, as is done throughout the paper, serves as an additional check on the effects of ignoring correlation. As shown in Section 7, the findings are qualitatively similar in both samples.

[81] In practice, there may be some dependencies between auction observations which these tests cannot account for. For example, if items sold are considered close substitutes by participants, inventory optimization behavior by buyers or sellers could lead to some valuations being dependent across auctions. While assuming independence of auction observations is common throughout the empirical auctions literature, it is particularly useful in the setting of this paper for honing in on the dynamics within an auction/bargaining sequence rather than between auctions. In conversations with auction participants, sellers claim to consider each auction in isolation, and buyers often have a specific retail consumer lined up for a specific car, lending some support to the assumption. Recent papers which have explored dynamic dependencies of valuations between multiple sealed-bid second price auctions include Zeithammer (2006), Ingster (2009), Sailer (2006), and Lewis and Backus (2010).

[82] As mentioned above, the measurement error technique of Li and Vuong (1998) would not yield an estimate of seller types (instead, it would only provide an estimate of the distribution of seller types), which is essential for the estimation of the dynamic mechanism described in Section 6.3.

9 Conclusion

This paper examines the efficiency of bargaining from a real-world setting with two-sided incomplete information. I develop a model and strategy for estimating the distributions of valuations on both sides of the market without relying on a particular structure or equilibrium for the bargaining game.

To estimate the distribution of buyer valuations, I use an order statistics approach. To estimate the distribution of seller valuations, I present a new approach using moment inequalities to bound the inverse of the seller's secret reserve price strategy. The bounds are consistent with the seller using a strictly increasing strategy, as implied by the model. I also use a bounds approach to obtain simple bounds on the costs of bargaining. I find these costs to be small and reasonable; the discount factor is close to one and the costs of making counteroffers are close to zero.

After estimating the distributions of buyer and seller types, I map these distributions into the static, direct revelation mechanism framework which traces out the efficient frontier derived in Myerson and Satterthwaite (1983) and Williams (1987). These mechanisms depend on the distributions of buyer and seller valuations. I measure the deadweight loss due to incomplete information as the gap between the first-best trade line and the second-best frontier. This gap is also small in this market. I find that the second-best mechanism achieves 98% of the first-best gains from trade.

I also find that the deadweight loss due to limited commitment is quite small. The dynamic mechanism achieves 88% of possible surplus for cars sold by used-car dealers and 96% of possible surplus for cars sold by large fleet and lease companies. 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."

Although the findings are suggestive that bargaining is quite efficient, improvements are still possible, particularly for cars sold by used-car dealers. The gains from trade and the probability of trade could increase by 14% and 8% if bargaining outcomes were to move to northeast-most point of the frontier. I demonstrate that these findings are robust to a variety of specifications.

I also compare the currently used dynamic mechanism to a $k = 1/2$ double auction, a simple, static alternative to dynamic, alternating-offer bargaining. I adapt methods from Satterthwaite and Williams (1989) to solve for equilibria of the $k = 1/2$ double auction given the estimated distributions. Findings indicate that the currently used dynamic mechanism performs about as well as the 0.75 quantile equilibrium examined here. I find that a wide range of other equilibrium outcomes is possible in this setting; some of these equilibria would be improvements over dynamic bargaining and others would not, and without a model of equilibrium selection it is impossible to say what would happen in practice if the $k = 1/2$ double auction were to be implemented.

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 an auction house is a platform in a two-sided market, required to attract both buyers and sellers, each of whom has a private 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.[83]

An auction followed by dynamic, alternating-offer bargaining, therefore, *can* potentially dominate a public reserve auction in terms of the buyer's payoff, total surplus, or the overall level of trade, and this paper provides evidence that the mechanism performs well in this market: gains from trade for the buyer are 61% and 46% higher in the dynamic mechanism than in a public reserve auction for used-car and fleet/lease sellers, respectively. Similarly, the volume of trade—an indicator used throughout the industry as a measure of auction house performance—is 8–18% higher under the currently used dynamic mechanism than it would be under a public reserve auction.

The nature of wholesale auto auctions as a two-sided market may also provide an explanation for the result found above that buyers appear to obtain a higher share of surplus than sellers. As discussed in Rochet and Tirole (2006), it is often the case in two-sided markets with network effects that participants on one side of the market obtain a higher share of surplus than those on the other side. In the case of auto auctions, a competing auction house which offered a higher payoff to sellers, for example, by switching to a public reserve auction, might be unable to attract a sufficient volume of buyers for the market to function.[84] Therefore, it may be the case that two-sided uncertainty and competition among auction houses plays an important role in a platform's choice of mechanism.

References

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

Ambrus, A., Chaney, E., and Salitsky, I. (2011). Pirates of the Mediterranean: An empirical investigation of bargaining with transaction costs. Working Paper, Harvard.

