

Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry*

Bradley J. Larsen[†] and Anthony Lee Zhang[‡]

November 29, 2021

Abstract

This study quantifies bargaining power in supply-side negotiations with incomplete information, where car dealers negotiate inventory prices with large sellers at wholesale used-car auctions. We measure an agent’s bargaining power in an incomplete-information setting as the fraction of the agent’s take-it-or-leave-it-offer payoff she receives. We propose a direct-mechanism method for estimating a seller’s private value, interpreting it as the gradient of a menu from which the seller chooses her secret reserve price. We find that, on average, dealerships (buyers) have a similar degree of bargaining power as sellers. For manufacturer sellers, or sales with substantial buyer competition, sellers’ bargaining power is much higher.

Keywords: Bargaining power, auto industry, incomplete information, vertical relationships, Revelation Principle

*This paper subsumes some results from an earlier methodological working paper titled “A Mechanism Design Approach to Identification and Estimation.” We thank Yuyang Jiang and Zichen Zhao for research assistance; Susan Athey, Lanier Benkard, Tim Bresnahan, Gabriel Carroll, Evgeni Drynkin, Matt Gentzkow, Han Hong, Jakub Kastl, Simon Loertscher, Leslie Marx, Andrey Ordin, Mike Ostrovsky, Karl Schurter, Paulo Somaini, Caio Waisman, and Ali Yurukoglu for helpful comments; and participants at Duke University, the FTC, Hitotsubashi University, Nagoya University, Northwestern University, Princeton University, University of Minnesota, University of Pennsylvania, University of Texas, University of Tokyo, University of Toronto, University of Wisconsin, Stanford University, the 2016 California Econometrics Conference, the 2016 Empirics and Methods in Economics Conference, IIOC 2018, and EARIE 2020 for helpful comments.

[†]Stanford University, Department of Economics and NBER; bjlarsen@stanford.edu

[‡]University of Chicago Booth School of Business; anthony.zhang@chicagobooth.edu

1 Introduction

The division of the gains from trade between negotiating parties is of interest in many settings, such as business-to-consumer negotiations, vertical contracting relationships, the division of rents among cartel members, or estimation of patent violation damages. The party taking home a larger share is traditionally referred to as having more *bargaining power*. A number of studies over the past decade have demonstrated the importance of accounting for bargaining power when examining counterfactual policies: ignoring bargaining power—or incorrectly modeling a buyer-seller relationship as though one party has all of the power—yields misleading welfare implications. In the existing literature, bargaining power is typically assumed to be an exogenously given weight in a *complete-information* Nash bargaining framework.¹ The Nash bargaining solution, however, abstracts away from an important feature of real-world negotiations: *private information*, in which a negotiating party does not know the willingness to pay or sell of other parties. Empirical analyses of bargaining power in private/incomplete-information settings are almost nonexistent.²

In this paper we study bargaining power in the wholesale used-car industry, where parties in a vertical supply relationship negotiate under incomplete information. In this market, large fleet-owning institutions (such as banks, rental car companies, or car manufacturers) sell cars to used-car dealers. Each car trades through a mechanism of a secret reserve price ascending auction followed by alternating-offer bargaining whenever the reserve price exceeds the auction price. The data consists of over 130,000 cars offered for sale through this mechanism. We observe actions taken by negotiating pairs even for cases where bargaining ends in disagreement. This feature

¹This is the case in many empirical studies of multiple simultaneous bilateral negotiations in a Nash-in-Nash framework, e.g., [Crawford and Yurukoglu \(2012\)](#) and subsequent studies.

²In the case of patent violation damages, for example, the standard in the courts for many years was to assume that, in the absence of infringement, parties would have split surplus according to a Nash bargaining solution (typically with a 50/50 split). In recent years, courts (e.g. *VirnetX, Inc. v. Cisco Systems, Inc.*, 2014) have criticized the Nash bargaining solution as detached from reality and have demanded better ways to identify bargaining power (rather than assuming it in an ad hoc fashion), but no standard approach exists.

not only makes the setting ideal for studying bargaining power in the presence of incomplete information, but, as we discuss below, this feature is necessary in *any* setting if a practitioner hopes to distinguish between Nash bargaining and incomplete-information bargaining. With this data, we address the question of how bargaining power of buyers compares to that of sellers in the wholesale used-car market, and how this depends on seller types and competition.

Bargaining power is of particular interest in the supply side of the U.S. car market. State laws have, for decades, prohibited manufacturers from distributing new cars directly to consumers, as well as from shutting down existing dealers. The effect of these laws on the manufacturer-dealer bargaining power has been a subject of debate; the bulk of economic theory and evidence suggests these restrictions give dealers more bargaining power (see [Lafontaine and Scott Morton 2010](#) for a review). In this paper we study an aspect of this vertical relationship that involves these same key players but is not subject to these same laws: the *secondhand* car market. Our data and methodology allow us to quantify the bargaining power of dealers and wholesalers in the supply side of the secondhand market, taking into account the private information of agents in this bargaining process. Accounting for incomplete information in this analysis is critical, as inventory is sold car-by-car, and agents frequently engage in negotiations that later fail.

The term *bargaining power* has no formal (or informal) definition in incomplete-information settings. Under Nash bargaining, in contrast, the term ubiquitously refers to an agent’s share of a commonly known total surplus. To remedy this, we propose a new measure of bargaining power under incomplete-information. Our bargaining power metric is the share of an agent’s best-case—i.e., take-it-or-leave-it-offer (TIOLIO)—surplus the agent achieves relative to what the agent would achieve under the opponent’s best-case scenario. This extends a traditional (complete-information, Nash bargaining) notion of power to the incomplete-information case.³ We denote a

³In Nash bargaining, an agent’s *share of the total surplus* and *share of her TIOLIO payoff* are

buyer’s bargaining power by α_B and seller’s by α_S . In the seller TIOLIO mechanism, $\alpha_S = 1$ and $\alpha_B = 0$, and in the buyer TIOLIO mechanism, $\alpha_B = 1$ and $\alpha_S = 0$. Any intermediate values are possible, as are negative weights. Unlike in Nash bargaining, under incomplete information the sum of these weights can be greater than 1: sellers and buyers can collectively achieve strictly greater expected utility than that available through any convex combination of TIOLIO mechanisms. The sum of α_B and α_S is informative about the *efficiency* of trade. These weights are thus a natural generalization of bargaining power to asymmetric information settings, giving information both about how the pie is split and also the size of the pie itself. This relationship between bargaining power and the size of the pie is a key point ignored by Nash bargaining.⁴

Next, we show how to estimate bargaining power under incomplete information. The bargaining theory literature shows that incomplete information gives rise to complications, such as multiple equilibria, delay, and inefficiency. [Ausubel et al. \(2002\)](#) highlight that different equilibria can have quite different properties and outcomes, and that no complete characterization of equilibria exists; this statement remains true twenty years later. As such, there is no off-the-shelf model for empiricists to bring to bargaining data to identify players’ private value distributions, unlike the now well-developed empirical literature on auctions (e.g. [Guerre et al. 2000](#)).⁵ Our paper offers a first step to addressing bargaining power empirically under incomplete

equivalent notions. Under incomplete information, however, these two notions are distinct. When both parties have private values on overlapping supports, the surplus available to the pair is unknown to both parties, and neither is able to extract the full surplus ([Myerson and Satterthwaite 1983](#)).

⁴[Loertscher and Marx \(2021\)](#) state this point as follows: “The complete information approach with efficient bargaining has the downside that shifts of bargaining power ... only affect the distribution of surplus and not its size since bargaining is, by assumption, efficient.”

⁵Unlike auction theory, where clean equilibrium results exist for settings suitable for empirical work, such as continuous values and incomplete information, bargaining theory is not immediately portable to empirical analysis. Several previous theoretical bargaining papers analyze an environment close to the environment we study—with continuous values, two-sided offers, and two-sided incomplete information about players’ values—but the equilibria derived in these studies are not suitable for structural estimation in our setting. For example, in [Perry \(1986\)](#) the game ends immediately and in [Cramton \(1992\)](#) at most two serious offers occur in equilibrium; neither of these possibilities can fully explain observations in our data.

information, focusing on the supply side of the U.S. used-car market.

In the wholesale used-car market, the primary challenge to identification is the distribution of *seller* values, F_S . Every choice of the seller—even the seller’s choice of secret reserve price in the pre-bargaining stage of the game—depends on the equilibrium of the post-auction bargaining subgame, and these equilibrium strategies are unknown to the econometrician. In contrast, the distribution of buyer values, F_B , can be identified from buyers’ auction bids using existing tools from the auction literature. These tools also allow us to handle game-level observable and unobservable heterogeneity.

We propose to estimate seller values based on an *empirical menu* approach. We show that the analyst can think of a seller of value v_S as choosing her secret reserve price, r , to maximize her expected payoff $v_S P_S(r) - T_S(r)$, where $P_S(r)$ is the seller’s expected probability of keeping the car and $T_S(r)$ is the expected transfer. Our identification argument is simple: a seller’s choice of reserve price is a choice from a convex equilibrium menu of possible (P_S, T_S) pairs, and the derivative of this menu, evaluated at the seller’s choice, corresponds precisely to that seller’s value. The data requirements to identify a seller’s value are observations of (i) the secret reserve price, (ii) the final allocation (i.e. an indicator for whether trade occurs), and (iii) the final payment. With these variables in hand, the objects $P_S(\cdot)$ and $T_S(\cdot)$ are essentially observed in the data, and derivatives of this menu correspond to agents’ values.

We apply these arguments to our data by estimating the trade-transfer menu faced by sellers in the wholesale used-car market. Our model implies two testable restrictions. First, the equilibrium menu must be convex. Second, the menu must satisfy individual rationality constraints for all agent types who participate. We impose both restrictions and find that they are not overly strong in our setting. With the estimated menu and distribution of values, we compute bargaining power. We find that car dealers (who are *buyers* in this supply-side market) exert a similar level of bargaining power as the large institutional sellers they purchase from: buyers achieve

a level of surplus that is 64% of the way between their TIOLIO payoff and what they would receive under the seller’s TIOLIO mechanism, whereas sellers’ surplus is only 62% of the way along their corresponding continuum.

We then decompose our results according to different seller categories, such as manufacturers (e.g., Ford, GM, or Chrysler), banks, fleet companies, or rental companies. We find that manufacturer sellers have substantially more bargaining power than buyers, achieving an outcome that is over 90% of what they would receive if they were to have all the bargaining power, and buyers at these sales have near-zero bargaining power (only 4% of their maximal payoff). At least part of this difference is explained by the fact that *competition* among dealerships (buyers) is much higher at manufacturer sales.

As highlighted in [Loertscher and Marx \(2019\)](#), how competition and bargaining power interact in settings with incomplete information is an open question of interest to antitrust and competition authorities. The empirical literature has studied the relationship of bargaining power to competition under assumptions of Nash bargaining (e.g. [Gowrisankaran et al. 2015](#)), but not under incomplete information. One would expect increased competition among buyers to increase the seller’s bargaining power, but it is unclear by how much. The seminal results of [Bulow and Klemperer \(1996\)](#) suggest that a seller would prefer increased competition to increased bargaining power, but this interpretation abstracts away from real-world negotiations, in which buyers may have some power. Our results suggest that, on average, buyers have a similar level of power to sellers in supply-side negotiations for used cars, but seller bargaining power increases drastically at high levels of buyer competition.

Our study relates to a growing body of structural work studying bargaining power in business-to-business settings, such as [Crawford and Yurukoglu \(2012\)](#), [Grennan \(2013\)](#), [Gowrisankaran et al. \(2015\)](#), and [Ho and Lee \(2019\)](#). We also contribute to the literature analyzing aspects of the vertical relationship between dealers and wholesalers in the automotive industry. [Lafontaine and Scott Morton \(2010\)](#) summa-

alize this literature and point to evidence that the current sea of state laws governing dealer-manufacturer relationships benefits dealers at the expense of manufacturers (and ultimately consumers). The implications of these laws is a key topic of interest for the Federal Trade Commission in recent years.⁶ [Murry and Zhou \(2020\)](#) analyze the effects in this market of manufacturers terminating dealer locations. [Donna et al. \(2021\)](#) study bargaining in vertical relationships in a separate industry (outdoor advertising), but discuss how direct-to-consumer sales in the auto industry, such as Tesla’s, could alter welfare in this market.

In contrast to previous work on vertical relationships, we study bargaining power without assuming complete information. We allow for agents to have private information about their willingness to pay and sell (and hence, incomplete information about their opponent’s value) and to be strategic in their bargaining behavior.⁷ Several structural studies of bargaining do allow for incomplete information.⁸ [Keniston \(2011\)](#) studies the question of whether welfare is higher under bargaining or a posted-price mechanism. [Larsen \(2021\)](#) analyzes some of the same used-car data we study, but focuses on the empirical implications of the main theorem of [Myerson and Satterthwaite 1983](#) and how efficient bargaining is relative to the theoretical second-best. [Freyberger and Larsen \(2021\)](#) study efficiency and impasse in bargaining on eBay.

We see our focus on *equity*—how the surplus is split—as a natural next question to address after efficiency. [Larsen \(2021\)](#) and [Freyberger and Larsen \(2021\)](#) derive partial identification results that yield bounds on surplus or trade probabilities, but do not address the question of surplus division. Indeed, these bounds, while informative

⁶See <https://www.ftc.gov/news-events/blogs/competition-matters/2015/12/ftc-opens-hood-automobile-distribution>.

⁷It is important to note that our approach is not a strict generalization of many complete-information (Nash or Nash-in-Nash) bargaining approaches. In particular, we specify agents’ bargaining surplus as quasilinear in price, whereas some complete-information studies of vertical bargaining allow the downstream firm to have a willingness to pay that depends on the price negotiated with the upstream firm, for example.

⁸A related theoretical study to ours is [Loertscher and Marx \(2021\)](#). The authors allow for incomplete information and propose measuring bargaining power as an agent’s weight in a weighted welfare maximization problem. Our definition instead quantifies an agent’s payoff relative to her TIOLIO payoff.

about inefficiency, are too wide to be informative about surplus division. In contrast, in this paper we obtain point estimates of this split. In doing so, we borrow some of the straightforward steps of [Larsen \(2021\)](#), including how we control for game-level heterogeneity and how we estimate the distribution of buyer values from auction prices, which are both tools from the auction literature. Our identification argument for seller valuations differs from that of [Larsen \(2021\)](#): we exploit optimality of the seller’s choice of secret reserve price, yielding point identification of the seller value CDF, whereas the former study exploits the seller’s choice to accept or reject the auction price, yielding only partial identification.

Our contribution to the structural methodology literature can be seen as generalizing the [Guerre et al. \(2000\)](#) first-price-auction method to bargaining games. In a related, contemporaneous study complementary to ours, [Kline \(2017\)](#) focuses on identification, but not estimation, in a class of games that overlaps with the class we study: trading games with monotone equilibria.⁹ As we emphasize in [Section 4](#), our identification results largely only require taking a stance on the structure of agents’ utility functions, not the specific rules of the game being played, and thus may be particularly valuable for studying bargaining, where researchers may observe negotiated prices without being able to fully characterize the equilibrium of the game generating those prices. In this sense, our work is an empirical analog of the theoretical mechanism design approach to bilateral bargaining (e.g. [Myerson and Satterthwaite 1983](#); [Williams 1987](#); [Loertscher and Marx 2021](#)), which abstracts away from extensive-form details.

⁹Related arguments are also used in [Perrigne and Vuong \(2011\)](#) and [Luo et al. \(2018\)](#). [Pinkse and Schurter \(2019\)](#) introduce efficient estimation procedures for auctions and related games, which, like ours, exploits convexity restrictions implied by bidders’ incentive compatibility conditions.

2 Background: Supply-Side Bargaining for Used Cars

The wholesale used-car industry—an industry with revenues above \$100 billion annually in the United States—operates through a network of several hundred auction house locations scattered throughout the country (and operations are similar internationally).¹⁰ These auction houses have been a part of the US used-car market for over seventy years. Over 15 million cars pass through auto auction houses annually. At each auction house, used-car dealers buy cars from large fleet companies, such as rental companies, banks with repossessed vehicles, or manufacturers with lease-buyback vehicles.¹¹ Sales at a given auction house typically take place once a week. A seller brings her car to the auction house several days before the sale and reports a secret reserve price to the auctioneer. On the day of the sale, buyers (used-car dealers) arrive, with many traveling long distances to attend. Remote bidders also participate virtually, watching the auction and bidding online. Cars are auctioned in the order they arrive, with multiple auctions running simultaneously in different lanes that divide the building where sales occur.

