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

Bradley J. Larsen† and Anthony Lee Zhang‡

May 26, 2025

Abstract

This study quantifies bargaining power in supply-side negotiations with incomplete information, where car dealers negotiate inventory prices with large wholesalers after an auction. We measure an agent's bargaining power by where the agent's expected surplus lies relative to a benchmark mechanism favoring the agent and one favoring the opponent. We consider second-best benchmarks, which account for information constraints, and first-best benchmarks, which do not, as well as benchmarks that account for the effect of competition on bargaining power. We propose a direct-mechanism method for estimating a seller's private value as the gradient of a menu from which she chooses a secret reserve price. Bargaining power weights offer insights about inefficiency, as bargaining is not a zero-sum game when agents have incomplete information. On average, dealers (buyers) have less bargaining power than sellers relative to a benchmark where dealers face no competition. Accounting for the direct effect of competition, dealers have more bargaining power than sellers, achieving close to the highest possible surplus given competition. This holds true when the seller is a manufacturer, a finding that is consistent with manufacturers' recent movement toward direct-to-consumer used-car sales.

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, Zichen Zhao, and Tianyi Zhang for research assistance; Susan Athey, Lanier Benkard, Tim Bresnahan, Gabriel Carroll, Evgeni Drynkin, Matt Gentzkow, Phil Haile, Han Hong, Jakub Kastl, Ashley Langer, Simon Loertscher, Leslie Marx, Andrey Ordin, Mike Ostrovsky, Hrishikesh Relekar, Yijing Ren, Karl Schurter, Paulo Somaini, Caio Waisman, and Ali Yurukoglu for helpful comments; and participants at CRA, Duke University, the FTC, Hitotsubashi University, Nagoya University, Northwestern University, Princeton University, University of Chicago, University of Minnesota, University of Pennsylvania, University of Texas, University of Tokyo, University of Toronto, University of Wisconsin, Stanford University, Yale University, the 2022 NBER Summer Institute IO Meetings, and other conferences for helpful comments.

†Washington University in St. Louis, Olin Business School and NBER; blarsen@wustl.edu.

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

1 Introduction

Surplus division between parties is of interest in many settings: business-to-consumer negotiations, vertical contracting, divisions of cartel rents, estimation of patent violation damages. The party taking home a larger share is traditionally referred to as having more *bargaining power*. Studies over the past decade have demonstrated the importance of accounting for bargaining power in 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 solution (or other complete-information models), 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]

We study bargaining power in the wholesale used-car industry, where parties in a vertical supply relationship negotiate under incomplete information. In this market, used-car dealers buy from large fleet-owning institutions, such as banks, rental car companies, or original equipment manufacturers (OEMs). Each car trades through a mechanism of a secret reserve price ascending auction followed by alternating-offer bargaining whenever the reserve price exceeds the highest bid from the auction (which we refer to as the *auction price*). The data consist of over 90,000 sale attempts. We observe actions taken by negotiating pairs even for cases ending in disagreement. This feature not only makes the setting appropriate for studying bargaining power under incomplete information, such data are necessary in *any* setting if one hopes to distinguish between Nash bargaining and incomplete-information bargaining. With these data, we address the question of how buyers' bargaining power compares to sellers' and how it compares for OEM vs. non-OEM sellers. Accounting for

[1] This 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.

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

incomplete information is critical, as inventory is sold car-by-car, and agents frequently engage in negotiations that later fail, a feature inconsistent with *complete*-information models (e.g., Nash bargaining).

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 weight in the joint product of common-knowledge surpluses. We propose new measures of bargaining power under incomplete-information that quantify where an agent's expected surplus in the real world lies between her expected surplus under a benchmark mechanism she would prefer and one her opponent would prefer. This extends a traditional (complete-information, Nash bargaining) notion of power to the incomplete-information case. We consider several sets of benchmarks, including first-best benchmarks, under which all efficient trades succeed (cases where the buyer values the car more than the seller), and second-best benchmarks, under which some efficient trades fail due to incomplete information. We also consider benchmarks that parse out the effect of bidder competition, which forces the final price to exceed the auction price; we refer to this as the *competition constraint*.