Andrews, D. and Barwick, P. (2011). Inference for parameters defined by moment inequalities: A recommended moment selection procedure. Working Paper. Forthcoming in *Econometrica*.

[83] 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 et al. 2009), and many procurement auctions (Huh and Park 2010). Other examples include privatization auctions in post-Soviet Europe (Dooley et al. (1993)) and French timber auctions (Elyakime et al. 1997).

[84] The original founder of the auction houses studied in this paper was vehemently opposed to post-auction bargaining. Eventually, however, the auction houses adopted the practice in order to be in line with the rest of the industry, suggestive that network effects in attracting buyers and sellers to the auction house may indeed be at play.

Andrews, D. and Soares, G. (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica*, 78(1):119–157.

Aradillas-Lopez, A., Ghandi, A., and Quint, D. (2012). Identification and inference in ascending auctions with correlated private values. Working Paper. Forthcoming in *Econometrica*.

Ashenfelter, O. (1989). How auctions work for wine and art. *The Journal of Economic Perspectives*, 3(3):23–36.

Ashenfelter, O. and Graddy, K. (2003). Auctions and the price of art. *Journal of Economic Literature*, 41(3):763–787.

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review*, 100(3):724–762.

Athey, S. and Haile, P. (2002). Identification of standard auction models. *Econometrica*, 70(6):2107–2140.

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

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

Ausubel, L. and Deneckere, R. (1993). Efficient sequential bargaining. *Review of Economic Studies*, 60(2):435.

Bajari, P. and Hortasu, A. (2003). The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *The RAND Journal of Economics*, 34(2):329–355.

Bazerman, M., Gibbons, R., Thompson, L., and Valley, K. (1998). Can negotiators outperform game theory? In Halpern, J. and Stern, R., editors, *Debating Rationality: Nonrational Aspects of Organizational Decision Making*, pages 79–99. Ithaca, NY: Cornell University Press.

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

Bontemps, C., Magnac, T., and Maurin, E. (2011). Set identified linear models. IDEI Working Paper 494. Forthcoming in *Econometrica*.

Bulow, J. and Klemperer, P. (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 Samuelson, W. (1983). Bargaining under incomplete information. *Operations Research*, pages 835–851.

Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, 6:5549–5632.

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

Chernozhukov, V., Lee, S., and Rosen, A. (2011). Intersection bounds: Estimation and inference. Working Paper, MIT. Forthcoming in *Econometrica*.

Chetverikov, D., Larsen, B., and Palmer, C. (2012). IV quantile regression for group-level treatments. Working Paper, MIT.

Cho, S.-J., Paarsch, H., and Rust, J. (2010). Is the 'Linkage Principle' valid? Evidence from the field. Center for Economic Institutions Working Paper Series No. 2010-4.

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.

Decarolis, F. (2009). When the highest bidder loses the auction: Theory and evidence from public procurement. *Bank of Italy Temi di Discussione (Working Paper) No*, 717.

Dooley, M., Friedman, D., and Melese-d'Hospital, S. H. (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 Siow, A. (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 Shannon, C. (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81(1):201–219.

Eklof, M. and Lunander, A. (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., Laffont, J., Loisel, P., and Vuong, Q. (1994). First-price sealed-bid auctions with secret reservation prices. *Annales d'Economie et de Statistique*, pages 115–141.

Elyakime, B., Laffont, J.-J., Loisel, P., and Vuong, Q. (1997). Auctioning and bargaining: An econometric study of timber auctions with secret reservation prices. *Journal of Business & Economic Statistics*, 15(2):209–220.

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

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

Gallant, A. and Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica*, pages 363–390.

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

Genesove, D. (1993). Adverse selection in the wholesale used car market. *The Journal of Political Economy*, 101(4):644–665.

Genesove, D. (1995). Search at wholesale auto auctions. *The Quarterly Journal of Economics*, 110(1):23–49.

Grant, S., Kajii, A., Menezes, F., and Ryan, M. (2006). Auctions with options to re-auction. *International Journal of Economic Theory*, 2(1):17–39.

Grether, D., Porter, D., and Shum, M. (2012). A field experiment of jump bidding in online automobile auctions. Working Paper, California Institute of Technology.

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

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

Hammond, R. and Morrill, T. (2012). Strategic exclusion of the highest-valued bidders in wholesale automobile auctions. Working Paper, North Carolina State University.

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

Holmström, B. and Myerson, R. (1983). Efficient and durable decision rules with incomplete information. *Econometrica*, pages 1799–1819.

Horstmann, I. J. and LaCasse, C. (1997). Secret reserve prices in a bidding model with a resale option. *American Economic Review*, 87(4):663–684.