The mechanism proceeds as follows: buyers participate in an ascending auction, indicating their willingness to pay the current price, with the bidding controlled by a human auctioneer who raises the price until only one bidder remains. The auction itself takes about 90 seconds (Lacetera et al. 2016). If the final auction price exceeds the secret reserve price (observed by the auctioneer but not the buyers), the high bidder takes the car. If not, the high bidder and seller enter alternating-offer bargaining, mediated by an auction-house employee over the phone.¹² If she chooses, the high bidder may opt out of bargaining before it begins.

Our data consists of 131,443 realizations of this mechanism from six auction houses

¹⁰https://www.naaa.com/pages/Auction_Industry_Survey/2021_Survey_MtIs/NAAA_2020_Industry_Survey_Slides.pdf.

¹¹Used-car dealers also operate as sellers in this marketplace. The data we use in this study comes only from sales of large fleet and lease companies. See Appendix B.4, as well as Larsen (2021), for an analysis of data from used-car dealer sales.

¹²For an analysis of these mediators, see Larsen et al. (2021).

Table 1: Descriptive Statistics (Sample Size = 131,443)

| A. | Mean | Standard Deviation | Seller Category | Fraction of Sample |
|----------------|--------|--------------------|-----------------|--------------------|
| Blue Book (\$) | 10,951 | 6,144 | Manufacturer | 0.1958 |
| Age (years) | 3.18 | 2.55 | Bank | 0.5423 |
| Mileage | 57,481 | 40,389 | Fleet Company | 0.0751 |
| Good Condition | 0.72 | 0.45 | Lease Company | 0.1143 |
| Num. Bidders | 25.99 | 14.71 | Rental Company | 0.0725 |

| B. | Conditional on Sale | | | | | Cond. on No Sale | |
|-----------------|---------------------|-------------|--------------------|--------------------|------------------|--------------------|--------------------|
| | Frac. of Sample | Frac. Agree | Auction Price (\$) | Reserve Price (\$) | Final Price (\$) | Auction Price (\$) | Reserve Price (\$) |
| End at Auction | 0.34 | 0.98 | 11,063 | 10,246 | 11,063 | 5,367 | 7,622 |
| Period 2 | 0.56 | 0.74 | 9,683 | 10,586 | 9,683 | 9,547 | 11,002 |
| Period ≥ 3 | 0.10 | 0.16 | 7,416 | 8,763 | 7,869 | 6,352 | 8,338 |

Notes: In panel A, “Blue Book” is an estimate of the car’s market value, provided by the auction house. “Good Condition” indicates average or above average car condition, based on auction house inspection. “Number of bidders” is an upper bound on the number of bidders, only observable in the bid log subsample (102,186 observations). “Seller Category” refers to type of company the seller is. Panel B shows statistics separately for games ending at the auction (through the auction price exceeding the reserve, or the buyer refusing to negotiate), games where the seller accepts or rejects the auction price (indicated by Period 2), or games ending after further bargaining (Period ≥ 3). Panel B shows average auction and reserve price separately for games ending in agreement/disagreement, and average final price for those ending in agreement.

from 2007–2010. For each realization, the primary variables we observe are the secret reserve price, final transaction price, final allocation (i.e. an indicator for whether the car sold), and auction price. We also observe a large set of characteristics, including features of the car and the auction house environment at the sale time.

Table 1 shows descriptive statistics. The average car has a blue book value (an estimate provided by the auction house) of \$10,951, is 3.18 years old (relative to its model-year), and has 57,481 miles on the odometer. The auction house provides a condition report for most cars, and 72% of cars are rated at average quality or above, which we indicate in panel A with “Good Condition.” Manufacturers, such as Ford, GM, and Chrysler, represent 20% of sellers in our data. Banks, such as Citibank or Bank of America, represent a slight majority, at 54%. Fleet companies (such as Wheels) represent 7.5%, rental companies (such as Budget Rental Car) represent a similar percentage, and lease companies represent 11%. Our data also contains detailed records (referred to as *bid logs*) of the bidding during the auction stage for

most observations (102,186). In this sample, we obtain bounds on the number bidders (N) in each auction, with an average upper bound of 26; see Section 5.1 and Appendix B.1 for details.

Approximately 30% of attempts to sell cars result in no trade. This large portion is inconsistent with a standard complete-information framework: under complete information, a buyer and seller would not engage in a trading game knowing a priori that they will disagree. Failed negotiations, however, are completely consistent with the presence of incomplete information (Myerson and Satterthwaite 1983; Perry 1986).

Panel B of Table 1 breaks down outcomes by how the game ends—with a sale (agreement) or no sale (disagreement). We report the primary variables that are required for our identification and estimation: the seller’s secret reserve price, final transaction price, final allocation, and auction price. The first row shows outcomes for games that end with no bargaining, which occurs in 34% of cases. In these cases the game either ends with the auction price exceeding the reserve price or with the buyer opting out of bargaining (which occurs 2% of the time). The second row, indicated by *Period 2*, refers to cases where the first action occurs that can be considered *bargaining*: the auction price falls below the reserve price, and the seller either accepts (74% of the time) or rejects (26% of the time) the auction price.¹³ The third row refers to games that end at some later period of the bargaining game, which occurs in 10% of the sample. When the game ends with a sale at the auction or in period 2, the final price naturally equals the auction price. When the game ends in a sale at a later stage of the game, the average auction price is \$7,416, the average reserve price is \$8,763, and the average final price is between the two, at \$7,869. When trade fails (the final two columns), the auction price is farther below the reserve price.

These final numbers illustrate an important point: it is a priori unclear how

¹³In this paper, we do not explicitly address the puzzle that sellers who end up accepting auction prices below their reserve prices could have potentially achieved that outcome by simply setting a lower reserve price upfront. Larsen (2021) and Goke (2021) offer some explanations for this puzzle.

to think of *bargaining power* in this context. It may be tempting to interpret the location of the final price relative to the auction and reserve prices as an indication of bargaining power. But this logic is flawed: a buyer's true value will be weakly higher than the auction price and a seller's weakly lower than the secret reserve price. These bounds say nothing about how the pie is split or what its size is; they do not rule out the possibility that the buyer's value is ∞ and the seller's is 0, for example, making it impossible to make inferences about bargaining power from these bounds alone. Our identification argument allows us to infer the distribution of buyer values from auction prices and seller values from reserve prices. From these primitives and trade outcomes we then quantify bargaining power.

3 Defining Bargaining Power Under Incomplete Information

Here we introduce our notion of bargaining power. Consider a seller with value v_S and buyer with value v_B who bargain over an indivisible good. The game in the wholesale used-car market is in fact a game between one seller and many buyers, but the mechanism boils down to bilateral trade between the seller and just the high bidder, as the auction serves to identify the highest-value buyer. We describe the full auction-plus-bargaining game in more detail in [Section 4](#).

Equilibria of a bilateral bargaining game under incomplete information can be complex to characterize theoretically, even for simple extensive forms such as alternating offers. This is because each offer signals information to the opposing party, who can then update her beliefs about the opponent's value. Belief updating following off-equilibrium offers can be used to sustain a large set of strategies in bargaining (see discussions in [Gul and Sonnenschein 1988](#) and [Ausubel et al. 2002](#)).

Rather than attempting to characterize equilibria of a given extensive form, we take a mechanism design approach. By the revelation principle ([Myerson 1979](#)), any equilibrium of a bilateral bargaining game has a corresponding *direct revelation*


mechanism made up of an allocation function describing the probability with which types v_S and v_B trade in equilibrium and a transfer function describing the expected transfer from the buyer to the seller. Let $\mathcal{M}(v_S, v_B)$ represent a particular mechanism.

Let $U_B(\mathcal{M})$ and $U_S(\mathcal{M})$ represent the expected surplus of the buyer and seller, respectively, under bargaining mechanism \mathcal{M} , where the expectation is taken over buyer and seller values; thus, $U_B(\mathcal{M})$ and $U_S(\mathcal{M})$ represent *ex-ante* surplus, in the terminology of [Holmström and Myerson 1983](#). [Williams \(1987\)](#), building on [Myerson and Satterthwaite \(1983\)](#), derives the Pareto frontier of bargaining mechanisms: the set of the *highest* possible combinations of buyer and seller surplus achievable by an incentive-compatible, individually rational, budget-balanced mechanism. This frontier is a convex function maximizing the weighted sum of welfare, $\eta U_S(\mathcal{M}) + (1 - \eta)U_B(\mathcal{M})$ for $\eta \in [0, 1]$. We illustrate this with the concave green line in [Figure 1](#). This welfare weight, η , might reasonably be thought of as one notion of bargaining power *among ex-ante efficient mechanisms*, but this notion would not be sufficient for our purposes; we seek a notion of bargaining power that can be applied to real-world bargaining situations, which will not necessarily correspond to points on the frontier.

Indeed, the endpoints ($\eta = 1$ or $\eta = 0$) are the only points on frontier known to be achievable by practical mechanisms in a general two-sided-uncertainty game. These endpoints consist of a TIOLIO by one party or the other. All other mechanisms along the frontier are, from a practitioner’s perspective, complicated black boxes, and are not necessarily achieved by any practical bargaining protocol, including the alternating-offer protocol of used-car markets.

Any real-world bargaining mechanism yields an expected buyer and seller surplus somewhere within this frontier, such as the point labeled “*” in [Figure 1](#). Our notion of bargaining power describes how good this outcome is for agents relative to what they would achieve if they could instead make a TIOLIO to the opposing party. We define the buyer’s bargaining power, α_B , in a given mechanism to be the fraction of the buyer’s TIOLIO surplus she receives relative to what she would receive if the

Figure 1: Illustration of Bargaining Power Measure



Notes: Visualization of bargaining power measures. The x-axis and y-axis are, respectively, the ex-ante expected utility of sellers and buyers. Our measures of bargaining power, α_B and α_S , represent points on a coordinate system where $(U_S(\mathcal{M}^1), U_B(\mathcal{M}^1))$ corresponds to $(\alpha_S = 1, \alpha_B = 0)$ (the seller's and buyer's expected payoff when the seller makes a TIOLIO), and $(U_S(\mathcal{M}^0), U_B(\mathcal{M}^0))$ corresponds to $\alpha_S = 0, \alpha_B = 1$. The point labeled "*" is an example of an arbitrary mechanism, mapping to $\alpha_S = 0.5, \alpha_B = 0.333$. The green line represents the second-best frontier. The red dashed line traces out mechanisms that are convex combinations of the seller- or buyer-optimal mechanisms.

seller were to instead make a TIOLIO. For the seller, α_S is defined similarly.

To define this more precisely, let \mathcal{M}^η for $\eta \in [0, 1]$ be a mechanism along the Pareto frontier. The mechanism \mathcal{M}^0 corresponds to the buyer TIOLIO case and \mathcal{M}^1 to the seller TIOLIO case. The buyer's payoff if the *buyer* makes a TIOLIO is then $U_B(\mathcal{M}^0)$, and the buyer's payoff if instead the *seller* makes a TIOLIO is $U_B(\mathcal{M}^1)$. We define the buyer's bargaining power α_B for any *arbitrary* bilateral trade mechanism \mathcal{M} as the weight in the convex combination $U_B(\mathcal{M}) = \alpha_B U_B(\mathcal{M}^0) + (1 - \alpha_B) U_B(\mathcal{M}^1)$. Similarly, α_S is the weight in $U_S(\mathcal{M}) = \alpha_S U_S(\mathcal{M}^1) + (1 - \alpha_S) U_S(\mathcal{M}^0)$.

The weights α_B and α_S are conveniently analogous to standard Nash bargaining weights applied in an incomplete-information world. Specifically, in Nash bargaining, a player with bargaining power 1 would receive her payoff from making a TIOLIO, and a player with bargaining power 0 would receive her payoff from the opponent making the TIOLIO, just as in our incomplete-information notion of bargaining power. However, unlike a complete-information setting, under incomplete information, a party's

TIOLIO payoff does not correspond to getting *all* of the surplus; some information rent is left for the player *receiving* the TIOLIO; this is the well-known efficiency/rent-extraction trade-off occurring in settings with incomplete information.

The set of mechanisms for which $\alpha_B + \alpha_S = 1$, illustrated with the red dashed line in Figure 1, is of particular interest. The surplus divisions implemented by these mechanisms can be achieved by randomly selecting one party to make a TIOLIO (what may be termed a *random ultimatum game*).¹⁴ For example, if $\alpha_B = 0.6$ and $\alpha_S = 0.4$, the utility outcomes are equal in expectation to what would arise in an ultimatum game that selects the buyer as the proposer with probability 0.6. The reason this mechanism is of interest is that it is simple to implement: the market designer simply randomly selects the buyer or seller as a TIOLIO proposer.

Unlike Nash bargaining, we do not require $\alpha_B + \alpha_S = 1$, as incomplete-information bargaining can yield combinations of buyer and seller surplus above or below the dashed line. The sum $\alpha_B + \alpha_S$, relative to 1, gives us some information about the efficiency of bargaining. Any mechanism for which $\alpha_S + \alpha_B < 1$ —those to the south-west of the red dashed line—are Pareto dominated in expectation by points lying on the $\alpha_B + \alpha_S = 1$ line. Relative to such a mechanism, a simple random ultimatum mechanism would achieve higher ex-ante expected utility for both players. Any mechanism with $\alpha_S + \alpha_B > 1$ is one that works relatively well compared to this set of simple-to-characterize mechanisms.

Computing α_B and α_S and making corresponding statements about bargaining power in a given real-world market requires knowledge of two key objects: the distributions of buyer values F_B and seller values F_S . These two primitives are therefore the primary goals of the identification and estimation results we describe below.

¹⁴This mechanism is also known as *final offer arbitration* or *baseball arbitration*, as it is used in salary negotiations of Major League Baseball contracts.

4 Identifying Negotiators' Private Value Distributions

We now offer a general model of the wholesale used-car market game. The game involves 1 seller and N potential buyers, where N is a random variable varying across instances of the game. For a given realization of $N = n$, we index the seller and buyers by $i \in \{S, B_1, \dots, B_n\}$. Agents are risk neutral and have values $V_i + Y\beta + W$, where $W \sim F_W$, $V_S \sim F_S$, and $V_i \sim F_B$ for all $i \in \{B_1, \dots, B_n\}$; thus, we impose bidder symmetry. Random variables W and V_i have bounded support.

The random variables V_i for $i \in \{S, B_1, \dots, B_n\}$ are players' private values.¹⁵ In our empirical analysis, Y is a vector observable to the econometrician and to players (such as the make and model of the car), whereas W represents game-level heterogeneity observed by the agents but not the econometrician, such as a dent or odor in the car. We assume $\{Y, W, N, V_i\}$ are mutually independent for all i in a given instance of the game.¹⁶ Thus, agents' overall values are correlated through game-level heterogeneity terms Y and W , but, conditional on the realizations of these terms, values are independent. To simplify exposition, we now condition on a realization of $N = n$ and return to this in Section 4.2. Likewise, we condition on a realization of $W = w$ and $Y = y$, omitting these variables from the discussion until Section 4.3.

The used-car auction/bargaining game begins with the seller choosing a secret reserve price, R . The bidders then participate in an ascending (button) auction. If the auction price, P^A , exceeds the reserve price, the high bidder and seller trade at price P^A . If $P^A < R$, the high bidder is given the opportunity to exit the game. If he chooses not to exit, the high bidder and seller engage in alternating-offer bargaining for up to τ periods, where we assume τ is finite but may be large. If the game ends with the seller and bidder i trading at some price P , the buyer receives a payoff of

¹⁵In our setting, heterogeneity in agents' private values arises from differences in location and inventory needs.

¹⁶To see the importance of this assumption, note that if we were to instead assume independence of N conditional W , this would allow for correlation between W and the auction price (because this price depends on N), and the game-level heterogeneity convolution arguments we invoke (akin to those used in [Krasnokutskaya 2011](#)) would not immediately apply.

$V_i - P$ and the seller receives a payoff of P , less any bargaining costs. If the two parties do not come to agreement, the seller receives a payoff of V_S (her value for keeping the car) and the buyer receives a payoff of 0, less any bargaining costs. We do not model costs explicitly, but assume that (i) they are nonzero, such that no party would bargain if it were common knowledge that no gains from trade exist, and (ii) they are small enough to be negligible from an estimation perspective.¹⁷

We focus on a restricted class of pure strategy Bayes Nash Equilibria (BNE). [Larsen \(2021\)](#) demonstrates the existence of such equilibria and proves several properties that we state here and exploit in identification/estimation:

Proposition 1. *Pure strategy BNE exist in which the following two restrictions hold: (i) the seller’s payoff in the post-auction bargaining subgame is continuous in P^A and (ii) each bidder drops out of the auction only when the auction price equals her value.*

In such equilibria, the following properties hold: (iii) secret reserve prices are strictly increasing in sellers’ values; (iv) auction, secret reserve, and final prices are additively separable in (and the probability of trade is invariant to) game-level heterogeneity; and (v) each seller type trades with all buyer types above a certain cutoff.