For each set of benchmarks, we consider one preferred by the seller and one preferred by the buyer. For a given pair of benchmarks m, let α_B^m be the buyer's bargaining power and α_S^m the seller's. In the seller's preferred mechanism, $\alpha_S^m = 1$ and $\alpha_B^m = 0$, and vice versa. Any intermediate values are possible, as are weights exceeding 1, which can arise from an opponent not fully exploiting her information rent, or negative weights, which can arise for the same reason or from an inefficient equilibrium.[3] Unlike in Nash bargaining, the sum of weights need not equal one.[4]

Bargaining theory shows that incomplete information gives rise to multiple equilibria, delay, and inefficiency, complicating empirical work. Ausubel et al. (2002) highlighted that

[3]Bargaining games with incomplete information have a continuum of equilibria that are qualitatively very different, with some being very inefficient (e.g., Ausubel and Deneckere 1992) and some very efficient (Ausubel and Deneckere 1993), and with different surplus splits.

[4]When measured relative to second-best benchmarks, the sum can exceed 1: agents can collectively achieve strictly greater expected utility than that available through any convex combination of agent-preferred mechanisms. The sum can be less than one in inefficient equilibria. These weights are thus informative about both the pie split and pie size. In tranferrable utility settings, any relationship between bargaining power and pie size is ignored by Nash bargaining. Loertscher and Marx (2022) explained, "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."

different equilibria can have quite different properties and outcomes, and that no complete characterization of equilibria exists; this statement remains true two decades later. As such, there is no off-the-shelf model for empiricists to bring to bargaining data to identify private value distributions, unlike well-developed auction methods, e.g., Guerre et al. 2000 (GPV).[5] In wholesale used-car markets, the primary identification challenge is the distribution of *seller* values, F_S. Every choice of sellers — even their choice of secret reserve price at the game's start — depends on the bargaining strategies, which are unknown to the econometrician. In contrast, the buyer value distribution, F_B, is identified from auction bids using existing auction methods. These tools also allow us to handle game-level heterogeneity.

We propose to estimate seller values based on an *empirical menu* approach generalizing GPV. We consider a seller of value v_S choosing her secret reserve price, r, to maximize her expected surplus $v_S P_S(r) - T_S(r)$, 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 that the seller chooses r 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 v_S. 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, $P_S(\cdot)$ and $T_S(\cdot)$ are essentially observed in the data, and derivatives of this menu correspond to sellers' values. Our model implies two restrictions that we impose in estimation: the equilibrium menu must be convex and satisfy individual rationality (IR).

Applying these arguments to our data, we estimate the trade-transfer menu faced by sellers in wholesale used-car markets. With the estimated menu and value distributions, we compute average bargaining power for a seller in this market (a large fleet owner, such as Ford, Bank of America, and Hertz) vs. the high-bidder (the buyer, a used-car dealer). Sellers, on average, obtain a surplus that represents 52.5% of the first-best surplus and buyers obtain 41.4%. The shortfall relative to 100% represents inefficiency due to incomplete information

[5] Unlike 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 environments close to the one we study — with continuous values, both parties having private values, and both parties making offers — but the equilibria derived in these studies are not suitable for 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 our data.

(Myerson and Satterthwaite (1983)).

Competition has a direct effect on bargaining outcomes, pushing up final prices. We quantify this by considering as a benchmark a specific split of the first-best surplus in which agents trade at the lowest price possible subject to the competition constraint: the price must exceed both the highest bid from the auction and the seller's value. We find that buyers' real-world surplus represents 93.6% of this first-best constrained benchmark. Sellers' real-world surplus is worse than this benchmark, suggesting that, in the real-world outcome, they are conceding some surplus to buyers. Together, these results imply that what can be termed sellers' *residual* or *constrained* bargaining power is quite low. In the language of some antitrust experts (e.g., Peters 2014), our results imply that sellers in this market have a high degree of *bargaining leverage* due to competition, but little residual bargaining power after accounting for the direct effect of competition. We find similar results when we compare agents' real-world surplus to second-best benchmarks: overall, sellers' surplus is closer than buyers to their preferred second-best mechanism payoff, and this is driven by bidder competition.