Hortaçsu, A., Matvos, G., Syverson, C., and Venkataraman, S. (2010). Are consumers affected by durable goods makers' financial distress? The case of auto manufacturers. Working Paper, University of Chicago.

Huh, W. and Park, K. (2010). A sequential auction-bargaining procurement model. *Naval Research Logistics (NRL)*, 57(1):13–32.

Ingster, A. (2009). A structural study of a dynamic auction market. Unpublished manuscript, New York University.

Ji, L. (2006). *Three essays on multi-round procurement auctions*. PhD thesis, Vanderbilt University.

Ji, L. and Li, T. (2008). Multi-round procurement auctions with secret reserve prices: Theory and evidence. *Journal of Applied Econometrics*, 23(7):897–923.

Judd, K. (1998). *Numerical Methods in Economics*. Cambridge: The MIT Press.

Katkar, R. and Reiley, D. (2007). Public versus secret reserve prices in eBay auctions: Results from a pokémon field experiment. *The BE Journal of Economic Analysis & Policy*, 6(2):7.

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

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

Kim, K. I. and Lee, J. (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.

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

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

Lacetera, N., Pope, D. G., and Sydnor, J. (2011). Heuristic thinking and limited attention in the car market. Working Paper. Forthcoming in *American Economic Review*.

Lacetera, N. and Sydnor, J. (2012). Would you buy a honda made in the U.S.? The impact of production location on manufacturing quality. NBER Working Paper No. 18005.

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

Lewis, G. and Backus, M. (2010). A demand system for a dynamic auction market with directed search. Working Paper, Harvard.

Li, H. and Tan, G. (2000). Hidden reserve prices with risk-averse bidders. Mimeo, Penn. State University.

Li, T. and Vuong, Q. (1998). Nonparametric estimation of the measurement error model using multiple indicators. *Journal of Multivariate Analysis*, 65(2):139–165.

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

Merlo, A., Ortalo-Magne, F., and Rust, J. (2008). The home selling problem: Theory and evidence. Working Paper, University of Maryland.

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

Myerson, R. and Satterthwaite, M. (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.

Pakes, A., Porter, J., Ho, K., and Ishii, J. (2011). Moment inequalities and their application. Working Paper, Harvard.

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

Roberts, J. (2010). Unobserved heterogeneity and reserve prices in auctions. Working Paper, Duke.

Rochet, J. and Tirole, J. (2006). Two-sided markets: A progress report. *The RAND Journal of Economics*, 37(3):645–667.

Rosenkranz, S. and Schmitz, P. W. (2007). Reserve prices in auctions as reference points. *The Economic Journal*, 117(520):637–653.

Roth, A. E. (1995). Bargaining experiments. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.

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

Sailer, K. (2006). Searching the ebay marketplace. CESifo Working Paper Series 1848, CESifo Group Munich.

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

Sawyers, A. (2012). Used-car drought relief? Next year. *Automotive News*. Url: http://www.autonews.com/apps/pbcs.dll/article?AID=/20120305/RETAIL07/303059967. Posted Mar. 5, 2012. Retrieved Jun. 25, 2012.

Sieg, H. (2000). Estimating a bargaining model with asymmetric information: Evidence from medical malpractice disputes. *The Journal of Political Economy*, 108(5):1006–1021.

Song, U. (2005). Nonparametric estimation of an eBay auction model with an unknown number of bidders. Economics Department Discussion Paper 05-14, University of British Columbia.

Tadelis, S. and Zettelmeyer, F. (2011). Information disclosure as a matching mechanism: Theory and evidence from a field experiement. Working Paper, UC Berkeley.

Tang, X. and Merlo, A. (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*, volume 18:1, pages 221–272. Elsevier.

Vincent, D. (1995). Bidding off the wall: Why reserve prices may be kept secret. *Journal of Economic Theory*, 65(2):575–584.

Wang, R. (2000). Bidding and renegotiation in procurement auctions. *European Economic Review*, 44(8):1577–1597.

Watanabe, Y. (2009). Learning and bargaining in dispute resolution: Theory and evidence from medical malpractice litigation. Working Paper, Northwestern University.

Williams, S. (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 Bewley, T., editor, *Advances in Economic Theory: Fifth World Congress*, pages 33–70. Cambridge: Cambridge University Press.

Zeithammer, R. (2006). Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476.

A Proofs

Proof of Proposition 1

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

$$\pi^B(\beta_j, b_i) = \delta_I \delta_E \Bigg((b_i - \beta_j) \Pr\left(d_2^S = A | \beta_j\right) + \delta_I \delta_E \Pr\left(d_2^S = C | \beta_j\right) E_{p_2^S} \left[\max\left\{ b_i - p_2^S, 0, V_3^B\left(b_i | \{\beta_j, p_2^S\}\right) \right\} \middle| \beta_j \right] \Bigg) - c_B$$

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 $M(\beta_j, b_i)$ 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 $b_i \leq \beta_j$, then $\beta_i < b_i \leq \beta_j$, so bidder i is not the high bidder, and would not have been even if he had bid $\beta_i = b_i$.