As highlighted in Section 3, the key objects to identify to evaluate bargaining power are F_S and F_B . With these objects, we can compute the expected TIOLIO payoff for each player and determine what fraction of these quantities each player receives in practice. The novel part of our identification is that of F_S . We dedicate Section 4.1 to this endeavor. Identification of F_B and incorporating game-level heterogeneity, on the other hand, rely largely on Proposition 1 and on prior results from the auction literature. We discuss these arguments in Sections 4.2–4.3.

¹⁷[Larsen \(2021\)](#) estimates bounds on expected bargaining costs in this market and finds that they are indeed negligible, (less than \$34 the buyers and \$5 for sellers).

4.1 Equilibrium Menus

A pure strategy BNE in this game is a complicated object. To describe it, let ζ_i to be the drop-out price of buyer i in the auction. Let $D_t^B \in \{A, C, Q\}$ represent the buyer's decision to accept, counter, or quit in odd periods t . Let P_t^B represent the buyer's counteroffer (if the buyer counters) in period t . Let D_t^S and P_t^S be defined similarly for even periods t . Let H_t represent the history of publicly observed actions up through period $t - 1$ of the game. These actions include the auction price and all previous bargaining offers and period-specific decisions.

The strategy of a buyer of type v_{B_i} is a history-contingent set of actions $\sigma_i^B(v_{B_i}) = \{\zeta_i, \{D_t^B|H_t\}_i, \{P_t^B|H_t\}_i\}$, where the decisions D_t^B and offers P_t^B included are those for periods in which it is the buyer's turn. The strategy of a seller of type v_S is a history-contingent set of actions $\sigma^S(v_S) = \{\rho, \{D_t^S|H_t\}, \{P_t^S|H_t\}\}$, where $\rho(v_S) = r$ is the seller's reserve price strategy, and where the decisions and offers are those for periods in which it is the seller's turn.¹⁸ A set of strategies $\{\sigma^{B*}(v_{B_i})\}_i^N$ for all buyers and $\sigma^{S*}(s)$ for the seller is a BNE if, for each player, her strategy is a best response to opponents' strategies and players update their beliefs about opponent values using Bayes rule at each history of the game reached with positive probability.¹⁹ We assume that the econometrician has access to data generated by a single such BNE.

Rather than working with this full set of strategies, we follow the mechanism design literature and analyze the game as a direct mechanism. Here we introduce explicit notation for the allocation function and transfer function, the two components of any mechanism \mathcal{M} . In a direct mechanism, each agent, $i \in \{S, B_1, \dots, B_n\}$, reports (or potentially misreports) her private value to a mechanism designer, who assigns

¹⁸Note that, for simplicity, we do not allow for equilibria that explicitly depend on N , W , or the drop-out prices of other bidders (other than the auction price).

¹⁹We do not impose any refinement, such as Perfect Bayes Equilibrium (PBE). These refinements have been shown to do little or nothing to restrict the set of equilibria in sequential bargaining games (see [Gul and Sonnenschein 1988](#) and discussions in [Ausubel et al. 2002](#)). This is because, even in a PBE, a large array of behavior can be sustained by carefully chosen beliefs following off-equilibrium actions. Importantly, however, all of our identification arguments apply regardless of whether we focus on BNE or a refinement such as PBE.

allocations according to a function $\mathbf{x}_i(v_S \dots v_{B_n})$. This function is equal to 1 for the agent allocated the car and zero for others. The mechanism designer allocates net payments made by i according to the function $\mathbf{t}_i(v_S \dots v_{B_n})$.

The expected outcome for a given agent can be described by *menus* of probability-transfer pairs. If player i behaves as if she is a type v'_i (potentially misreporting her type), she attains an expected outcome $(P_i(v'_i), T_i(v'_i))$, defined as

$$P_i(v'_i) \equiv E[\mathbf{x}_i(v'_i, V_{-i})], \quad T_i(v'_i) \equiv E[\mathbf{t}_i(v'_i, V_{-i})].$$

$P_i(v'_i)$ and $T_i(v'_i)$ are, respectively, the expectation of i 's allocation $\mathbf{x}_i(v'_i, V_{-i})$ and transfer $\mathbf{t}_i(v'_i, V_{-i})$, over values V_{-i} of other players $-i$, which are random variables from i 's perspective.

The expected utility of i when she has value v_i but plays as if it were v'_i is

$$v_i P_i(v'_i) - T_i(v'_i). \quad (1)$$

In any incentive compatible (IC) mechanism, v_i (i 's true value) maximizes (1):

$$v_i P_i(v_i) - T_i(v_i) \geq v_i P_i(v'_i) - T_i(v'_i) \quad \forall v'_i. \quad (2)$$

These IC conditions offer immediate bounds on the value of agent i :

Theorem 1. *For any agent i , v_i must satisfy*

$$\begin{aligned} v_i &\geq \frac{T_i(v_i) - T_i(v'_i)}{P_i(v_i) - P_i(v'_i)} \quad \forall v'_i : P_i(v'_i) < P_i(v_i) \\ v_i &\leq \frac{T_i(v'_i) - T_i(v_i)}{P_i(v'_i) - P_i(v_i)} \quad \forall v'_i : P_i(v'_i) > P_i(v_i). \end{aligned}$$

Proof. Follows immediately from (2). □

In any BNE in which i 's strategy involves an action that is *one-to-one* with her

value, v_i , Theorem 1 can be restated in terms of that action, rather than in terms of types v'_i that i could mimic. In our game, by Proposition 1.iii, the seller's secret reserve price is such an action: in any BNE, $r = \boldsymbol{\rho}(v_S)$ (the secret reserve price function) is strictly increasing in v_S . We combine Theorem 1 with this strict monotonicity property to obtain the following corollary specific to the seller's value; as such, we state it only for $i = S$:

Corollary 1. *If, in equilibrium, $r = \boldsymbol{\rho}(v_S)$ is one-to-one with v_S , v_S must satisfy*

$$v_S \geq \frac{T_S(r) - T_S(r')}{P_S(r) - P_S(r')} \quad \forall r' : P_S(r') < P_S(r) \quad (3)$$

$$v_S \leq \frac{T_S(r') - T_S(r)}{P_S(r') - P_S(r)} \quad \forall r' : P_S(r') > P_S(r). \quad (4)$$

Note that Corollary 1 adopts a slight modification of notation, which we will maintain moving forward, in which we write P_S and T_S as functions of r directly rather than v_S . This is without loss of generality, as r is one-to-one with v_S , and is less cumbersome than writing P_S and T_S as functions of $v_S = \boldsymbol{\rho}^{-1}(r)$. We also write the seller's allocation and transfer functions as $\mathbf{x}_S(r, V_{-S})$ and $\mathbf{t}_S(r, V_{-S})$, respectively.

Figure 2, in the left panel, illustrates a hypothetical equilibrium menu faced by the seller. We consider a case where the seller's possible choices of secret reserve prices are $r' \in \{r^1, \dots, r^5\}$. Indifference curves in this figure are straight lines, with the seller's utility being higher toward the southeast of the figure. To interpret, consider a seller choosing r^3 . Compared to r^3 , points r^4 and r^5 have higher probability $P_S(r)$ of the seller keeping the good and higher transfer $T_S(r)$ from the seller to buyer.²⁰ If the seller prefers r^3 to r^4 or r^5 , her value must be lower than the average cost of *purchasing* this additional probability. That is, $v_S \leq \frac{T_S(r') - T_S(r^3)}{P_S(r') - P_S(r^3)}$ for items $r' \in \{r^4, r^5\}$.

Similarly, compared to point r^3 , r^1 and r^2 have lower transfers and lower prob-

²⁰Here we follow the mechanism design literature in modeling the expected transfer an agent makes, rather than receives. For the seller, this can be counterintuitive: a higher transfer from the seller to the buyer is in fact a lower payment received by the seller, as all payments flow from buyer to seller in practice. The second panel of Figure 2 provides an alternative illustration.

ability of keeping the good. If the seller prefers point r^3 , her value must be higher than the average cost of purchasing the additional probability offered by r^3 . That is, $v_S \geq \frac{T_S(r^3) - T_S(r')}{P_S(r^3) - P_S(r')}$ for items $r' \in \{r^1, r^2\}$. Thus, the bounds in (3) and (4) imply that the value of any seller type choosing r^3 must lie between the slopes of the blue lines labeled $\underline{v}(r^3)$ and $\bar{v}(r^3)$ in Figure 2.

Alternatively, equilibrium menus faced by the seller can also be thought of in terms of the probability that the seller sells the car and the expected revenue received by the seller. The right panel of Figure 2 shows the equilibrium menu in terms of these quantities. Sellers effectively choose a probability-revenue pair from a concave menu: when sellers set lower reserve prices, they sell more often and get higher expected revenues, but *marginal* revenue from selling with higher probability is lower with lower reserve prices. If we observe a seller choosing r^3 , the secant lines on the menu at r^3 bound the seller's value for the car.

This leads to our main (and most powerful) result: for a sufficiently smooth and continuous game, the bounds in Theorem 1 collapse to yield point identification.

Corollary 2. *Suppose $P_S(r)$ is continuous and strictly monotone with derivative bounded away from 0; $P_S(r), T_S(r)$ are both continuously differentiable; and $\rho(v_S)$ is continuous and strictly monotone in v_S . Then $\mathbf{v}_S(r)$, the inverse of $\rho(v_S)$, satisfies $\mathbf{v}_S(r) = \frac{T'_S(r)}{P'_S(r)}$ for all reserve prices r played by some type v_S in equilibrium.*

Corollary 2 is essentially the smooth, continuous-action analog of what is illustrated in Figure 2. In words, the required smoothness conditions are that different types of sellers play different reserve prices, different reserve prices lead to different probabilities of trade, and all functions are differentiable. Under these conditions, the upper and lower bound slopes in Figure 2 collapse to a line, and the seller's value is exactly the slope of that line—the tangent line at the point the seller chooses.²¹ We

²¹Corollary 2 applies only to strictly separating equilibria in the sense that each seller type chooses a different reserve price and each reserve price leads to a different trade probability. In some bargaining games, equilibria feature partial-pooling regions where, for example, all sellers with

maintain these smoothness conditions throughout the rest of the paper. Note that these are conditions assumed directly on equilibrium objects (the equilibrium menu and the secret reserve price function) rather than on primitives.

Corollary 2 can be considered a generalization of the seminal first price auction result of Guerre et al. (2000) (GPV).²² The GPV argument, however, is specific to the first-price auction game. In contrast, our approach can be applied to any game in which the econometrician observes whether trade was successful, the transfers paid between traders, and some variable that is one-to-one with agents' types. This feature of our menu approach may be advantageous in other bargaining settings as well, where equilibria may be difficult to characterize, and probability and transfer rules may be different from standard well-studied auction formats.

The BNE framework also imposes two restrictions on equilibrium menus, which we utilize in identification and estimation. First, in addition to incentive compatibility, equilibrium menus must satisfy *individual rationality*, implying that the seller's expected payoff under the optimal choice of r must exceed v_S , the seller's value of keeping the car herself:

$$\max_{r'} v_S P_S(r') - T_S(r') \geq v_S \quad \forall v_S \quad (5)$$

values above some cutoff never trade. In such settings, values for agents who pool would not be point identified; however, the arguments in Theorem 1 would still yield a one-sided bound on these agents' values by considering deviations to the nearest interior action. We do not need to appeal to these arguments here because our data appear to be described well by a strictly separating equilibrium: we estimate a strictly increasing $P_S(r)$ function, and a convex menu with no mass points.

²²To see this, consider an n -bidder first-price auction in a symmetric independent private values environment. Let the distribution of bids be written $G(\cdot)$, with density $g(\cdot)$. In a first-price auction, the expected probability of winning, (P_i) , and expected transfer (T_i) for bidder i bidding bid b_i are known transformations of G , given by $P_i(b_i) = G(b_i)^{n-1}$ and $T_i(b_i) = b_i G(b_i)^{n-1}$. Applying our menu approach to differentiate P_i and T_i as in Corollary 2, player i 's value is given by

$$v_{B_i} = \frac{\frac{dT_i(b_i)}{db_i}}{\frac{dP_i(b_i)}{db_i}} = \frac{b_i(n-1)G(b_i)^{n-2}g(b_i) + G(b_i)^{n-1}}{(n-1)G(b_i)^{n-2}g(b_i)} = b_i + \frac{G(b_i)}{(n-1)g(b_i)}.$$

This last expression is equivalent to that derived in the identification argument of GPV.

Second, equilibrium menus must be convex, which we state as a theorem:

Theorem 2. *The graph of $\{(P_S(r), T_S(r))\}$ is convex.*

The intuition for Theorem 2 can be seen in the left panel of Figure 2. Every action played in equilibrium must be optimal for some type, so the upper and lower bounds in Theorem 1 must intersect at some point. Any point interior to the menu’s convex hull is dominated: no type would find it optimal to play such actions.²³ We impose convexity and IR constraints on our estimated menu. These constraints serve as tests of the model: violations suggest that BNE behavior does not rationalize the data well.

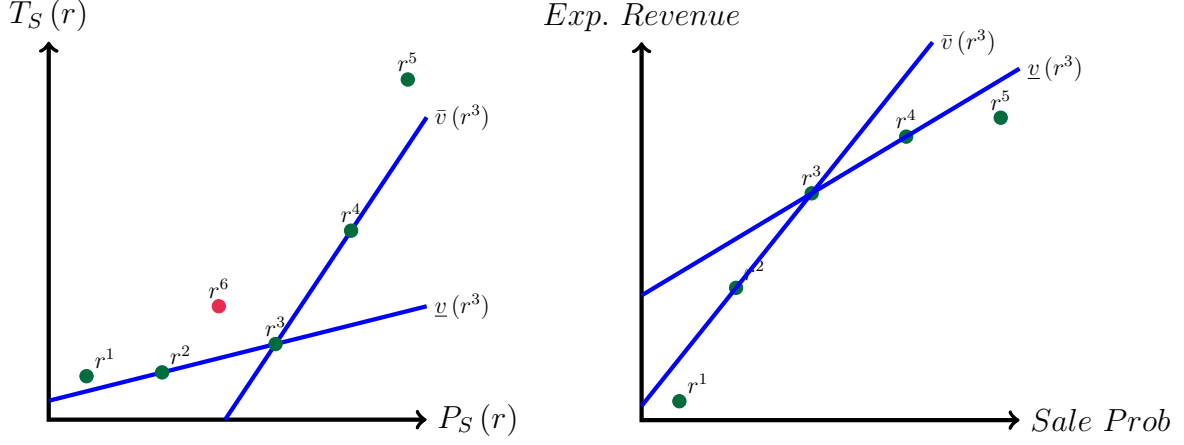
One key point about this identification argument is that it relies on variation in the probability that a game ends in agreement; the argument is not useful if the researcher observes no cases where parties disagree. Indeed, in bilateral bargaining, if the researcher only observes successful trades, the researcher cannot reject the possibility that complete information (e.g. Nash bargaining) is actually the correct behavioral model. It is data on failed attempts to trade that are essential for rejecting a complete-information environment and for identifying agents’ private values.

4.2 Identification of the Distribution of Buyer Values

Identification of F_B is relatively standard. For any continuation game after the auction, dropping out when the auction price reaches the bidder’s value is a weakly dominant strategy (Proposition 1.ii). Recall that N , the number of bidders, is a random variable varying across instances of the game. Let $\Pr(N = n)$ denote the probability mass function of N . F_B is identified via an order statistics inversion using $\Pr(N = n)$ and the distribution of auction prices $F_{PA}(\cdot)$. In Section 5 we discuss how

²³Note that the menu agents are *offered* does not necessarily correspond to the *equilibrium* menu observed by the econometrician, and the offered menu need not be convex. Throughout the paper, we use the term *menu* to refer to this *equilibrium* menu.

Figure 2: Hypothetical Equilibrium Menu



Notes: The left panel shows a hypothetical equilibrium menu. The slopes of the blue lines are upper and lower bounds for the value of an agent choosing action r^3 . The right panel shows the menu in terms of the seller's probability of sale, $1 - P_S(r)$, and expected revenue, $-T_S(r)$. Again, the slopes of the blue lines are upper and lower bounds for the value of a seller choosing action r^3 .

we estimate $\Pr(N = n)$ and $F_{PA}(\cdot)$. For any y , the following holds:

$$F_{PA}(y) = \sum_n \Pr(N = n) [nF_B(y)^{n-1} - (n-1)F_B(y)^n]. \quad (6)$$

The right-hand side of (6) is monotonic in y , and thus F_B is identified. The distribution of the highest-value bidder (the maximum order statistic) is then given by

$$F_B^{(1)}(y) = \sum_n \Pr(N = n) F_B(y)^n \quad (7)$$

This is the value distribution for the bidder who potentially bargains with the seller.