We then estimate bargaining power separately for cases where the seller is an OEM, such as Ford or GM, vs. a non-OEM (banks, fleet companies, or rental companies). Since the inception of the wholesale used-car market in the mid-twentieth century, OEMs have only engaged in the used market by selling leased vehicles at the end of the leasing term (referred to as *off-lease* vehicles) to dealers, not directly to consumers. In recent years, OEMs have opened direct-to-consumer (DTC) channels. Whether OEMs will continue this trend may depend in part on how they view their bargaining power in the traditional wholesale market. We quantify this by estimating our model separately for OEM and non-OEM sellers.

We find that the total surplus available is higher for OEM sales, and OEM sellers achieve a higher fraction of that surplus, implying that OEM sellers have more overall bargaining power relative to buyers than do non-OEM sellers. The same is true relative to the seller's second-best benchmark: OEM sellers achieve a payoff that lies 91.8% (and non-OEM sellers 63.1%) of the way between their preferred second-best mechanism payoff and that of buyers. As in the full data sample, however, we find that this effect is driven purely by bidder competition: sellers are unable to push the price significantly higher than the auction price and are even conceding to buyers some of the surplus gained from bidder competition.

Our study relates to a growing body of structural work studying bargaining power in vertical business-to-business settings, such as Crawford and Yurukoglu (2012), Grennan (2013), Gowrisankaran et al. (2015), and Ho and Lee (2019). In contrast to most previous work on vertical relationships, we allow for incomplete information and allow agents to be strategic in their bargaining behavior.[6] Previous work on vertical relationships between dealers and OEMs has focused on new cars (Lafontaine and Scott Morton 2010; Murry and Zhou 2020), whereas we offer insights into the used-car arm.[7]

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 competition authorities. The empirical literature has studied related ideas under assumptions of Nash bargaining (e.g. Capps et al. 2003; Gowrisankaran et al. 2015), but not under incomplete information. One would expect increased competition among buyers to increase the sellers' (and decrease buyers') bargaining power, but it is unclear by how much, and whether sellers have any additional bargaining power beyond that afforded them by competition. The seminal result of Bulow and Klemperer (1996) suggests 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, and therefore the benefits to a seller of increased competition may be even greater than suggested by Bulow and Klemperer (1996).[8] Our study contributes to this literature by studying a real-world setting in which both parties potentially have bargaining power and quantifying the degree to which that power is driven by competition.

[6]Relative to empirical work with complete-information, ours is not a strict generalization, and vice-versa. For example, many studies using Nash bargaining, such as Crawford and Yurukoglu (2012), are more general in the dimension of modeling *non-transferable utility* (allowing the downstream firm to have a willingness to pay that depends on the price negotiated with the upstream firm). The model of Ho and Lee (2019) (Nash-in-Nash with threat of replacement) is more general than ours in the dimension of endogenizing agents' outside options. In contrast, our model is more general than both in the dimension of allowing for incomplete information, but we model *transferable utility*, and outside options are fixed: the seller's *private value* is her outside option and the buyer's is her willingness to pay net of any outside option. See Section 4.

[7]Donna et al. (2024) studied vertical bargaining in advertising and the welfare effects of DTC sales.

[8]Consistent with this, Bulow and Klemperer (1996) stated, "No amount of bargaining power is as valuable to the seller as attracting one extra bona fide bidder," and then concluded, "Our analysis assumed that a seller could negotiate optimally, making credible commitments of the sort that might not be possible in real life, and we also assumed that bidders had no bargaining power in a negotiation. We therefore believe that our basic result does not overstate the efficacy of auctions relative to negotiations."