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

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

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

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 $\beta_j \leq b_i$, then $\beta_j \leq b_i < \beta_i$, then bidder i is the high bidder and gets payoff $M(\beta_j, b_i)$, 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 $\beta_j > b_i$, then the following is true:

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

   (b) If $b_i < \beta_j < \beta_i$, 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, $V_t^S(s|h^t)$ 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 decline. 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

$$\delta_I W_T^S(s,p|h^{T-1}) - c_S \equiv \delta_I \left(p\delta_E \Pr(d_T^B = A|h^T) + s\left(\delta_E(1-\Pr(d_T^B = A|h^T)) + 1 - \delta_E\right)\right) - c_S$$

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

$$V_{T-1}(s,s'|h^{T-1}) = \delta_I \left(p^*(s'|h^{T-1})\delta_E \Pr(d_T^B = A|h^T) + s(\delta_E(1-\Pr(d_T^B = A|h^T)) + 1 - \delta_E)\right) - c_S,$$

and

$$V_{T-1}(s',s'|h^{T-1}) = \delta_I \left(p^*(s'|h^{T-1})\delta_E \Pr(d_T^B = A|h^T) + s'(\delta_E(1-\Pr(d_T^B = A|h^T)) + 1 - \delta_E)\right) - c_S$$

Thus,

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

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 $V_{T-(t-1)}^S(s|h^{T-(t-1)})$ denote the seller's counteroffer payoff with $t-1$ periods remaining, and suppose $V_{T-(t-1)}^S(s|h^{T-(t-1)})$ is weakly increasing 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

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

Note that

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

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 $V^B_3(\cdot)$ 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 $c_B > 0$.

Recall that ζ represents the demand uncertainty which is unknown at the time the seller chooses R. Let Z represent the support of ζ. Let $F_b(\cdot;\zeta)$ and $F_{p^A}(\cdot;\zeta)$ denote the distributions (and similarly for the densities) of buyer valuations and high bids conditional on ζ. Also note that before ζ is realized, χ will depend on ζ, as will the seller's bargaining payoff, π^S.

The seller's ex-ante expected payoff can then be written as

$$
\begin{aligned}
&E_\zeta\left[\int_R^{\bar{b}} p^A f_{p^A}(p^A;\zeta)dp^A + \int_{\underline{b}}^R \left[\int_{p^A}^{\chi^{-1}(p^A;\zeta)} s f_b(b;\zeta)db + \int_{\chi^{-1}(p^A;\zeta)}^{\bar{b}} \pi^S\left(p^A, s;\zeta\right) f_b(b;\zeta)db\right] \frac{f_{p^A}(p^A;\zeta)}{1-F_b(p^A;\zeta)}dp^A\right] \\
=&E_\zeta\left[\int_R^{\bar{b}} p^A f_{p^A}(p^A;\zeta)dp^A + \int_{\underline{b}}^R \left[s\left(F_b(\chi^{-1}(p^A;\zeta);\zeta) - F_b(p^A;\zeta)\right) + \pi^S\left(p^A,s;\zeta\right)\left(1-F_b(\chi^{-1}(p^A;\zeta);\zeta)\right)\right]\right. \\
&\left.\times \frac{f_{p^A}(p^A;\zeta)}{1-F_b(p^A;\zeta)}dp^A\right]
\end{aligned}
$$

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

$$
\frac{\partial}{\partial R} = -RE_\zeta\left[f_{p^A}(R;\zeta)\right] + sE_\zeta\left[\frac{F_b(\chi^{-1}(R;\zeta)) - F_b(R;\zeta)}{1-F_b(R;\zeta)} f_{p^A}(R;\zeta)\right] + E_\zeta\left[\pi^S\left(R,s;\zeta\right)\frac{1-F_b(\chi^{-1}(R;\zeta))}{1-F_b(R;\zeta)} f_{p^A}(R;\zeta)\right]
$$

Lemma A implies that $\pi^S\left(p^A, s; \zeta\right)$ is weakly increasing in s, and thus $\frac{\partial}{\partial R}$ will be strictly increasing in s because $F_b(\chi^{-1}(R;\zeta)) > F_b(R;\zeta)$ given that $\chi^{-1}(\cdot;\zeta)$ is strictly increasing. Given that $\frac{\partial}{\partial R}$ is strictly increasing in s, the Edlin and Shannon (1998) Theorem implies that, as long as the optimal $R^*(s)$ lies on the interior of the support of R, $R^*(s)$ will be strictly increasing in s. The support of R is the real line (nothing prevents sellers from reporting a negative reserve), 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 multiplicative 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 multiplicatively separable in $\Gamma(X)$. That the auction high bid will be multiplicatively separable in $\Gamma(X)$ is obvious, given that the bidding function is the identity function. To demonstrate that bargaining offers are also multiplicatively 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 decline. Suppose the buyer's type is given by $b\Gamma(X) \sim F_{b\Gamma(X)}(b\Gamma(X))$, with density $\frac{1}{\Gamma(X)} f_{b\Gamma(X)}(b\Gamma(X))$. The seller's type is given by $s\Gamma(X)$, and the per-period bargaining disutility by $c_S\Gamma(X)$.