4.3 Identification Under Game-Level Heterogeneity

We now describe how we incorporate game-level heterogeneity. Recall that game-level heterogeneity is captured by $W + Y\beta$, where Y is observable to the econometrician and W is not. By Proposition 1.iv, the game is *location invariant*: P^A , R , and final prices shift additively with $W + Y\beta$, and the trade probability is invariant. For our

discussion, we condition on a realization of Y and focus here on W ; we discuss Y in Section 5.

BNE in this game requires that agents' strategies constitute a BNE conditional on any value of w . Define the expected probability and transfer the seller achieves when playing reserve price $r + w$ in equilibrium, when $W = w$, as $P_S^w(r + w) \equiv E[\mathbf{x}_S(r + w, V_{-S} + w) \mid W = w]$ and $T_S^w(r + w) \equiv E[t_S(r + w, V_{-S} + w) \mid W = w]$, with expectations taken over other agents' values (and hence their equilibrium actions). Due to location invariance, equilibrium menus are fully characterized by probabilities and transfers conditional on $w = 0$, $P_S^0(r), T_S^0(r)$. These objects are not immediately identified from conditional expectations in the data because we only observe realizations of *noisy* reserve prices $\tilde{R} \equiv R + W$, which are confounded with unobserved heterogeneity W . Rather, we can identify probabilities and transfers conditional on realizations of $\tilde{R} = \tilde{r}$, which we denote $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$, and then identify $P_S^0(\cdot)$ and $T_S^0(\cdot)$ from these objects. As an intermediate step, we also identify the densities f_W , f_R , and f_{PA} by a standard convolution argument exploited elsewhere in the auction literature. This requires several technical assumptions on characteristic functions, stated in Appendix A.3.

Our identification result is the following theorem:

Theorem 3. $P_S^0(\cdot), T_S^0(\cdot)$ are identified from $\tilde{P}_S(\tilde{r}), \tilde{T}_S(\tilde{r})$, and the joint distribution of noisy reserve prices and auction prices, $\tilde{R} = R + W$ and $\tilde{P}^A = P^A + W$.

The proof of Theorem 3 demonstrates that $P_S^0(\cdot)$, the underlying expected allocation function purged of unobserved heterogeneity, solves

$$\tilde{P}_S(\tilde{r}) = \frac{\int P_S^0(r) f_R(r) f_W(\tilde{r} - r) dr}{\int f_R(r) f_W(\tilde{r} - r) dr}, \quad (8)$$

and $T_S^0(r)$ solves

$$\tilde{T}_S(\tilde{r}) - E[W \Delta P_S \mid \tilde{r}] = \frac{\int T_S^0(r) f_R(r) f_W(\tilde{r} - r) dr}{\int f_R(r) f_W(\tilde{r} - r) dr}, \quad (9)$$

where

$$E[W\Delta P_S | \tilde{r}] \equiv \frac{\int (\tilde{r} - r) (P_S^0(r) - 1) f_R(r) f_W(\tilde{r} - r) dr}{\int f_R(r) f_W(\tilde{r} - r) dr}. \quad (10)$$

We describe in Section 5 how we exploit (8)–(10) to estimate $P_S^0(\cdot)$ and $T_S^0(r)$.

5 Estimation of Menus and Private Values

5.1 Estimation Details

Our estimation follows the above identification arguments closely. Let j denote an observation in the data, consisting of the allocation (who ends up with the car), the transfer (zero if car the does not sell and the final price if it does), the seller’s secret reserve price, and the auction price. We observe a large set of other characteristics for each game j , which we denote y_j . We describe each estimation step in turn below.

Observable Heterogeneity. We specify the total game-level heterogeneity for observation j to be $y_j'\beta + w_j$, where β is a vector of parameters to be estimated. Let r_j^{raw} and $p_j^{A,raw}$ be the reserve and auction prices for observation j before removing any game-level heterogeneity. We control for observable heterogeneity through a standard *homogenization* regression (Haile et al. 2003). We run the following regression:

$$\begin{bmatrix} r_j^{raw} \\ p_j^{A,raw} \end{bmatrix} = \begin{bmatrix} y_j'\beta \\ y_j'\beta \end{bmatrix} + \begin{bmatrix} \tilde{r}_j \\ \tilde{p}_j^A \end{bmatrix}, \quad (11)$$

where $\tilde{r}_j \equiv r_j + w_j$, $\tilde{p}_j^A \equiv p_j^A + w_j$. To control for as much variation as possible, y_j includes a rich set of observables: dummies for each make-model-year-trim-age combination (the age of the vehicle in years), dummies for the interaction of mileage with car-make dummies, dummies for 32 vehicle damage categories, and more.²⁴

²⁴Other controls are fifth-order polynomials in the auction house’s blue book estimate and the odometer reading; the number of previous attempts to sell the car; the number of pictures of the car on the auction house’s website; a dummy for whether the odometer reading is considered accurate, and the interaction of this dummy with the odometer reading; dummies for condition report grade (ranging from 1–5); dummies for the year-month combination of the sale date and for auction house

To improve our estimates in these regressions, we use our main sample of 131,443 observations, which contain both auction and reserve prices, as well as an additional 80,213 observations for which we observe only a reserve or auction price but not both.²⁵ The R^2 from this regression is 0.95, suggesting that most of the variation in these prices arises from observable differences across cars. For the sake of concreteness, we refer to the predicted value $y'_j\hat{\beta}$ as the *market value* of the car.

Unobserved Heterogeneity Distribution Estimation. The residuals from the above regression, \hat{r}_j and \hat{p}_j^A , are contaminated with unobserved heterogeneity, $W \sim F_W$. We estimate the marginal distributions F_W , F_R , and F_{PA} and their corresponding densities via maximum likelihood (MLE). The contribution of observation j to this likelihood is given by

$$\int f_{PA}(\hat{p}_j^A - w; \theta_{PA}) f_R(\hat{r}_j - w; \theta_R) f_W(w; \theta_W) dw \quad (12)$$

The objects θ_{PA} , θ_R , and θ_W are parameter vectors to be estimated. We specify the density of each random variable $Z \in \{P^A, R, W\}$ as $N(\mu_Z, \sigma_Z)$.²⁶ An alternative approach would be to approximate characteristic functions and perform a Fourier transform, as in [Li and Vuong \(1998\)](#) or [Krasnokutskaya \(2011\)](#). Either approach is consistent; we found the likelihood approach (used also in [Athey et al. 2011](#)) most straightforward.

Estimating Buyer Value Distribution. We estimate F_B by solving (6) on a grid

location interacted with hour of sale; dummies for each seller appearing in at least 500 observations; dummies for discrete odometer bins; and several measures of the thickness of the market during a given sale and of the order the cars were run (see [Larsen 2021](#) for details on their construction).

²⁵The additively separable structure we impose is testable. In particular, we could instead estimate β in (11) using only auction prices or only reserve prices. We find that doing so yields predicted values of $y'_j\hat{\beta}$ that are highly correlated. However, because our estimation involves a number of subsequent steps, such differences can matter for our final estimates of bargaining power. We discuss these differences in Appendix B.1. We choose to pool reserve and auction prices to use all available information.

²⁶We explored a more flexible approximation for the densities in this step using fifth-degree Hermite polynomials, as in [Gallant and Nychka \(1987\)](#). We found that a likelihood ratio test failed to reject the more parsimonious Normal approximation.

of values for the buyer. This requires an estimate of F_{PA} , which comes from the MLE estimates above, and an estimate of $\Pr(N = n)$, which we obtain from the subsample of the data with detailed bid logs. This data does not record the actual number of bidders, but an auction-by-auction upper bound on the number of bidders can be imputed based on the sum of two objects that are observable: the total number of *bids* placed by bidders who were physically present for a given auction and the number of bidders registered to participate *online* in the bidding of a given auction house lane.²⁷ After estimating F_B , we then construct the value distribution for the bidder who wins the auction (and potentially enters bargaining), $F_B^{(1)}$, using (7).

Local Linear Regressions for $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$. We estimate the noisy expected allocation function $\tilde{P}_S(\tilde{r})$ through a local linear regression of the allocation in game j on the noisy reserve price, i.e. the residual from the homogenization step, $\hat{\tilde{r}}_j$ for game j . We estimate the noisy expected transfer $\tilde{T}_S(\tilde{r})$ analogously through a local linear regression of the transfer in game j on $\hat{\tilde{r}}_j$. For these regressions, we use a Gaussian kernel and bandwidth of \$500. For comparison, the mean and standard deviation of reserve prices in the data are \$10,385 and \$5,805.

Estimating $P_S^0(r)$ via Spline-Fitting. We parameterize the function $P_S^0(r)$ as a quadratic I-spline (Ramsay 1988) with 5 knots, constrained to be nondecreasing in r .²⁸ Denote this $P_S^0(r; \theta_P)$. We estimate the spline coefficients θ_P as the solution to

$$\min_{\theta_P} \int \left[\left(\frac{\int P_S^0(r; \theta_P) \hat{f}_R(r) \hat{f}_W(\tilde{r} - r) dr}{\hat{h}(\tilde{r})} \right) - \hat{P}_S(\tilde{r}) \right]^2 \hat{h}(\tilde{r}) d\tilde{r} \quad (13)$$

In words, (13) numerically solves (8) by minimizing the distance between the esti-

²⁷In Appendix B.1, we discuss alternative choices for $\Pr(N = n)$.

²⁸We choose these knots to be uniformly spaced between -3000 and 5000. The choice of 5 as the number of knots (for P_S^0 and for \tilde{T}_S^0 , described below) was driven by an attempt to remain flexible while avoiding over-fitting, which required some degree of experimentation. Avoiding over-fitting at this stage of the estimation is important, as these objects are inputs in the subsequent stage where we differentiate to obtain estimates of F_S . For all integrals against the density of r or \tilde{r} , such as (13), we approximate the integral using a uniformly spaced grid of \tilde{r} values from -5000 to 7000 in intervals of 10. These end points correspond approximately to the 0.00001 and 0.99999 quantiles of the density of \tilde{r} .

mated function $\hat{P}_S(\tilde{r})$ and the convolution of $P_S^0(r; \theta_P)$ and $\frac{\hat{f}_R(r)\hat{f}_W(\tilde{r}-r)dr}{\hat{h}(\tilde{r})}$, weighting by $\hat{h}(\tilde{r}) \equiv \int \hat{f}_R(r) \hat{f}_W(\tilde{r}-r) dr$, which is the estimated density function of \tilde{r} .²⁹ We denote the estimate $P_S^0(r; \hat{\theta}_P)$ by $\hat{P}_S^0(r)$. The square root of the value of (13) at the optimum constitutes one measure of fit—a root weighted mean squared error (RMSE). Because we estimate a probability in this step, the RMSE naturally lies in $[0,1]$. We estimate a RMSE of 0.012, suggesting that our probability estimates differ from the local linear regression fit by only 1.2 percentage points on average.

Estimating $T_S^0(r)$ via Spline-Fitting. Using the estimated probability $P_S^0(r; \hat{\theta}_P)$ function, we parameterize the expected transfer function as a convex regression spline (C-spline; see Meyer 2008) in the probability rather than as a function of r directly. We denote this composition by $\check{T}_S^0(P_S^0(r; \hat{\theta}_P); \theta_T)$.³⁰ This type of spline approximation allows us to directly constrain the transfer-probability menu to be convex. We estimate the spline coefficients θ_T as the solution to

$$\min_{\theta_T} \int \left[\left(\frac{\int \check{T}_S^0(P_S^0(r; \theta_P); \theta_T) \hat{f}_R(r) \hat{f}_W(\tilde{r}-r) dr}{\hat{h}(\tilde{r})} \right) - \left(\tilde{T}_S(\tilde{r}) - \hat{E}[W\Delta P_S | \tilde{r}] \right) \right]^2 \hat{h}(\tilde{r}) d\tilde{r}.$$

subject to the constraint that $\check{T}_S^0(p)$ is weakly convex. This exercise also requires an estimate of $E[W\Delta P_S | \tilde{r}]$, which we construct using (10). We denote the value of $\check{T}_S^0(P_S^0(r; \hat{\theta}_P); \hat{\theta}_T)$ at the estimated parameters, when written as a function of r , by $\hat{T}_S^0(r)$. The RMSE from this step represents the amount in dollars by which our fit is off; we find this number to be quite low (\$3) relative to prices in this market, indicating a good fit.

Imposing convexity and individual rationality. As discussed in Section 4.1, the BNE model imposes two key restrictions on equilibrium menus: menus must be convex, and all menu points must satisfy individual rationality (IR), shown in (5). We

²⁹Weighting by $\hat{h}(\tilde{r})$ does not matter asymptotically. In finite samples it has the effect of down-weighting estimation where \tilde{r} has low density—where $\hat{P}_S(\tilde{r})$ may be less accurately estimated.

³⁰For the spline approximation of \check{T}_S^0 we again use 5 knots, placing more knots where more of the mass lies, with knots at $\{0, 0.1, 0.25, 0.4, 0.6\}$.

impose both constraints on the menu during estimation. We impose convexity on the menu through the constraints of the C-spline approximation $\check{T}_S^0(p)$. The convexity constraint is binding only on a set of measure 0: we find that one C-spline coefficient is equal to 0, implying that the second derivative of the menu is equal to 0 for exactly one value of P . To impose IR, for any agents choosing reserve prices r such that (5) does not hold, we hold $\hat{P}_S^0(r)$ fixed, and set $\hat{T}_S^0(r)$ to the value that makes (5) hold with equality. This treats these agents as trading with probability $\hat{P}_S^0(r)$ but having zero expected surplus, making them exactly indifferent between participating or not. We find that, even without enforcing them, IR constraints are satisfied for an overwhelming majority of sellers (90.4%). For the remaining 9.6%, the IR constraint binds. These facts—convexity of the estimated menu and IR constraints being non-binding for most of the estimated menu—offer some evidence of good model fit.


Estimating the Seller Value CDF and Reserve Price Function. From the final estimated menu, we then construct, for a grid of values for r , an estimate of the corresponding type v_S that would choose each r . Specifically, for each r , we obtain the inverse mapping $v_S(r)$ by evaluating the derivative $\frac{d\hat{T}}{dp}$, which has a closed form given our spline approximation, at $\hat{P}_S^0(r)$. This also yields an estimate of the mapping $\hat{\rho}(v_S)$. An estimate of F_S is then given by $\hat{F}_S(v_S) = \hat{F}_R(\hat{\rho}(v_S))$ for any v_S . Note that this exploits the estimate of \hat{F}_R from the MLE unobserved heterogeneity step. In what follows, we sometimes exploit the seller’s expected allocation and transfer as a function of her value rather than her reserve price. Estimates of these objects are given by $\hat{P}_S^0(\hat{\rho}(v_S))$ and $\hat{T}_S^0(\hat{\rho}(v_S))$.

5.2 The Estimated Menu and Value Distributions

We now display, in Figure 3.A, the estimated menu, with the estimate of $P_S^0(r)$ on the horizontal axis and $T_S^0(r)$ on the vertical axis.³¹ Each point on the menu corresponds to the expected payoff for the seller from choosing a given secret reserve price. We note

³¹Appendix B.1 provides estimates of the initial estimation steps described in Section 5.1.

Figure 3: Estimated Menu

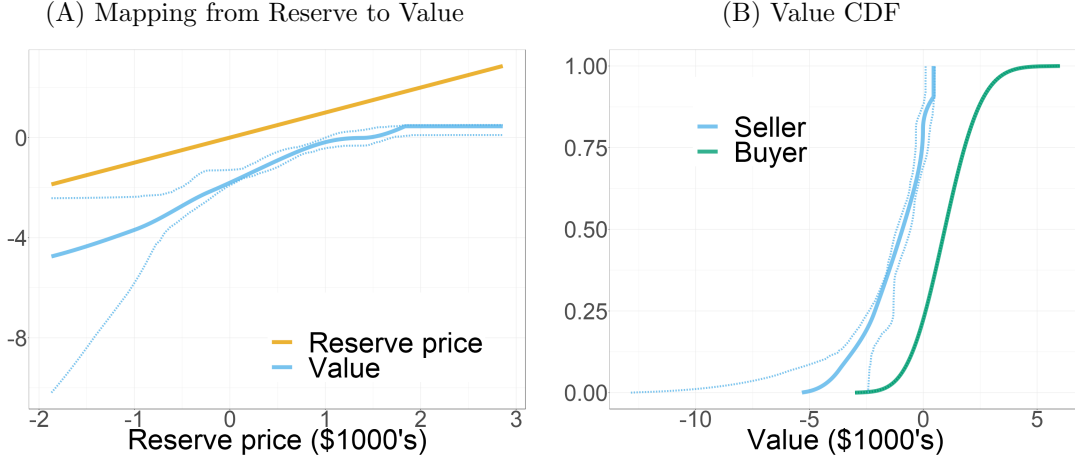


Notes: Panel A displays the final estimated menu, $(\hat{P}_S^0, \hat{T}_S^0)$. Dashed lines show pointwise 95% confidence bands from bootstrapping with 200 replications. Units on the vertical axis are \$1,000 relative to the market value estimate. The points marked A–D are discussed in the body of the paper. Right panel shows sale probabilities, $1 - P_S^0(r)$, and expected revenues, adjusted upward by \$10,000, approximately the mean estimated market value.

here that our formulation for payoffs described in (1) in Section 4, $v_S P_S^0(r) - T_S^0(r)$, means $P_S^0(r)$ represents the probability of the seller *keeping* the good (so $1 - P_S^0(r)$ is the probability that a sale occurs) and $T_S^0(r)$ is the expected transfer *paid by the seller* (so $-T_S^0(r)$ is the expected payment received by the seller). The units for the vertical axis are \$1,000, and these numbers can be negative in panel A because they are the result of subtracting off the market value estimate $y_j' \hat{\beta}$; these numbers thus indicate where a monetary amount lies relative to the car's market value.