Several empirical bargaining studies contain structural models of two-sided incomplete information. Keniston (2011) studied welfare under bargaining vs. a posted price. Kong et al. (2024) studied arbitration in union wage negotiations. Larsen (2021) analyzed empirical implications of the Myerson-Satterthwaite Theorem using a superset of the data from our paper. Freyberger and Larsen (2025) studied inefficient impasse in eBay bargaining. We see our focus on *equity* — how the surplus is split — as a natural next question to address after the *efficiency* questions of the latter two studies, which derived bounds on surplus or trade probabilities but did not address the question of surplus division. Indeed, their bounds, while informative about inefficiency, are too wide to be informative about the division of surplus. In contrast, our paper derives point estimates of this split.[9]

Our contribution to the structural methodology literature can be seen as generalizing GPV to bargaining games. In a related study, Kline (2023) focused on identification, but not estimation, in a class of games that overlaps with the class we study: trading games with monotone equilibria. Agarwal et al. (2023) derived results that nest a number of related identification arguments, including ours.[10] 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 bargaining (e.g., Myerson and Satterthwaite 1983; Williams 1987; Loertscher and Marx 2022), which abstracts away from extensive-forms.

2 Background: Supply-Side Bargaining for Used Cars

The U.S. wholesale used-car industry — with revenues above \$110 billion annually — operates through a network of several hundred auction house locations scattered throughout

[9] We borrow some straightforward steps of Larsen (2021), including how we control for game-level heterogeneity and estimate buyer values. Our identification of seller values differs from Larsen (2021): we exploit optimality of the seller's choice of secret reserve price, yielding point identification, whereas the former study exploited the seller's choice to accept or reject the auction price, yielding bounds.

[10] Related arguments also appear in Perrigne and Vuong (2011) and Luo et al. (2018). Pinkse and Schurter (2019) introduced efficient estimation procedures for auctions and related games, which, like ours, exploit convexity restrictions implied by incentive compatibility.

the country (and operations are similar internationally).[11] These auction houses have been a part of the U.S. used-car market for over seventy years. Each year, over 12 million cars pass through these auction houses, where used-car dealers buy from rental companies, banks with repossessed vehicles, or manufacturers with off-lease vehicles.[12]

Several days before a weekly sale, a seller brings her car to the auction house and reports a secret reserve price; this reserve price is typically never revealed to bidders, even after the auction. On the sale day, buyers (dealers) arrive, with many traveling long distances to attend. Remote bidders participate virtually. Cars are auctioned in their arrival order. Multiple auctions run simultaneously in different lanes dividing the auction house. Each car's auction takes about 90 seconds (Lacetera et al. 2016), with bidding run by a human auctioneer raising prices until one bidder remains. If the auction price (the highest bid) exceeds the secret reserve price, the high bidder takes the car. If not, the high bidder may opt out of bargaining. If he does not opt out, he and the seller enter alternating-offer bargaining, mediated by an auction-house employee over the phone.[13] Bargaining continues until one party accepts or quits.

Our data consist of 91,743 realizations of this mechanism from six auction houses owned by the same parent company spanning 2007–2010. For each realization, the key observables are the auction price, secret reserve price, final transaction price, and an indicator for whether the car sold. We also observe a large set of characteristics, including features of the car and the auction house environment at the sale time. Our 91,743 observations are those that remain after the data cleaning steps in Larsen (2021) as well as one additional step in which we keep only the first appearance of a given vehicle at a given auction house: if the sales attempt does not result in trade, the vehicle may appear at a future auction sale, and we drop these future attempts. This drops 28% of sales attempts. Appendix B.1 discusses additional data cleaning steps.

Table 1 shows descriptive statistics. The average car has a book value (an estimate pro-

[11] https://www.naaa.com/pdfs/auction_industry_surveys/2023_NAAA_Auction_IndustrySurvey.pdf.

[12] Sellers can also be dealers, but not in the data we use herein. See Appendix B.8.

[13] Larsen et al. (2024) studied these mediators. If the auction price is far below the reserve price, the auctioneer does not proceed to bargaining; we treat these cases as though the seller rejects the auction price (because we have no data to distinguish the difference).