In the final period, the buyer will accept a price, $\tilde{p}$, if and only if $\tilde{p} \leq b\Gamma(X)$. The seller's payoff from countering at a price of $\tilde{p}$ is then

$$\delta_I\left(\tilde{p}\delta_E(1 - F_{b\Gamma(X)}(\tilde{p})) + s\Gamma(X)\left(\delta_E F_{b\Gamma(X)}(\tilde{p}) + 1 - \delta_E\right)\right) - c_S\Gamma(X)$$

which has a first order condition of

$$\tilde{p} = s\Gamma(X) + \Gamma(X)\frac{1 - F_{b\Gamma(X)}(\tilde{p})}{f_{b\Gamma(X)}(\tilde{p})}$$

Dividing by $\Gamma(X)$ and changing variables from $b\Gamma(X)$ to b yields

$$\frac{\tilde{p}}{\Gamma(X)} = s + \frac{1 - F_b(\tilde{p}/\Gamma(X))}{f_b(\tilde{p}/\Gamma(X))}$$

Thus, the penultimate bargaining offer in the heterogeneous setting is simply the bargaining offer from the homogeneous good setting scaled by $\Gamma(X)$. Note that it is also the case that the seller's maximized payoff in the heterogeneous setting is simply the payoff from the homogeneous good setting scaled by $\Gamma(X)$.

To complete the proof by induction, suppose that offers and payoffs in periods $T-(t-1)$ and $T-(t-2)$ are simply the homogeneous good objects scaled by $\Gamma(X)$. It remains to be shown that this would imply 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:

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

Note that 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 given the induction assumption that payoffs and offers are scaled by $\Gamma(X)$ in this period. This can be seen from the buyer's payoffs for each action:

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

Therefore, the $\tilde{p}$ in the seller's counteroffer payoff can be replaced with $p\Gamma(X)$, and the seller's choice over $\tilde{p}$ can be replaced with a choice over p. Dividing by $\Gamma(X)$ then yields the same payoffs and offers as in the homogeneous good auction.

Finally, observe that the seller's secret reserve price will then also be scaled by $\Gamma(X)$ in the heterogeneous goods auction. That is, this result requires the assumption that the seller's uncertainty about demand, ζ, referred to in the the proof of Proposition 2, does not enter $\Gamma(X)$ The first order condition for her secret reserve in this case is the following:

$$\frac{\partial}{\partial R} = -R\Gamma(X)E_\zeta\left[f_{p^A}(R;\zeta)\right] + s\Gamma(X)E_\zeta\left[\frac{F_b(\chi^{-1}(R;\zeta)) - F_b(R;\zeta)}{1-F_b(R;\zeta)} f_{p^A}(R;\zeta)\right] + \Gamma(X)E_\zeta\left[\pi^S(R,s;\zeta)\frac{1-F_b(\chi^{-1}(R;\zeta))}{1-F_b(R;\zeta)} f_{p^A}\right.$$

□

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 $x : [p^A, \overline{b}] \times [\underline{s}, \overline{s}] \to [0, 1]$ to maximize

$$\int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} x(s,b) f_s(s) f_b(b) ds db$$

subject to the players' participation constraint, which is

$$0 \leq \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(\phi_2(b,1) - \phi_1(s,1)\right) x(s,b) f_s(s) f_b(b) ds db$$

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

$$\int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(1 + \lambda \left(\phi_2(b,1) - \phi_1(s,1)\right)\right) x(s,b) f_s(s) f_b(b) ds db$$

For any $\lambda \geq 0$, the Lagrangian is maximized when $x(s,b) = 1$ if and only $\left(1 + \lambda \left(\phi_2(b,1) - \phi_1(s,1)\right)\right) \geq 0$. 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

$$H(\kappa) = \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{\overline{s}} \left(\phi_2(b,1) - \phi_1(s,1)\right) x^{\kappa}(s,b) f_s(s) f_b(b) ds db$$

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 $\alpha < \kappa$, $H(\alpha)$ will differ from $H(\kappa)$ 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, $H(\kappa)$ is decreasing in κ.