With this interpretation in mind, we compare several points along the menu, r_A, \dots, r_D , where $r_A < r_B < r_C < r_D$, in Figure 3.A. Points A and B lie along the downward-sloping portion of the menu. Choice r_A yields a lower probability of keeping the good and a lower expected transfer (i.e. a less negative T_S^0) than would choice r_B . Therefore, a seller who chooses r_A wants to get rid of the car more than a seller who chooses r_B , implying that the former seller has a lower value (lower v_s) than the latter. This is precisely what Figure 3.A shows: the derivative of the menu at r_A is more negative than at r_B , and these derivatives (by Corollary 2) reveal sellers' values, so a seller choosing r_A must therefore have a value that is farther below the market value of the car than does a seller choosing r_B . Points C and D lie along the upward-

Figure 4: Value Mapping and Value CDFs



Notes: Left panel shows the estimated mapping from reserve prices to values (blue line) and the reserve price (yellow line, $y = x$). Right panel displays estimated CDF of seller values, F_S (blue line), and estimated CDF of maximum order statistic of buyer values, $F_{B(1)}$ (green line). Dashed lines show 95% confidence bands from bootstrapping with 200 replications. Units on horizontal axes (and on vertical axis of panel A) are \$1,000 relative to market value estimate.

sloping portion of the menu. Choice r_C yields a lower probability of keeping the good but a higher expected transfer to the seller (i.e. a more negative T_S^0). Therefore, a seller choosing r_D wants to keep the good more (i.e. have a higher v_S) than a seller choosing r_C , manifest by a derivative that is more positive at D than at C .

In Figure 3.B, we offer an alternative version of this menu in terms of expected revenues and sale probabilities, analogous to the right panel of Figure 2. To improve readability, we adjust expected revenues by adding \$10,000, which is approximately the average market value ($y'_j \hat{\beta}$) for cars in our data.³² The plot displays several possible reserve price choices (in units of \$1,000). As in Figure 2, by choosing lower reserve prices, the seller sells with higher probability and attains higher expected revenue, but the marginal sale revenue decreases as r decreases.

Taking the menu's derivatives at each point r yields the mapping between the reserve price and the inferred value v_S . We display this mapping with a solid blue line in Figure 4.A, with reserve prices on the horizontal axis and values on the vertical

³²This adjustment does not merely shift the menu in Figure 3.B upward but also rotates it. The explanation of the rotation introduced by shifts in game-level heterogeneity is found in the proof of Theorem 3.

axis. The units for each axis are \$1,000. The dashed lines indicate a pointwise bootstrapped 95% confidence band. The yellow line shows the 45 degree line (the reserve price itself). To interpret, consider a particular point on this mapping at about $r = -\$300$ (-0.3 on the horizontal axis). We see that the corresponding inferred value for such a seller is about $v_S = -\$2,000$ (a value of -2 on the vertical axis). Therefore, a seller who chooses a reserve price that is \$300 below the market value is actually willing to let the car go for up to \$2,000 below the market value.

The estimated reserve-value mapping, combined with the distribution \hat{F}_R , gives us an estimated distribution \hat{F}_S of sellers' values, and we plot this in Figure 4.B. We also plot the estimated distribution of the highest-value buyer—the buyer who potentially ends up in bargaining. The distributions indicate that 75% of sellers have values less than the market value, while 75% of buyers have values *above* this amount. This suggests that there are typically gains from trade in this market, which is to be expected given that auction houses have been functioning well as market makers in this industry for three quarters of a century. However, panel B also shows overlap between seller and buyer distributions, which, by the [Myerson and Satterthwaite \(1983\)](#) Theorem, can lead to inefficiency, with some trades failing to occur even when the highest-value buyer values the car more than the seller.

6 Quantifying Bargaining Power

In this section, we use the estimated menu and value distributions to compute bargaining power weights (α_S and α_B). This requires estimating buyer and seller surplus in the real-world mechanism, as well as surplus under the buyer-optimal and seller-optimal mechanisms. We then show how bargaining power differs between buyers and sellers, how these weights differ for cases in which the seller is a manufacturer or not, how competition among buyers affects bargaining power, and how bargaining power relates to the size of the pie.

6.1 Computing Surplus and Bargaining Power Weights

We first discuss how we calculate the ex-ante trade surplus, or *gains from trade*, for the real-world mechanism (the bargaining observed in the data). To calculate total surplus, we exploit the property that a seller of type v_S trades with all buyer types above a certain cutoff (Proposition 1.v).³³ Here we invoke the additional simplifying assumption that these cutoffs do not vary with the auction price, p^A .³⁴ Consequently, the $P_S^0(\boldsymbol{\rho}(v_S))$ function and the value distribution for the highest-value bidder, $F_{B^{(1)}}(\cdot)$, together imply cutoffs $\tilde{c}(v_S) = \left\{v_B : 1 - F_B^{(1)}(v_B) = P_S^0(\boldsymbol{\rho}(v_S))\right\}$. This allows us to calculate the total surplus from the actual mechanism as

$$\int \int (v_B - v_S) 1\{v_B \geq \hat{c}(v_S)\} d\hat{F}_{B^{(1)}}(v_B) d\hat{F}_S(v_S). \quad (14)$$

We can also calculate the ex-ante surplus for the seller, which is the expected payment to the seller minus the expected value of the seller conditional on trading. Because our estimation procedure produces estimates of the probabilities $P_S^0(\boldsymbol{\rho}(v_S))$ and transfer $T_S^0(\boldsymbol{\rho}(v_S))$ for all seller values, we can calculate the surplus achieved by any seller type as $v_S \hat{P}_S^0(\hat{\boldsymbol{\rho}}(v_S); \hat{\theta}_P) - \hat{T}_S^0(\hat{\boldsymbol{\rho}}(v_S))$. The ex-ante surplus for the seller is then given by integrating this quantity over v_S . Buyer gains from trade are obtained by subtracting the seller surplus from the total surplus.

We then compute the buyer-optimal mechanism, which is a TIOLIO price $p^*(v_B)$ maximizing $(v_B - p) \hat{F}_S(p)$. Total surplus is given by

$$\int \left[v_B \hat{F}_S(\hat{p}^*(v_B)) + \int v_S 1\{v_S \geq \hat{p}^*(v_B)\} d\hat{F}_S(v_S) \right] d\hat{F}_{B^{(1)}}(v_B) \quad (15)$$

Expected seller surplus is then total surplus minus buyer surplus.

³³This property is also satisfied in many other bilateral trade settings, and is referred to by [Ausubel and Deneckere \(1993\)](#) as the “Northwestern Criterion.”

³⁴We have performed estimation allowing cutoffs to vary with p^A and obtained results that are similar qualitatively and quantitatively. We adopt this simplification to reduce the computational burden.

To calculate the seller-optimal mechanism, we use the well-known result from auction theory that, when buyers' values are symmetric, seller revenue is maximized by a second-price auction with a public reserve price. As shown in [Aradillas-López et al. \(2013\)](#), for a seller with value v_S , the seller's maximized surplus can be written

$$\max_p \int \max\{p, v\} dF_{PA}(v) - v_S - F_{B(1)}(p)(p - v_S) \quad (16)$$

We plug in our estimated distributions and solve (16) for each seller type to find optimal public reserve prices $p^*(v_S)$ and seller surplus in the seller-optimal mechanism. Analogous to (15), total surplus is

$$\int \left[v_S \hat{F}_{B(1)}(\hat{p}^*(v_S)) + \int v_B 1\{v_B \geq \hat{p}^*(v_S)\} d\hat{F}_{B(1)}(v_B) \right] d\hat{F}_S(v_S),$$

and expected buyer surplus is calculated as total surplus minus seller surplus.

With estimates of surplus for the buyer and seller under each of these mechanisms, we compute bargaining power weights. In the notation of Section 3, the seller-optimal mechanism is \mathcal{M}^1 and the buyer-optimal mechanism is \mathcal{M}^0 . Let the real-world mechanism be denoted \mathcal{M}^{RW} . Bargaining power is then given by

$$\hat{\alpha}_B = \frac{\hat{U}_B(\mathcal{M}^{RW}) - \hat{U}_B(\mathcal{M}^1)}{\hat{U}_B(\mathcal{M}^0) - \hat{U}_B(\mathcal{M}^1)} \quad (17)$$

$$\hat{\alpha}_S = \frac{\hat{U}_S(\mathcal{M}^{RW}) - \hat{U}_S(\mathcal{M}^0)}{\hat{U}_S(\mathcal{M}^1) - \hat{U}_S(\mathcal{M}^0)} \quad (18)$$

where $\hat{U}_B(\mathcal{M})$ is the estimate of the ex-ante expected surplus of the buyer under mechanism \mathcal{M} , as in Section 3, and similarly for $\hat{U}_S(\mathcal{M})$.

6.2 Overall Estimates of Bargaining Power Between Buyers and Sellers

In this subsection, we focus on our estimates of bargaining power for our full sample. We discuss various subsamples in the subsections that follow. Table 2 shows our

Table 2: Components for Measuring Bargaining Power

| | (1) | (2) | (3) | (4) | (5) | (6) |
|----------------|------------------|-----------------------|------------------------|-----------------------|-----------------------|----------------------|
| | Trade Prob. | Total Surplus (\$) | Seller Surplus (\$) | Seller Margin (\$) | Buyer Surplus (\$) | Buyer Margin (\$) |
| Real-world | 0.763 (0.001) | 2,208 (143) | 1,251 (136) | 1,640 (176) | 957 (15) | 1,254 (19) |
| Buyer-optimal | 0.549 (.028) | 1,933 (141) | 799 (119) | 1,455 (224) | 1,135 (94) | 2,067 (141) |
| Seller-optimal | 0.712 (.011) | 2,178 (156) | 1,531 (145) | 2,150 (192) | 647 (22) | 908 (18) |

Notes: Sale probability, total surplus, seller surplus (and margin), and buyer surplus (and margin) in the buyer-optimal, seller-optimal, and real-world mechanisms. Seller and buyer margins are calculated by dividing seller and buyer surplus, respectively, by trade probability. Standard errors from 200 bootstrap samples are in parentheses

estimates of trade probabilities and surplus in the real-world, seller-optimal, and buyer-optimal mechanisms. In levels, we find that sellers achieve \$1,251 of surplus in the real-world mechanism, compared to \$1,531 in the seller-optimal mechanism, and \$799 in the buyer-optimal mechanism. We also calculate average margins for sellers by dividing seller utility by the trade probability; this is the average markup (in dollars) of the sale price over the seller's private value for successful trades. The average margin is \$1,640 in the real-world mechanism, compared to \$1,455 in the buyer-optimal mechanism and \$2,150 in the seller-optimal mechanism. Given that average car prices are approximately \$10,000, this implies that sellers' expected markups over their values V_S , in percentage terms, are approximately 16.4% of the price of a car, compared to 14.6% and 21.5% at the buyer- and seller-optimal mechanisms, respectively.

Buyer surplus, in levels, is \$957 in the real-world mechanism, compared to \$647 and \$1,135 in the seller-optimal and buyer-optimal mechanisms, respectively. Buyers' margins are \$1,254 in the real-world mechanism (roughly 12.5% of the price of a car), compared to \$2,067 and \$908 in the buyer- and seller-optimal mechanisms.

The six numbers in columns 3 and 5 of Table 2 are the building blocks to plug into (17) and (18) to construct estimates of bargaining power, α_B and α_S . The yellow and

blue objects in the first column of Figure 5 (“Full Sample”) display these estimates, surrounded by 95% confidence intervals. Our point estimates are $\hat{\alpha}_S = 0.618$ and $\hat{\alpha}_B = 0.636$. In words, the seller, on average, achieves an outcome that is 61.8% of the way between her maximal payoff—her payoff from making a TIOLIO—and the payoff she would receive if instead the buyer were to make a TIOLIO. The buyer, on average, achieves an outcome that is 63.6% of the way between his maximal payoff and what he would receive if the seller were to make a TIOLIO. We do not find evidence in the full sample of one side of the market having more bargaining power than another: the confidence intervals of buyer and seller bargaining power overlap.

The sum of $\hat{\alpha}_S$ and $\hat{\alpha}_B$ offers two measures of efficiency. First, we can compare this sum to 1 to evaluate how well the real-world bargaining performs relative to a simple random combination of the buyer- and seller-optimal mechanisms. For our full sample, we find that this sum is significantly greater than 1 (shown in green in the first column of Figure 5). This suggests that the real world achieves higher surplus than could be achieved by giving either party all of the bargaining power.

Second, we can compare $\hat{\alpha}_B + \hat{\alpha}_S$ in the real-world mechanism to what it would be in the counterfactual first-best (ex-post) efficient outcome, which consists of the buyer and seller trading whenever the buyer’s value exceeds the seller’s. We show this estimate in purple in the first column of Figure 5.³⁵ This first-best mechanism can also be thought of as the outcome that would occur under *vertical integration* between the upstream supplier (the company selling the car) and the downstream buyer (the dealership), assuming there are no remaining bargaining frictions within such an integrated firm. Without vertical integration, the first-best efficient surplus is generally unattainable in bilateral bargaining (Myerson and Satterthwaite 1983). In our full sample, however, we find that the real-world mechanism lies close to the

³⁵Note that the division of surplus is not unique within the first-best mechanism, so the first-best sum of α_S and α_B is not uniquely pinned down. For plotting the first-best in Figures 5, 6, and A.5, we calculate the sum $\alpha_B + \alpha_S$ corresponding to the same ratio of bargaining power α_B/α_S as in the real-world mechanism. Graphically, this is equivalent to the division of surplus in the first-best that lies on a ray connecting the point $\alpha_B = 0, \alpha_S = 0$ to the real-world α_B, α_S in Figure 1.

vertically integrated outcome, suggesting that efficiency losses due to incomplete-information are not large in the full sample.³⁶


6.3 Manufacturer vs. Non-manufacturer Sellers

We next analyze whether bargaining power differs between sellers who are car *manufacturers* and those who are not, and whether these manufacturer sellers have more bargaining power than the dealers with whom they negotiate. Manufacturers, such as Ford, General Motors, and Daimler Chrysler (the three largest U.S. companies), are often major clients of used-car auctions, typically selling cars that have been leased for a period and then re-purchased by the manufacturer, at which point the manufacturer takes care of vehicle resale. As described in Section 2, other major sellers in the secondhand wholesale market include fleet companies (such as Wheels or Orix), banks (such as Bank of America or Wells Fargo), or rental companies (such as Budget or Enterprise). Manufacturers have a more complicated relationship with franchised car dealers (who can be buyers at auctions) than do non-manufacturers, with many state-specific laws governing relationships.

In the second and third columns of Figure 5, we repeat our analysis of $\hat{\alpha}_S$ and $\hat{\alpha}_B$ separately for observations where the seller is a manufacturer vs. not. We find that, in contrast to our full-sample results, in manufacturer sales sellers have much more bargaining power than buyers: $\hat{\alpha}_S$ is 0.91 and $\hat{\alpha}_B$ is 0.045. Indeed, manufacturers' achieve a surplus that is very close to what they would receive if they had *all* of the bargaining power. Given that the outcome is so close to the seller-optimal mechanism,

³⁶This result contrasts with that of [Larsen \(2021\)](#), who documented some inefficiency in bargaining in wholesale used-car markets. The difference is driven by several factors. First, we analyze the full real-world mechanism (the auction plus bargaining), whereas the former only studies the post-auction bargaining. Second, we study only fleet/lease sales, and the bounds on efficiency loss in Larsen's earlier work are indeed closer to zero for such sales. Third, our point estimates of \hat{F}_S lie close to—but slightly outside of (stochastically dominating)—the Larsen's bounds, which can happen because the two approaches rely on different assumptions. Our approach exploits optimality of reserve prices, whereas the former exploits optimality of the seller's choice to accept or reject the auction price. A stochastically higher seller value CDF corresponds to a smaller gap between the real-world mechanism and the first-best (see Proposition 6 of [Larsen 2021](#)).