To prove the $H(\kappa)$ 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 $H(\kappa)$ can be written as

$$H(\kappa) = \int_{p^A}^{\overline{b}} \int_{\underline{s}}^{h(b,\kappa)} \left(\phi_b(b,1) - \phi_s(s,1)\right) f_s(s) f_b(b) ds db$$

where $h(b,\kappa)$ is continuous in b and κ, so $H(\kappa)$ is continuous. Note also that $H(0) >= 0$, and $\lim_{\kappa \to 1} H(\kappa) = -\infty$. Therefore, there exists a unique $\kappa \in [0, 1)$ such that $H(\kappa) = 0$.

By Theorem 1 of Williams (1987), the transfer function of this mechanism will by given by (17) and (18). □

B Construction of Number of Bidders and Multiple Bids

This section describes how the number of bidders at each auction sale is estimated and how the sample with multiple bids is created. To estimate the number of bidders, I first divide the auction sale day (typically a span of three hours) into 100 equal time bins.

Each bidder receives a time stamp when arriving at the auction (or, for online bidders, when logging into the website). The data does not record when bidders exit. Therefore, to estimate the actual number of bidders present at the auction house, I assume that no bidders exit until at least one-third of the way though the auction sale day. This is consistent with my own observations from auction sales, as well as auction activity observed in a second dataset. Specifically, this second dataset records the number of bids placed by floor bidders (those physically present at the auction house), the number placed by online bidders, and the number of online bidders participating. Examining each of these quantities over the course of the auction day yields clear evidence that activity increases until about the one-third mark, at which it point it slowly declines. I measure the rate of decline using the following regression:

$$y_t = \delta_0 + \delta_1 t + u_t$$

where $t = 33, ..., 100$ represents the 100 time bins starting at the one-third point; y_t is either the number of floor bids, online bids, or online bidders; and u_t is an error term. Each of the three y_t yields an estimate of about $\hat{\delta}_1 = -0.008$. This number is then used to calculate the total number of bidders present at time bin t as:

$$N_t^{Present} = N_t^{Arrived}(1 + \hat{\delta}_1 * 1\{t > 33\})$$

where $N_t^{Arrived}$ is the total number of bidders who have arrived up to point t. $N_t^{Present}$ is then divided by the number of active lanes at time t in order to estimate the number of bidders per lane.

To construct the sample of multiple bids discussed in Section 8.1, I use a third dataset which contains information on auctions broadcast online ("simulcast" auctions). The simulcast data records all of the bids which were placed, but the identity of the a bidder is only recorded if the bidder is participating online. Therefore, multiple order statistics are distinguishable only in a small sample of auctions. For example, suppose the final four bidders recorded in the bid log for a given auction are "Floor, Floor, Simulcast 413, Floor," and the final bids are "1000, 1100, 1200, 1300." This implies that the fourth and third highest bids were placed by an unspecified floor bidders (those who are physically present), the second to last bid was placed by a specific online bidder, number 413, and the final bid was placed another floor bidder, who could have been the same as the bidder who placed the third highest bid. Therefore, while is is clear that 1300 can be interpreted as the second order static of valuations (given that the players are following the button auction model as described), it is by no means clear whether 1200 is indeed the third order statistic. Cases in which the third order statistic is clearly observable include the following:

1. The last three bidders are unique (e.g., "Floor, 413, 117").

2. The second-to-last and third-to-last, or the last and second-to-last, bidders are floor bidders and the other of the last three is a simulcast bidder.

Cases satisfying these conditions yield a sample size of 6,362 observations.

C Additional Computational Details

C.1 Optimal Weighting of Multiple Estimates of Buyer Distribution

This section describes how the distribution of buyer valuations is estimated using an optimal weighting of the estimates coming from auctions with different numbers of bidders. In practice, in each data sample, I use auction observations in which N, the number of bidders, was among the ten most common values for N observed in the sample. Depending on the sample, these ten values range from about 45 to 60.

Let $\{u_1, u_2, ..., u_M\}$ be a set of M grid points on which $F_b(b)$ will be estimated. Let K be the number of distinct subsamples of auctions where all auctions in a given subsample had the same number of bidders present. I use K=10. Let N_k be the number of bidders present in auction k, where $k = 1, ..., K$. At each grid point, u_m, equation (16) yields an estimate of $F_b(u_m)$. Call this estimate $\hat{F}_b^k(u_m)$. Let $\{\alpha_m^k : k = 1, ..., K; m = 1, ..., M\}$ be weights to be determined, and let the vectors

$$\begin{aligned}
\alpha_m &= [\alpha_m^1 \alpha_m^2 ... \alpha_m^K]' \\
\alpha &= [\alpha_1' \alpha_2' ... \alpha_M']' \\
\hat{F}_{b,m} &= \left[\hat{F}_b^1(u_m) \quad \hat{F}_b^2(u_m) \ldots \hat{F}_b^K(u_m)\right]'
\end{aligned}$$