Figure 5: Bargaining Power Estimates



Notes: Estimates of $\hat{\alpha}_B$ (yellow), $\hat{\alpha}_S$ (blue), the sum $\hat{\alpha}_B + \hat{\alpha}_S$ (green), and the sum $\hat{\alpha}_B + \hat{\alpha}_S$ in the first-best mechanism (purple), for the full sample (leftmost), manufacturer and non-manufacturer sales, and two comparison samples: non-manufacturer sales with similar average mileage and average age to manufacturer sales, and non-manufacturer sales with a similar number of bidders as the manufacturer sample. To construct the first-best sum $\alpha_B + \alpha_S$, within each subsample, we take first-best total surplus, and divide it between buyers and sellers such that the ratio α_S/α_B is the same as in the real-world mechanism. Points represent estimates in the baseline sample, and confidence bars represent 95% confidence intervals from 200 nonparametric bootstrap replications.

the real-world bargaining in this sample performs close to what could be achieved by a random ultimatum game (the 95% confidence interval includes 1), and thus is relatively efficient from this perspective. We can also measure efficiency relative to the first-best, vertically integrated counterfactual. Note that the vertical integration between dealers and manufacturer sellers also captures a notion of *direct-to-consumer* sales for used-car markets, as it removes the wholesale-market negotiation. Relative to this benchmark, manufacturer sales also do well, with the 95% confidence interval for the first best sum of bargaining power weights overlapping that of the real-world.

In the non-manufacturer sample we find starkly contrasting results: $\hat{\alpha}_S$ is 0.44 and $\hat{\alpha}_B$ is 0.81. Thus, buyers have much more bargaining power when facing non-manufacturer sellers. In this sample, we can reject the possibility that a random-proposer game would perform better—the sum of the α s is 1.25—but we also find evidence that the bargaining in this sample falls short of the first-best outcome, where $\alpha_B + \alpha_S$ is 1.38 (and the confidence intervals between the real-world and first-best sum do not overlap).

These differences in bargaining power between manufacturer and non-manufacturer sales may be driven by differences in the composition of the two samples. To investigate this possibility, we create a sample of non-manufacturer sales which are of similar age and mileage to manufacturer cars by limiting the sample of non-manufacturer sales to those with age and mileage values below particular cutoffs such that this sample has similar average characteristics to those of manufacturer cars. We refer to this as our *comparison* sample. We show the results using this characteristics-comparison sample in the fourth column of Figure 5.³⁷ In the comparison sample, $\hat{\alpha}_S$ is approximately 0.30 and $\hat{\alpha}_B$ is approximately 0.68. These are quite different—and oppositely ranked—from the estimates of the manufacturer sample, suggesting that these characteristics do not explain the difference between manufacturer and non-manufacturer sales.³⁸

Manufacturer sales also tend to have more bidders than non-manufacturer sales.³⁹ We construct a comparison sample based on the number of bidders in the auction, constructed using a cutoff such that the average number of bidders is the same in the manufacturer and comparison non-manufacturer samples. The results using this bidders-comparison sample are shown in the fifth column of Figure 5. Here we find that $\hat{\alpha}_S$ is much higher, approximately 0.83, and $\hat{\alpha}_B$ is approximately 0.31. Thus, in this bidders-comparison sample, sellers have slightly lower bargaining power than in the manufacturer sample, and buyers have much higher bargaining power. Hence, increased competition appears to partially explain sellers' high bargaining power in manufacturer sales, but does not fully explain why buyers' surplus is so low in the manufacturer sample.

³⁷Note that, by construction, the comparison samples have no overlap with the manufacturer sample, because the comparison sample consists entirely of non-manufacturer sales.

³⁸In Appendix B.3, we explore other sample splits, showing how bargaining power varies with characteristics such as car age, condition, and mileage, as well as seller experience.

³⁹The number of bidders in manufacturer sales tends to be high even though these sales typically have a restriction in place that only franchised dealers are allowed to attend and bid at the sale. The large number of bidders suggests that these dealers compete heavily for the chance to have late-model, used inventory of a given make.


Our finding that, on average, dealers have higher bargaining power than wholesalers in the secondhand market, and that the opposite is true among manufacturer sales, lends credence to previous claims that state laws governing the dealer-manufacturer relationship generally favor dealers (Lafontaine and Scott Morton 2010). Specifically, these laws govern the relationship on the supply side for *new* car transactions. For used cars, where manufacturers are not subject to these laws, we find that they hold the lion’s share of bargaining power. Lafontaine and Scott Morton (2010) also argue that dealer-manufacturer relationships in new markets are inefficient, with these state laws being the driver of the inefficiency and with consumers being the primary losers. Consistent with their arguments, our results suggest that manufacturer sales for used cars, where these laws do not apply, are not inefficient.

6.4 How Does Competition Affect Bargaining Power and Efficiency?

We now explore in more detail the effects of competition on bargaining power. To do so, we divide the full sample into thirds based on the upper bound on the number of bidders in the auction from the bid log data described in Section 5. We estimate bargaining power separately in each of these three subsamples. The results are shown in Figure 6. We find a clear monotonic relationship between the number of bidders and bargaining power: $\hat{\alpha}_S$ is higher and $\hat{\alpha}_B$ lower when there are more bidders.

Moreover, when the number of bidders is low, bargaining is very inefficient: the sum $\hat{\alpha}_S + \hat{\alpha}_B$ is far below 1, and particularly far below the vertically integrated, first-best outcome. Seller bargaining power $\hat{\alpha}_S$ is statistically indistinguishable from 0, so sellers achieve surplus similar to what they would in the buyer-offer mechanism, but $\hat{\alpha}_B$ is around 0.71, meaning that buyers are getting less surplus than they would in the buyer-offer mechanism. In the medium- and high-bidder samples sellers do better and buyers do somewhat worse, but the sum $\hat{\alpha}_S + \hat{\alpha}_B$ is significantly greater than 1 and its confidence interval overlaps with the vertically integrated outcome in both cases.

Figure 6: Bargaining Power Estimates, Splitting by Number of Bidders



Notes: Estimates of $\hat{\alpha}_B$ (yellow), $\hat{\alpha}_S$ (blue), the sum $\hat{\alpha}_B + \hat{\alpha}_S$ (green), and the sum $\alpha_B + \alpha_S$ in the first-best mechanism (purple), for the main dataset (leftmost), and three tercile samples of the data, based on the number of bidders in the auction. To construct the first-best sum $\alpha_B + \alpha_S$, within each subsample, we take first-best total surplus, and divide it between buyers and sellers such that the ratio α_S/α_B is the same as in the real-world mechanism. Points represent estimates in the baseline sample, and the confidence bars represent 95% confidence intervals from 200 nonparametric bootstrap replications.

Our finding in the low- N case is worth discussion. Recall from Section 3 that any outcome with $\alpha_S + \alpha_B = 1$ can be implemented by a *random ultimatum* game in which, with probability α_S , the seller gets to make a TIOLIO, and with probability $\alpha_B = 1 - \alpha_S$ the buyer gets to make a TIOLIO. This mechanism is a useful benchmark in that it is simple to implement from a practical perspective. In the full sample, as well as in the subsamples with medium and high numbers of bidders, we find $\hat{\alpha}_S + \hat{\alpha}_B$ to be significantly greater than 1, suggesting that the real-world bargaining does significantly better than this random ultimatum benchmark. However, when competition is low (the second column of Figure 6), $\hat{\alpha}_S + \hat{\alpha}_B$ is approximately 0.74, so surplus would be higher if the mechanism designer were to implement this randomized take-it-or-leave-it mechanism instead of the mechanism in the data. Vertical integration would also significantly improve efficiency in this case where there is low competition on the buyers' side of the market.

7 Conclusion

This study provides an empirical analysis of bargaining power under asymmetric information. We focus on negotiations between buyers (car dealers) and sellers (large institutions, such as manufactures or fleet-owning companies) in the supply side of the U.S. used-car market. These negotiations are facilitated by wholesale used-car auction platforms, who first run an ascending auction, and then facilitate bargaining between the seller and highest bidder whenever the auction price falls below the seller’s secret reserve price.

The private value distribution of the buyer can be easily estimated from data on auction prices. The private value distribution of the seller is much more complex to identify and estimate, but we show how this can be achieved by applying a revelation-principle-like argument, interpreting the seller’s choice of secret reserve price as a choice from a menu of expected probabilities of keeping the car and expected transfers—a menu that constitutes the marginal direct mechanism the seller faces. The derivative of this menu evaluated at the point chosen by the seller corresponds to the seller’s privately known value. These two key objects—the value distribution of the seller and buyer—can then be used to evaluate the seller’s and buyer’s bargaining power.

As bargaining power is not a well-studied concept in incomplete-information bargaining, we propose a new definition: an agent’s bargaining power is the share of the agent’s best-case (i.e. TIOLIO) surplus the agent achieves, relative to what the agent would achieve under the opponent’s best-case scenario. This extends a traditional (complete-information, Nash bargaining) notion of bargaining power to the incomplete-information setting.

We estimate the buyer and seller value distributions and their corresponding bargaining power in this market. We find that, overall, car dealers (buyers) have a similar degree of bargaining power to sellers (large companies). However, focusing on sales

by manufacturers vs. non-manufacturers, we find that manufacturers have substantially more bargaining power than other sellers, and substantially more power than the dealers with whom they negotiate. This strong bargaining position is driven in part by a high level of competition among buyers for manufacturer sales. In settings with fewer buyers competing in the auction, buyers' bargaining power is substantially higher than sellers and the bargaining is quite inefficient, both relative to a simple-to-implement random-proposer alternative and relative to a benchmark of vertical integration (the first-best outcome). We see these results as a first step toward understanding bargaining power in the supply side of used-car markets, and, more broadly, in quantifying bargaining power in industries where asymmetric information plays a role.

References

- Aradillas-López, A., Gandhi, A., and Quint, D. (2013). Identification and inference in ascending auctions with correlated private values. *Econometrica*, 81(2):489–534.
- Athey, S., Levin, J., and Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics*, 126(1):207–257.
- 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.
- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.
- Cramton, P. (1992). Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59(1):205–225.
- Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multi-channel television markets. *American Economic Review*, 102(2):643–685.
- Donna, J. D., Pereira, P., Trindade, A., and Yoshida, R. C. (2021). Direct-to-consumer sales by manufacturers and bargaining. Working Paper, Florida State.
- Evdokimov, K. and White, H. (2012). Some extensions of a lemma of Kotlarski. *Econometric Theory*, 28(4):925–932.

- Freyberger, J. and Larsen, B. (2021). How well does bargaining work in consumer markets? A robust structural analysis. NBER WP 29202.
- Gallant, A. and Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica*, 55(2):363–390.
- Goke, S. (2021). Non-binding secret reserve prices: The case of wholesale used-car auctions. Working Paper, Stanford University.
- Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.
- Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–177.
- Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first price auctions. *Econometrica*, 68(3):525–574.
- Gul, F. and Sonnenschein, H. (1988). On delay in bargaining with one-sided uncertainty. *Econometrica*, 56(3):601–611.
- Haile, P., Hong, H., and Shum, M. (2003). Nonparametric tests for common values at first-price sealed-bid auctions. NBER WP 10105.
- Ho, K. and Lee, R. S. (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, 109(2):473–522.
- Holmström, B. and Myerson, R. (1983). Efficient and durable decision rules with incomplete information. *Econometrica*, pages 1799–1819.
- Keniston, D. (2011). Bargaining and welfare: A dynamic structural analysis. Working Paper, Yale.
- Kline, B. (2017). Identification of incomplete information games in monotone equilibrium. Working Paper, UT Austin.
- Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies*, 78(1):293–327.
- Lacetera, N., Larsen, B. J., Pope, D. G., and Sydnor, J. R. (2016). Bid takers or market makers? The effect of auctioneers on auction outcomes. *American Economic Journal: Microeconomics*, 8(4):195–229.
- Lafontaine, F. and Scott Morton, F. (2010). Markets: State franchise laws, dealer terminations, and the auto crisis. *Journal of Economic Perspectives*, 24(3):233–50.

- Larsen, B. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies*, 88(2):851–882.
- Larsen, B., Lu, C., and Zhang, A. L. (2021). Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations. NBER WP 29159.
- Li, T. and Vuong, Q. (1998). Nonparametric estimation of the measurement error model using multiple indicators. *Journal of Multivariate Analysis*, 65(2):139–165.
- Loertscher, S. and Marx, L. M. (2019). Merger review for markets with buyer power. *Journal of Political Economy*, 127(9).
- Loertscher, S. and Marx, L. M. (2021). Incomplete information bargaining with applications to mergers, investment, and vertical integration. Working Paper, Duke.
- Luo, Y., Perrigne, I., and Vuong, Q. (2018). Structural analysis of nonlinear pricing. *Journal of Political Economy*, 126(6):2523–2568.
- Meyer, M. C. (2008). Inference using shape-restricted regression splines. *Annals of Applied Statistics*, 2(3):1013–1033.
- Murry, C. and Zhou, Y. (2020). Consumer search and automobile dealer colocation. *Management Science*, 66(5):1909–1934.
- 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.
- Perrigne, I. and Vuong, Q. (2011). Nonparametric identification of a contract model with adverse selection and moral hazard. *Econometrica*, 79(5):1499–1539.
- Perry, M. (1986). An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica*, pages 313–321.
- Pinkse, J. and Schurter, K. (2019). Estimation of auction models with shape restrictions. *arXiv preprint arXiv:1912.07466*.
- Ramsay, J. O. (1988). Monotone regression splines in action. *Statistical Science*, pages 425–441.
- Williams, S. R. (1987). Efficient performance in two agent bargaining. *Journal of Economic Theory*, 41(1):154–172.

Appendix for Online Publication

A Proofs

A.1 Proof of Corollary 2

Proof. We have assumed that $\rho(v_S)$ is strictly increasing—different types of sellers play different reserve prices—and that $P_S(r)$ is strictly monotone, meaning that different reserve prices lead to different probabilities of trade. Without loss of generality, suppose $P_S(r)$ is strictly increasing; the argument when $P_S(r)$ is decreasing is analogous.

Note that the support of R can be treated as bounded given that the support of V_B is assumed to be bounded. To see this, let $[\underline{M}, \overline{M}]$ denote the support of V_B . Choosing any secret reserve price r below \underline{M} is a dominated action for the seller given that every buyer has a value of at least \underline{M} . Moreover, the seller is indifferent between a secret reserve price of \overline{M} and any secret reserve price higher than this because no buyer would ever be willing to pay more than \overline{M} . Thus, the support of R can be treated as being bounded within $[\underline{M}, \overline{M}]$.

Consider some value v_S strictly in the interior of the interval $[\underline{M}, \overline{M}]$. First, we apply the bounds in (4), comparing $\rho(v_S)$, the reserve price chosen by v_S , to $\rho(v_S + \delta)$, the reserve price chosen by type $v_S + \delta$. We have, for any δ ,

$$v_S \leq \frac{T_S(\rho(v_S + \delta)) - T_S(\rho(v_S))}{P_S(\rho(v_S + \delta)) - P_S(\rho(v_S))} \quad (19)$$

where the right-hand side of (19) always exists, because by assumption both $\rho(\cdot)$ and $P_S(\cdot)$ are strictly monotone. Now, let $\delta \rightarrow 0$. Because ρ is strictly monotone and $T_S(\cdot)$ and $P_S(\cdot)$ are differentiable, we have

$$\lim_{\delta \rightarrow 0} \frac{T_S(\rho(v_S + \delta)) - T_S(\rho(v_S))}{P_S(\rho(v_S + \delta)) - P_S(\rho(v_S))} = \frac{T'_S(\rho(v_S))}{P'_S(\rho(v_S))} \quad (20)$$

The ratio of derivatives on the right-hand side of (20) always exists, because by assumption $T_S(\cdot)$ is differentiable, and $P_S(\cdot)$ is strictly monotone and differentiable, so $P'_S(\boldsymbol{\rho}(v_S)) \neq 0$ for all v_S . Thus, the bound in (19) becomes:

$$v_S \leq \frac{T'_S(\boldsymbol{\rho}(v_S))}{P'_S(\boldsymbol{\rho}(v_S))} \quad (21)$$

Next, applying the bound in (3), we also have, for any δ ,

$$v_S \geq \frac{T_S(\boldsymbol{\rho}(v_S)) - T_S(\boldsymbol{\rho}(v_S - \delta))}{P_S(\boldsymbol{\rho}(v_S)) - P_S(\boldsymbol{\rho}(v_S - \delta))} \quad (22)$$

Analogously, taking the limit as $\delta \rightarrow 0$, we have:

$$v_S \geq \frac{T'_S(\boldsymbol{\rho}(v_S))}{P'_S(\boldsymbol{\rho}(v_S))} \quad (23)$$

Combining (21) and (23), along with the fact that $\boldsymbol{v}_S(r)$ is the inverse of $\boldsymbol{\rho}(v_S)$, yields the desired result for any v_S in the interior of $[\underline{M}, \bar{M}]$.