Let $\hat{F}_b(u_m) = \alpha_m \hat{F}_{b,m}$ and let $\hat{\Omega}_m$ be an estimate of the variance-covariance matrix of $\hat{F}_{b,m}$, which I estimate through 50 bootstrap draws, and let $\hat{\Omega}$ denote a $KM \times KM$ block diagonal matrix with the Ω_m matrices along the diagonal. The weights, α_m^k, can then be determined by solving the following quadratic programming problem:

$$\begin{aligned}
&\min_{\alpha} \alpha' \hat{\Omega} \alpha \\
&\text{s.t. } \hat{F}_b(u_m) \leq \hat{F}_b(u_{m+1}), m = 1, ..., M-1 \\
&\sum_{k=1}^{K} \alpha_m^k = 1, m = 1, ..., M
\end{aligned}$$

C.2 Gauss-Chebyshev Quadrature

The counterfactual analysis in this paper requires the evaluation of a significant number of integrals, such as (10). 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.

C.3 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 et al. (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 $\phi_s(s)$ and $\phi_b(b)$ obtained by plugging in the estimated distributions and densities from Sections 5.2-5.3 evaluated at 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 $F_b(x)$. The above integration can be performed simply and accurately by choosing the grid z^S to be the Chebyshev nodes on $[\underline{s},\overline{s}]$ and performing Gauss-Chebyshev quadrature. I use $K = 200$ nodes in each dimension, s and b. See Appendix C.2 for more details. Thus,

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

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

C.4 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 6.

Theorem (from Williams 1987)**.**

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

1. *For $0 \leq \gamma < 1/2$, if $G(x^{\bar{\alpha}_1,0}) \geq 0$ for $\bar{\alpha}_1 = 1 - \gamma/(1-\gamma)$, then $(x^{\bar{\alpha}_1,0}, 0)$ is the unique solution maximizing (9) for this γ; 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 $(\alpha_2 - 1) = (\alpha_1 - 1)(1-\gamma)/\gamma$, and $(x^{\alpha_1^*,\alpha_2^*}, 0)$ is the unique solution maximizing (9) for this γ.*

2. *For $1/2 < \gamma \leq 1$, if $G(x^{0,\bar{\alpha}_2}) \geq 0$ for $\bar{\alpha}_2 = 1 + (\gamma - 1)/\gamma)$, then $(x^{0,\bar{\alpha}_2}, G(x^{0,\bar{\alpha}_2}))$ is the unique solution maximizing (9) for this γ; 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 $(\alpha_2 - 1) = (\alpha_1 - 1)(1-\gamma)/\gamma$, and $(x^{\alpha_1^*,\alpha_2^*}, 0)$ is the unique solution maximizing (9) for this γ.*

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

$$p_s(s) = q + s\int_{p^A}^{\overline{b}} x(s,b)\frac{f_b(b)}{1 - F_b(p^A)}db + \int_{s}^{\overline{s}}\int_{p^A}^{\overline{b}} x(u,b)\frac{f_b(b)}{1 - F_b(p^A)}f_s(u)dbdu \tag{17}$$

$$p_b(b) = G(x) - q + b\int_{\underline{s}}^{\overline{s}} x(s,b)f_s(s)ds + \int_{p^A}^{b}\int_{\underline{s}}^{\overline{s}} x(s,u)\frac{f_b(u)}{1 - F_b(p^A)}f_s(s)dsdu \tag{18}$$

C.5 Estimation of Equilibria in k Double Auctions

Satterthwaite and Williams (1989) demonstrated that a continuum of differentiable equilibria exist satisfying (11) and (12), and also satisfying several regularity conditions (see Satterthwaite and Williams (1989) for these conditions). They demonstrated how these equilibria can be solved for numerically in the case where both buyer and seller types have support on $[0, 1]$. Only a slight modification is required for the case of asymmetric supports.

Let p represent the bid/offer of a player and consider p to be a parameter. Write the player's inverted strategy as the player's type as a function of p, i.e. $s(p)$ and $b(p)$. Let $\dot{s} \equiv ds/dp$ and $\dot{b} \equiv db/dp$. Then $s = p^{S(-1)}(p)$, $b = p^{B(-1)}(p)$, $p^{S'}(s) = 1/\dot{s}$, and $p^{B'}(b) = 1/\dot{b}$. Substituting into (11)-(12) yields

$$\dot{s} = k\frac{\phi_s(s) - s}{b - p}$$

$$\dot{b} = (1-k)\frac{\phi_b(b) - b}{s - p}$$

$$\dot{p} = 1$$

Let $\mathcal{P} = \{(s,b,p) : \underline{s} \leq s \leq p \leq b \leq \overline{b}, s \leq \overline{s}, b \geq \underline{b}\}$. Any point in $\mathcal{P}$ has an equilibrium which passes through it. To solve for an equilibrium, pick any point in $(s,b,p) \in \mathcal{P}$ and move in the direction $(\dot{s},\dot{b},\dot{p})$ for a small step size, repeat the process until the path leaves $\mathcal{P}$. Then begin again at (s,b,p), moving in the opposite direction until the path leaves $\mathcal{P}$. This path defines an approximate solution, $p(s)$ and $p(b)$. Graphically, this procedure traces out a one-dimensional manifold in $\mathbb{R}^3$.