Now, for $v = \bar{M}$, recall that we have assumed $\boldsymbol{\rho}(v_S)$ is continuous, and thus

$$\boldsymbol{\rho}(\bar{M}) = \lim_{\epsilon \rightarrow 0} \boldsymbol{\rho}(\bar{M} - \epsilon) = \lim_{\epsilon \rightarrow 0} \frac{T'_S(\bar{M} - \epsilon)}{P'_S(\bar{M} - \epsilon)}$$

We have assumed T_S and P_S are continuously differentiable, hence,

$$\lim_{\epsilon \rightarrow 0} \frac{T'_S(\bar{M} - \epsilon)}{P'_S(\bar{M} - \epsilon)} = \frac{T'_S(\bar{M})}{P'_S(\bar{M})}$$

We have thus shown that $\boldsymbol{\rho}(\bar{M}) = \frac{T'_S(\bar{M})}{P'_S(\bar{M})}$. The proof that $\boldsymbol{\rho}(\underline{M}) = \frac{T'_S(\underline{M})}{P'_S(\underline{M})}$ is analogous. This completes the proof of Corollary 2. \square

A.2 Proof of Theorem 2

Proof. We prove this result by contradiction. Suppose that the graph of $\{(P_S(r), T_S(r))\}$ is not convex. Then there exists a triple r, r', r'' , all of which are played with positive probability in equilibrium, such that

$$\gamma P_S(r') + (1 - \gamma) P_S(r'') = P_S(r) \quad (24)$$

for $0 \leq \gamma \leq 1$, and

$$T_S(r) > \gamma T_S(r') + (1 - \gamma) T_S(r'') \quad (25)$$

Consider the type v_S whose optimal action is r . By playing r , her expected utility is

$$v_S P_S(r) - T_S(r) \quad (26)$$

If she were to instead play r' with probability γ and r'' with probability $(1 - \gamma)$, her expected utility would be

$$v_S [\gamma P_S(r') + (1 - \gamma) P_S(r'')] - [\gamma T_S(r') + (1 - \gamma) T_S(r'')] \quad (27)$$

Plugging (24) into (27) yields $v_S P_S(r) - [\gamma T_S(r') + (1 - \gamma) T_S(r'')]$, which, by (25), is strictly greater than (26). Because r is optimal for type v_S , this yields a contradiction, and thus the graph of $\{(P_S(r), T_S(r))\}$ is convex. \square

A.3 Proof of Theorem 3

Before proving this result, we first state the technical conditions required for the convolution argument. These are that (i) the characteristic functions of f_R and f_W have only isolated real zeros and (ii) the real zeros of the characteristic function of f_{P^A} and the real zeros of its derivative are disjoint. These are weak conditions derived in [Evdokimov and White \(2012\)](#)—weaker than those of [Li and Vuong \(1998\)](#)

or [Krasnokutskaya \(2011\)](#) while still yielding the same identification result.

We also introduce one piece of notation: Let $\boldsymbol{\rho}^w(v_S)$ represent the secret reserve price chosen by a seller of type v_S in a game where the realization of $W = w$.

Proof. By Proposition 1.iv, we have that the seller's secret reserve price increases additively with unobserved heterogeneity; that is,

$$\tilde{r} = \boldsymbol{\rho}^w(v_S) = \boldsymbol{\rho}^0(v_S) + w$$

Because the seller's value V_S is independent of buyers' bids conditional on W , $\boldsymbol{\rho}^0(V_S)$ is also independent of auction prices P^A conditional on W . Thus, by [Evdokimov and White \(2012\)](#), F_W , F_R , and F_{P^A} are identified from the joint distribution of $\tilde{R} = R + W$ and $\tilde{P}^A = P^A + W$.

Next, we show that $P_S^0(r)$ and $T_S^0(r)$ are identified from $\tilde{P}_S(\tilde{r})$, $\tilde{T}_S(\tilde{r})$, F_W , and F_R . We describe the identification steps separately for probabilities and transfers.

Probabilities. The probability of trade contaminated with W , $\tilde{P}_S(\tilde{r})$, can be written as:

$$\tilde{P}_S(\tilde{r}) = E_{R,V_{-S},W} [\boldsymbol{x}_S(R + W, V_{-S} + W) \mid R + W = \tilde{r}] \quad (28)$$

$$= E_{R,V_{-S},W} [E_{V_{-S}} [\boldsymbol{x}_S(R + W, V_{-S} + W) \mid W, R + W = \tilde{r}] \mid R + W = \tilde{r}] \quad (29)$$

$$= E_{R,W} [P_S^W(R + W) \mid R + W = \tilde{r}] \quad (30)$$

$$= E_{R,W} [P_S^0(R) \mid R = \tilde{r} - W] \quad (31)$$

Expression (29) follows from applying the law of iterated expectations to (28). (30) follows from taking the expectation over V_{-S} , and using the definition of $P_S^w(\cdot)$ in Section 4.3, and using that R and W are constant after conditioning on W and $R + W = \tilde{r}$. The equality between (30) and (31) follows because of location invariance

in Proposition 1.iv, which implies that

$$P_S^w(r+w) = P_S^0(r) \quad \forall w, r;$$

that is, the probability of trade attained by setting reserve price $r+w$ when unobserved heterogeneity is $W = w$, is the same as the probability attained by setting reserve price r when unobserved heterogeneity is $W = 0$. This allows us to replace $P_S^W(R+W)$ in (30) with $P_S^0(R)$ in (31).

In integral form, expression (31) is (8) in the main text. In words, (8) shows that $\tilde{P}_S(\tilde{r})$ is essentially a noisier version of $P_S^0(r)$: it is a combination of values of $P_S^0(r)$, for \tilde{r} close to r , equal to $P_0^S(r)$ convolved against the function

$$\frac{f_R(r) f_W(\tilde{r} - r)}{\int f_R(r) f_W(\tilde{r} - r) dr}. \quad (32)$$

Because the distributions of R and W both have bounded support, we can set $P_0^S(r)$ to 0 for all r outside the support of R . Thus, both $P_0^S(r)$ and (32) are in L^1 , and hence the convolution theorem applies, meaning that the convolution of $P_0^S(r)$ and (32) is invertible and hence $P_S^0(r)$ is identified from (32) and $\tilde{P}_S(\tilde{r})$.

Transfers. Let $price^w(r)$ represent the average trade price (conditional on trade) when $W = w$ and $R = r$. From Proposition 1.iv, the game is location-invariant; thus, average prices increase one-for-one with changes in w ; that is,

$$price^w(r+w) = price^0(r) + w \quad \forall w, r \quad (33)$$

The seller is paid $price^w(r+w)$ with probability $1 - P_S^w(r+w)$, i.e., one minus the probability $P_S^w(r+w)$ that the seller *keeps* the car. Thus, the expected transfer the seller “pays”, $T_S^w(r+w)$, can be written in terms of the average trade price and the probability that the seller keeps the good as $T_S^w(r+w) = price^w(r) (P_S^w(r+w) - 1)$.

From (33), and from $P^0(r) = P^w(r + w) \forall w, r$, we have

$$T_S^w(r + w) = (\text{price}^0(r) + w) (P_S^0(r) - 1) = T_S^0(r) + w (P_S^0(r) - 1) \quad \forall w, r \quad (34)$$

The transfer function contaminated with W , $\tilde{T}_S(\tilde{r})$, can then be written as

$$\tilde{T}_S(\tilde{r}) = E_{R, V_{-S}, W} [\mathbf{t}_S(R + W, V_{-S} + W) \mid R + W = \tilde{r}] \quad (35)$$

$$= E_{R, V_{-S}, W} [E_{V_{-S}} [\mathbf{t}_S(R + W, V_{-S} + W) \mid W, R + W = \tilde{r}] \mid R + W = \tilde{r}] \quad (36)$$

$$= E_{R, W} [T_S^W(R + W) \mid R + W = \tilde{r}] \quad (37)$$

$$= E_{R, W} [T_S^0(R) + (\tilde{r} - R) (P_S^0(R) - 1) \mid W = \tilde{r} - R] \quad (38)$$

These equations are similar to expressions (28) to (31) above. (36) follows from applying the law of iterated expectations to (35). (37) then follows from the definitions of $T_S^w(\cdot)$ and $P_S^w(\cdot)$, noting again that R and W are constant after conditioning on W and $R + W = \tilde{r}$. Finally, (38) follows from applying (34) to (37).

In integral form, (38) is equivalent to (9) and (10) from the main text. In (9), the term $\tilde{T}_S(\tilde{r})$ is identified in the data and the term $E[W \Delta P_S \mid \tilde{r}]$ can be calculated for any \tilde{r} using (10) because $P_S^0(r)$ is identified; thus, the left-hand side of (9) is known. The right-hand side of (9) is a convolution of $T_S^0(r)$ against (32). By the convolution theorem, this is invertible, and thus $T_S^0(r)$ is identified. \square

B Additional Estimation Results and Discussion

B.1 Additional Details from Main Estimation Steps

The first step of our estimation regresses auction prices and reserve prices on observable features of the game. This step could have been performed instead using only auction prices or only reserve prices as the left-hand-side variable of interest, and our model suggests that we should obtain equivalent estimates of the predicted

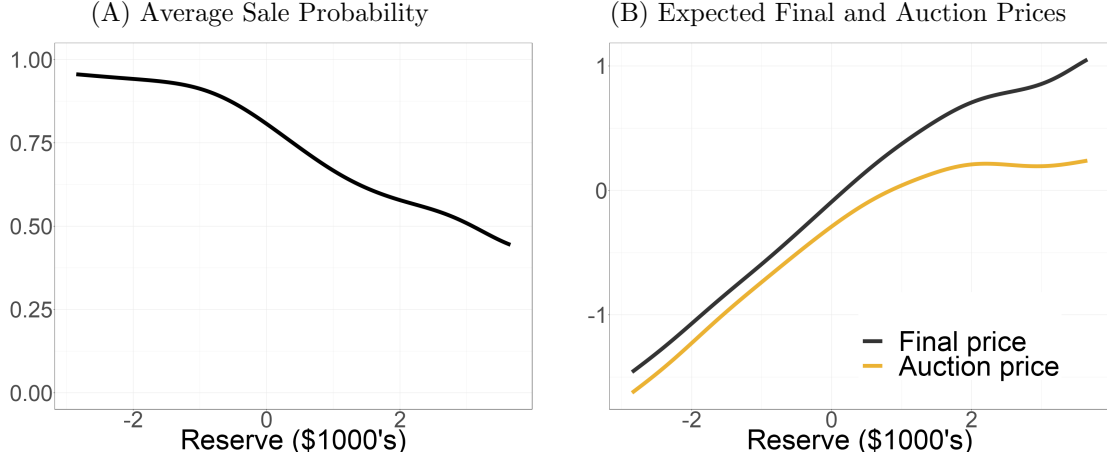
market value $Y\hat{\beta}$ regardless. When limiting to one or the other of these prices, we find estimated predicted values that are highly correlated with our main estimate (a correlation coefficient above 0.98). When we instead examine the correlation of residual reserve prices from (11) under these alternative approaches, we find that the residuals have a lower correlation coefficient (0.78) across methods. The results are similar for residual auction prices (a correlation of 0.79). Because this step is followed by a number of other estimation steps, differences in early stages of estimation can still matter for our final results, and indeed we find that our main estimates of $(\hat{\alpha}_B, \hat{\alpha}_S)$, which are (0.64, 0.62), change to (0.53, 0.62) if we use only auction prices in the observable heterogeneity step and to (0.13, 0.78) if we use only reserve prices. We choose to use the pooled sample in our regressions to include all available information.

We now examine the residuals from this regression in more detail. In the left panel of Figure A.1 we show the probability of sale as a function of sellers' residualized reserve prices, \tilde{R} , estimated via a local linear regression. This probability corresponds to $1 - \tilde{P}_S(\tilde{r})$, as $\tilde{P}_S(\tilde{r})$ is the probability of the seller *keeping* the good. The units for the horizontal axis are \$1,000, and these numbers can be negative because they are the result of subtracting off the market value estimate $y'_j\hat{\beta}$; these numbers can thus be thought of indicating where the reserve price lies relative to the market value estimate of the car. The right panel of Figure A.1 displays, on the vertical axis, the auction price and the final negotiated price, again from a local linear regression against reserve prices.⁴⁰ Here we observe that higher reserve prices are associated with expected final prices that represent a higher markup over the auction price.

Figure A.1.B demonstrates that, after accounting for observable heterogeneity, auction prices remain correlated with reserve prices, suggesting that it is important to account for unobserved heterogeneity in our setting. We also see in panel A that

⁴⁰Recall that auction prices and final prices will not necessarily coincide with one another because of the bargaining component of the mechanism. Also, the final price exists only for observations of the game that end in trade, whereas the auction price always exists. The regressions in each panel use observations with reserve prices lying between the 0.01 and 0.99 quantiles of empirical reserve prices.

Figure A.1: Sale Probability, Auction Price, and Final Price by Reserve Price



Notes: Panel A displays local linear regression estimates of an indicator for whether the car sold regressed on the secret reserve price residual. Panel B contains similar local linear regression estimates for the auction price (the high bid from the auction) in yellow and the final price (conditional on a sale occurring) in black. Units are in terms of \$1,000, relative to the market value estimate. Uses main sample used in the body of the paper.

sellers who post higher reserve prices sell with lower probabilities, but are able to attain higher prices conditional on sale. In particular, the difference between the average final price conditional on sale and the average auction price is increasing in the reserve price. The average auction price roughly measures the value of unobserved car-level heterogeneity conditional on the reserve price.⁴¹ The fact that the *difference* between the final price and the auction price is *increasing* in the reserve price suggests that sellers who post high reserve prices are forgoing some probability of selling the good in order to obtain a higher sale price, as the menu approach requires. We contrast these results to those using observations in which the seller is a car dealership (which we do not use in our main sample) in Appendix B.4.

⁴¹To see this, note that (i) $\tilde{P}^A = P^A + W$; (ii) $\tilde{R} = R + W$; and (iii) P^A, R, W mutually independent together imply that:

$$E[\tilde{P}^A | \tilde{R}] = E[P^A + W | R + W] = E[P^A | R + W] + E[W | R + W] = E[P^A] + E[W | \tilde{R}]$$

That is, $E[\tilde{P}^A | \tilde{R}]$ is equal to $E[W | \tilde{R}]$, the conditional expectation of unobserved heterogeneity W given \tilde{R} , plus the constant $E[P^A]$.


In estimating the buyer value distribution, we use an auction-by-auction upper bound on the number of bidders from our bid log sample to construct an estimate of $\Pr(N = n)$. Here we differ from [Larsen \(2021\)](#), who uses instead an auction-by-auction lower bound on N .⁴² Our choice is motivated by our goal to explore how bargaining power varies with the degree of bidder competition, and the lower bound on N offers far less variance in this dimension than the upper bound—it nearly always equals 2 or 3. However, [Larsen \(2021\)](#) offers empirical evidence that, while the underlying estimate of F_B is naturally sensitive to the choice of $\Pr(N = n)$, the estimated *maximum order statistic distribution* $\hat{F}_B^{(1)}(y)$ (which is what we use in constructing bargaining power estimates) implied by the estimated \hat{F}_B and by $\widehat{\Pr}(N = n)$ is relatively insensitive to this choice.⁴³

To see whether this choice for $\Pr(N = n)$ drives our results, we estimated our model using instead the lower bound on the number of bidders. The results are shown in [Figure A.2](#), analogous to those in [Figure 5](#) in the body of the paper. We do not show results in which we split by the number of bidders because of the lack of variance in the N lower bound. These results differ slightly from those in the body of the paper in that here we observe a slightly lower α_S and slightly higher α_B , and the confidence intervals surrounding the two bargaining power estimates are disjoint, unlike those in the body of the paper. In both [Figures A.2](#) and [5](#), the bargaining is more efficient than a random ultimatum game. The implications from the comparison of manufacturers to non-manufacturers are the same in [Figures A.2](#) and [5](#): manufacturer sellers have more bargaining power than the buyers they negotiate with, and more bargaining

⁴²As described in Larsen’s work, this lower bound, in a given auction, is the sum of the number of unique bidders who bid online (bidder identities are observed for online bidders) plus 1 if the the bid log records any physically present bidders (bidder identities are not recorded for these bidders) or plus 2 if the log records two consecutive physical bids. This lower bound treats all physically present bids as having come from a single bidder unless there are two such bids in a row, motivated by the intuition that no bidder should bid against herself and so two consecutive physical bids must correspond to an auction with at least two bidders physically present.

⁴³[Larsen \(2021\)](#) shows that this insensitivity is not just an empirical artifact: for some choices of $\Pr(N = n)$ (Poisson), the inferred $F_B^{(1)}$ has a derivative with respect to changes in $\Pr(N = n)$ that is identically zero, meaning that counterfactual exercises that rely on the maximum order statistic distribution rather than on F_B can be insensitive to how $\Pr(N = n)$ is specified.

Figure A.2: Bargaining Power Estimates, Adopting Number of Bidders Lower Bound

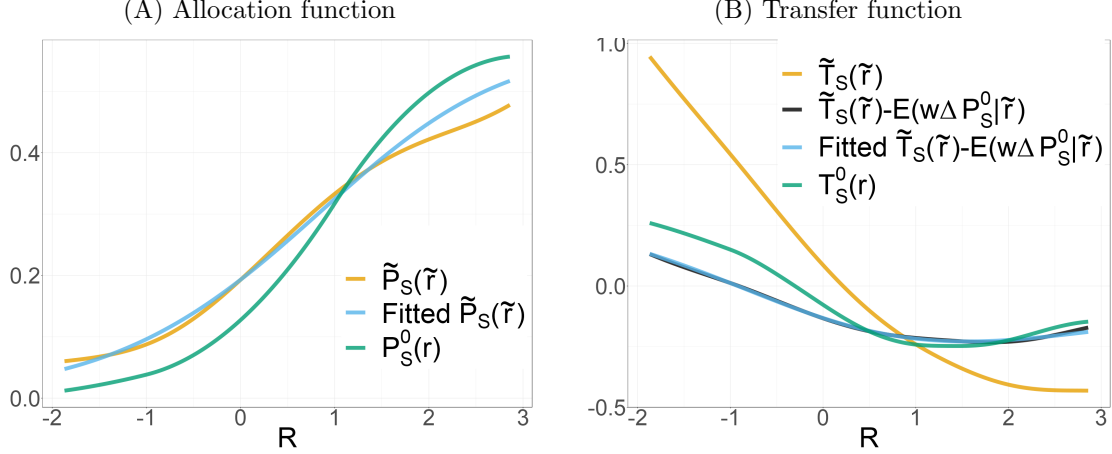


Notes: Estimates as in Figure 5 but using lower bound on N instead of upper bound. Figure shows estimates of $\hat{\alpha}_B$ (yellow), $\hat{\alpha}_S$ (blue), the sum $\hat{\alpha}_B + \hat{\alpha}_S$ (green), and the sum $\alpha_B + \alpha_S$ in the first-best mechanism (purple), for the main dataset (leftmost), manufacturer and non-manufacturer sales, and a comparison sample: non-manufacturer sales with similar average mileage and average age to manufacturer sales. To construct the first-best sum $\alpha_B + \alpha_S$, within each subsample, we take first-best total surplus, and divide it between buyers and sellers such that the ratio α_S/α_B is the same as in the real-world mechanism. Points represent estimates in the baseline sample, and the confidence bars represent 95% confidence intervals from 200 nonparametric bootstrap replications.

power than non-manufacturer sellers.

We next examine the local linear regression and spline estimation steps. In Figure A.3, we show the local linear estimates of $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$, as well as the heterogeneity-corrected estimates $\hat{P}_S^0(r)$ and $\hat{T}_S^0(r)$. We also display intermediate steps in this unobserved heterogeneity correction to illustrate the procedure. For probabilities, the $\tilde{P}_S(\tilde{r})$ function is essentially a noisy version of the $P_S^0(r)$ function; thus, correcting for unobserved heterogeneity will yield an estimate of $P_S^0(r)$ that is steeper than $\tilde{P}_S(\tilde{r})$. This can be seen in panel A by comparing the $P_S^0(r)$ line to the $\tilde{P}_S(\tilde{r})$ line. For transfers, unobserved heterogeneity necessitates two corrections to the $\tilde{T}_S(\tilde{r})$ function. First, we subtract from $\hat{\tilde{T}}_S(\tilde{r})$ the term $\hat{E}[W\Delta P_S | \tilde{r}]$, which represents the expected value of the unobserved heterogeneity conditional on \tilde{r} . Intuitively, for higher values of \tilde{r} , we will observe that trades tend to happen at higher prices, but much of this is due to the unobserved heterogeneity term W being higher on average rather than the transfer $T_S^0(r)$ being higher. In panel B, comparing the

Figure A.3: Removing Unobserved Heterogeneity from Allocation/Transfer Functions



Notes: Figure displays heterogeneity correction for allocation function (Panel A) and transfer function (Panel B). Yellow lines display the original uncorrected estimates of $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$ from local linear regressions, and green lines display final, corrected estimates, $\hat{P}_S^0(r)$ and $\hat{T}_S^0(r)$. In panel B, the black line (which is very close to the blue line) displays estimates from intermediate step subtracting off mean of unobserved heterogeneity, $\hat{T}_S(\tilde{r}) - \hat{E}(W\Delta P_S^0 | \tilde{r})$. In each panel, the blue line displays the fitted value for comparison. Units on the horizontal axis (and vertical axis of panel B) are \$1,000, relative to the market value estimate.

The $\tilde{T}_S(\tilde{r})$ line to the $\hat{T}_S(\tilde{r}) - E[W\Delta P_S | \tilde{r}]$ line shows that this correction makes the slope of the expected transfer function significantly less negative. Secondly, $\hat{T}_S^0(r)$ is essentially a noise-corrected version of $\hat{T}_S(\tilde{r}) - \hat{E}[W\Delta P_S | \tilde{r}]$, and thus the slope and concavity of $\hat{T}_S^0(r)$ are both larger in absolute value than the noisy version. The net effect is that $\hat{T}_S^0(r)$ is much less negatively sloped—and somewhat more concave—than the original nonparametric estimate $\hat{T}_S(\tilde{r})$. In each panel, the blue line displays the fitted estimates, constructed by the convolution of the estimated allocation or transfer function against \hat{F}_W ; in each case, the estimate aligns closely with the local linear estimates. Quantitatively, the RMSE of the fitted $\hat{\tilde{P}}_S(\tilde{r})$ function is 0.012, and the RMSE of the $\hat{T}_S(\tilde{r}) - \hat{E}[W\Delta P_S | \tilde{r}]$ function is \$3.

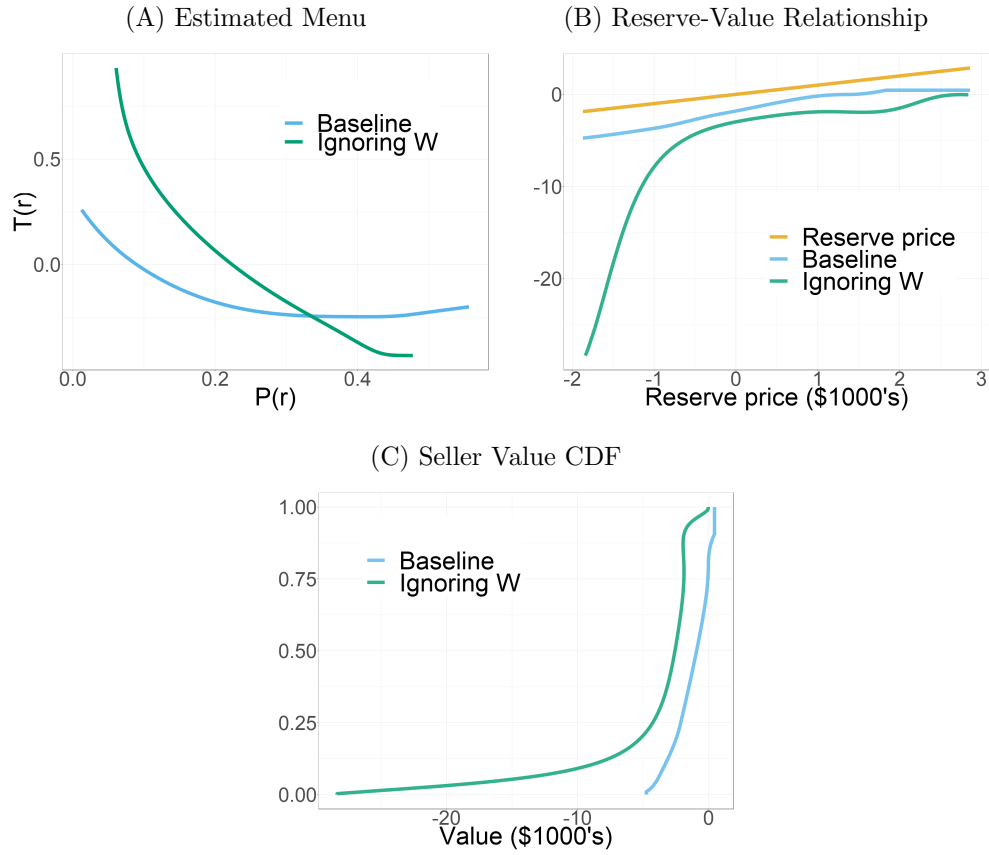
B.2 Evaluating the Impact of the Unobserved Heterogeneity Correction

In this section we evaluate the importance of the unobserved heterogeneity correction for our analysis. First, recall that Figure A.1.B in Appendix B.1 demonstrates that, after accounting for observable heterogeneity, auction prices are still highly correlated with reserve prices, suggesting that it is important to account for unobserved heterogeneity in our setting. To analyze this in more depth, we repeat our full analysis but ignoring unobserved heterogeneity, treating the estimated expected allocation and transfer functions, $\hat{\hat{P}}_S(\tilde{r}), \hat{\hat{T}}_S(\tilde{r})$, as if they constitute the true menu. We then proceed as in the main estimation steps, numerically differentiating this menu to estimate the distribution of seller values. Figure A.4 shows results analogous to those from Figures 3.A, 4.A, and 4.B, comparing estimates accounting for and ignoring unobserved heterogeneity.

Figure A.4.A shows how the unobserved heterogeneity correction affects the seller menu. The menu is much steeper without the unobserved heterogeneity correction. This is because when we observe in the data that a seller sets a higher reserve price (and thus achieves a lower probability of sale), the average sale price is higher for two reasons. First, the seller is higher up on the menu (P_S^0, T_S^0) , attaining higher prices in exchange for lower probabilities of sale. Second, the conditional expectation of the unobserved heterogeneity component W is larger, and hence buyers' bids are also higher in dollar terms because such cars are better in a way which is observed by sellers and buyers, but not by the econometrician. Recall that the difference between the slopes of the yellow and black lines in Figure A.1.B (in Appendix B.1) measures the size of this second force.

Panel B of Figure A.4 shows the implied mappings from reserve prices to sellers' values, and panel C shows the implied distributions of seller values. Because sellers' values are simply derivatives of the seller menu, without the unobserved heterogeneity correction, we infer that sellers' values are much lower: almost all sellers' values are negative, implying that they value the car less than the predicted market value of

Figure A.4: Estimation With and Without Unobserved Heterogeneity Correction



Notes: Estimates of the seller menu and seller value distributions, with and without the unobserved heterogeneity correction. Panel A shows the baseline menu estimate (blue) and the menu estimate using the local polynomial estimates of allocation and transfer functions, $\hat{P}_S(\bar{r})$, $\hat{T}_S(\bar{r})$, without the unobserved heterogeneity correction (green). Panel B shows the mapping from reserve prices to seller values from the two menus, and panel C shows the CDFs of seller values.

the car. Without the unobserved heterogeneity correction, we estimate that sellers' average values are -\$4,307, with a nontrivial fraction of sellers having values from -\$30,000 to -\$10,000, which appear to be unrealistically far below the market value of the car. Thus, the unobserved heterogeneity correction appears to be important to account for in order to obtain reasonable estimates of sellers' reservation values.


B.3 Other sample splits

We now explore how agents' bargaining power varies with characteristics of the game or players, including the car's age, mileage, and condition, and the seller's experience. Car age is measured as the difference between the year the car is offered for sale by the auction house and the model-year of the car. Condition is based on an inspection performed by the auction house prior to the auction. We measure a seller's experience as the cumulative number of times to date (at the time the seller offers a given car for sale) that the seller has participated in the mechanism.

Separately for each of these characteristics, we split the data into terciles based on the values of the characteristic (except condition, where we split into high and low) and then run our estimation routine within a given subsample to obtain estimates of bargaining power. The results are shown in Figure A.5. We find that, relative to the full sample estimates, sellers have less bargaining power than buyers in cases where the car is in poor condition. Buyer bargaining power has a non-monotonic relationship with age and mileage of cars, whereas seller bargaining power tends to increase with age and mileage. Transactions of the oldest cars exhibit high seller bargaining power and low buyer bargaining power. Figure A.5 demonstrates that sales with more experienced sellers do not exhibit more seller bargaining power, but instead buyers in these transactions have more bargaining power.

In most subsamples of the data, we find that $\hat{\alpha}_S + \hat{\alpha}_B > 1$, suggesting that the efficiency of the auction-plus-bargaining mechanism relative to the random-proposer game is quite robust. The one exception is low-age (new) cars, where the 95% con-

Figure A.5: Bargaining power estimates, split samples



Notes: Estimates of $\hat{\alpha}_B$ (yellow), $\hat{\alpha}_S$ (blue), the sum $\hat{\alpha}_B + \hat{\alpha}_S$ (green), and the sum $\alpha_B + \alpha_S$ in the first-best mechanism (purple), for different data subsamples. The x-axis shows different sample splits. The leftmost set of bars shows the full sample, and the other bars show results for different sample splits, by the car's condition, mileage, or age, and by the seller's experience. To construct the first-best sum $\alpha_B + \alpha_S$, within each subsample, we take first-best total surplus, and divide it between buyers and sellers such that the ratio α_S/α_B is the same as in the real-world mechanism. Points represent estimates in the baseline sample, and the confidence bars represent 95% confidence intervals from 200 nonparametric bootstrap replications.


confidence interval for this sum lies below 1. The confidence interval surrounding the first-best, vertically integrated outcome overlaps that of the real-world bargaining in most subsamples, but not for poor-condition cars, low- or medium-age cars, or transactions with experienced sellers, where the estimates suggest some efficiency loss relative to the fully efficient outcome.

B.4 Dealers vs. Fleet/Lease Sellers

Our study focuses on cars sold by fleet/lease sellers because we are particularly interested in bargaining power between dealers and manufacturers or other large institutional sellers. In Figure A.1 we provided some descriptive evidence that the menu approach is indeed appropriate in this setting, i.e., different secret reserve prices yield different payoffs for sellers, and hence can serve to help separate seller types as our method requires. In Figure A.6, we show these same descriptive results using instead the sample of cars sold by dealers, which we do not use anywhere in the body of the paper. Here we observe a relatively flat probability of trade and an expected final price that represents a relatively constant markup over the auction price, regardless of which reserve price the seller chooses. These results suggest that it would be challenging to use our menu approach to identify seller values in this sample, as this approach exploits a tradeoff sellers face between probability of trade and transfers at different reserve prices—a tradeoff that does not jump out from Figure A.6, unlike Figure A.1.

Consistent with this evidence, when we estimate our model on the dealers sample, we find that menu convexity binds and that, prior to enforcing the IR constraint, the constraint is violated for 44% of sellers, unlike in the fleet/lease sample, where menu convexity does not bind except at a single point and only a small fraction of sellers (9.6%) require the IR-enforcement step of our estimation. This suggests that it would be unwise to attempt to infer seller values based on an assumption of optimally chosen secret reserve prices in the dealers sample. [Larsen \(2021\)](#) takes a different approach,

Figure A.6: Dealer Sellers: Sale Prob., Auction Price, and Final Price by Reserve Price



Notes: Panel A displays local linear regression estimates of an indicator for whether the car sold regressed on the secret reserve price. Panel B contains similar local linear regression estimates for the auction price (the high bid from the auction) in yellow and the final price (conditional on a sale occurring) in black. Units are in terms of \$1,000, relative to the market value estimate. Uses dealers sample.

only partially identifying seller values by imposing a weak rationality assumption on the seller's choice to accept or reject the auction price in the first stage of the bargaining game. The assumptions in [Larsen \(2021\)](#) yield bounds on the distribution of seller valuations that are, unfortunately, too wide to be informative about bargaining power. As highlighted above, we focus only on the fleet/lease sample in this paper.