C.6 Estimation of Buyer Distribution When Number of Bidders is Unknown

As mentioned in Section 8.1, Song (2005) demonstrated that the density of the k^{th} order statistic conditional on the j^{th} order statistic (where $k < j$) is given by

$$f^{k|j}(y|x) = \frac{(j-1)!(F(y)-F(x))^{j-k-1}(1-F(y))^{k-1}f(y)}{(j-k-1)!(k-1)!(1-F(x))^{j-1}} \tag{19}$$

Following Kim and Lee (2008), I approximate $f^{2|3}(y|x)$ using the space of normalized Hermite polynomials, given by

$$\left\{ f_{(\Delta(J))} : f(z,\theta) = \left(\sum_{i=1}^{\Delta(J)} \theta_i H_i(z) \right)^2 + \epsilon_0 \phi(z), \sum_{i=1}^{\Delta(J)} \theta_i^2 + \epsilon_0 = 1 \right\}$$

where ϵ_0 is a small positive constant, $\phi(\cdot)$ is a standard normal pdf, $H_i(z)$ s defined recursively by

$$H_1(z) = \frac{1}{(\sigma^2\sqrt{2\pi})^{1/2}} \exp -\frac{1}{4}\left(\frac{z-\mu}{\sigma}\right)^2$$
$$H_2(z) = \frac{z}{(\sigma^2\sqrt{2\pi})^{1/2}} \exp -\frac{1}{4}\left(\frac{z-\mu}{\sigma}\right)^2$$
$$H_i(z) = \frac{1}{\sqrt{j}} z H_{i-1}(z) - \sqrt{i-1} H_{i-2}(z), \text{ for } i = 3, ..., \Delta(J)$$

The parameters μ and σ are location and scale parameters. Using (19), I form the likelihood function of the second order statistic conditional on the third and estimate the parameters μ, σ, and θ. $\Delta(J)$ is assumed to increase in J, the number of auctions. The estimates described in Section 8.1 are for $\Delta = 5$, but the approach yielded similar estimates with $\Delta = 6$, 7, or 8.

D 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 (8). 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 7: Performance of broker-revenue maximizing mechanism

	Used-car dealers sample		*Fleet/lease sellers sample*	
	Broker-optimal	Dynamic mechanism	Broker-optimal	Dynamic mechanism
Auction house revenue	0.131	0.046	0.087	0.035
	(0.002)	(0.000)	(0.003)	(0.000)
Expected gains from trade	0.304	0.388	0.244	0.334
	(0.004)	(0.007)	(0.004)	(0.003)
Buyer gains	0.084	0.256	0.070	0.203
	(0.002)	(0.007)	(0.002)	(0.003)
Seller gains	0.083	0.087	0.060	0.096
	(0.001)	(0.001)	(0.002)	(0.001)
Payment from auction house to seller	0.331	0.688	0.351	0.763
	(0.006)	(0.001)	(0.006)	(0.001)
Payment from buyer to auction house	0.467	0.745	0.466	0.799
	(0.006)	(0.001)	(0.006)	(0.001)
Probability of trade	0.362	0.715	0.377	0.774
	(0.005)	(0.002)	(0.005)	(0.002)

Notes: Estimates of welfare measures under the broker-optimal mechanism and dynamic mechanism. For the dynamic mechanism, auction house revenue comes from observed buyer and seller fees when trade occurs. The broker-optimal mechanism comes from Myerson and Satterthwaite (1983). Standard errors are from 200 bootstrap replications. Units (other than for probability of trade) are in terms of the observable value of the car.

The performance of this mechanism relative to the dynamic mechanism is shown in Table 7. The expected revenue for the auction house in the dynamic mechanism was estimated using data on fees when trade occurred. Table 7 demonstrates that the broker-optimal mechanism would result in auction house revenues of 0.131 of the market value for cars sold by used-car dealers and 0.087 of the market value for cars sold by fleet/lease sellers, between two to three times as much as under the currently used dynamic mechanism. Both buyer gains and seller gains would be lower under the broker-optimal mechanism than under the currently used mechanism, although less so for sellers.

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.