Table 1: Descriptive Statistics (Sample Size = 91,743)

A.	Mean	Standard Deviation	Seller Category	Frac. of Sample
Book value ($)	11,030	6,284	OEM	0.2838
Age (years)	3.26	2.52	Bank	0.5027
Mileage	56,761	39,557	Fleet Company	0.0817
Good condition	0.73	0.44	Lease Company	0.0757
$\underline{N}\mid\underline{N} \geq 2$	3.00	0.45	Rental	0.056
Agree	0.86	0.35		

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	0.43	0.99	11,053	10,197	11,053	5,402	7,547
Period 2	0.51	0.82	9,855	10,672	9,855	9,793	11,165
Period ≥ 3	0.06	0.18	7,465	8,724	7,929	6,411	8,266

Notes: In panel A, "Book value" is an estimate of the car's book value, provided by the auction house. "Good condition" indicates average or above average car condition, based on auction house inspection. "$\underline{N}\mid\underline{N} \geq 2$" is an auction-by-auction lower bound on the number of bidders, only observable in the bid log subsample, and limited to cases where this lower bound is at least 2 (71,870 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 $\geq$ 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.

vided by the auction house) of $11,030, is 3.26 years old (relative to its model-year), and has 56,761 miles on the odometer. The auction house provides a condition report for most cars, and 73% of cars are rated at average quality or above, which we indicate in panel A with "Good condition." Our data also contain detailed records (referred to as *bid logs*) of the bidding during the auction stage for most observations (73,100, with 71,870 having a lower bound of at least two). In this sample, we obtain bounds on the number of bidders (N) in each auction, with an average lower bound of 3 bidders (see Section 5.1 and Appendix B.3). The mean of the "Agree" variable (an indicator for whether trade occurs) is 0.86, implying that 14% of sales attempts result in no trade. These failed trades are inconsistent with a standard complete-information framework, where a buyer and seller do not engage in a trading game knowing a priori that they will disagree. Failed negotiations, however, are consistent with the presence of incomplete information (Myerson and Satterthwaite 1983; Perry 1986). The final column in panel A shows that OEMs, such as Ford, represent 28.38% of sales. Banks, such as Citibank or Bank of America, represent a slight majority, at 50.27%. Fleet companies (such as Wheels) represent 8.17%, rental companies (such as Budget) 5.6%, and other lease companies 7.57%.

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: agreement, and the auction, secret reserve, and final prices. The first row shows outcomes for games ending without bargaining, which occurs in 43% 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 1% of the time).[14] The second row, indicated by *Period 2*, refers to cases where the reserve price exceeds the auction price, which the seller either accepts (51% of the time) or rejects (49%). The third row refers to games that end at some later period (6% of observations).[15] 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 $7,465, the average reserve price is $8,724, and the average final price is between the two, at $7,929. When trade fails (the final two columns), the auction price is farther below the reserve price.

These final numbers in the preceding paragraph illustrate an important point: it is a priori unclear how 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, preventing inferences about bargaining power from these bounds alone. Our identification argument infers the distribution of buyer values from auction prices and seller values from reserve prices, trade probabilities, and final prices. With these distributions, we then quantify bargaining power.

[14]The auction house attempts to prevent these opt-outs when the seller is a large fleet/lease institution, as in our data, but we find that it still occurs in practice.

[15]Our identification and estimation use all observations — those that end at the auction or at a later period.

3 Defining Bargaining Power Under Incomplete Information

Our goal is to quantify bargaining power of a seller and buyer who potentially trade. The seller's private value is V_S and the buyer's is $V_{B^{(1)}}$. The wholesale used-car market game is between one seller and many buyers, but, because the auction identifies the highest-value bidder, it is only the seller and high bidder whose surpluses are affected by the game outcome. As such, we use the terms *buyer* and *high bidder* interchangeably. We describe below the role the auction plays in our bargaining power definition.

Incomplete-information bargaining games — where agents' values are private information — are complex to model theoretically, even for seemingly simple extensive forms like alternating offers. Each offer can signal information to the opposing party, who can then update beliefs about the opponent's value. Under standard equilibrium concepts, belief updating after off-path actions can sustain a large set of strategies in sequential 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 has a corresponding *direct mechanism*. Let $\mathcal{M}_{RW}$ be the mechanism corresponding to a real-world equilibrium.[16]

Let $U_B(\mathcal{M})$ and $U_S(\mathcal{M})$ represent the expected surplus of the buyer and seller, respectively, under an arbitrary 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. We propose a two-dimensional measure that quantifies the power of the buyer and seller in $\mathcal{M}_{RW}$ relative to some *reference* or *benchmark* mechanisms: one that is favored by the buyer (denoted $\mathcal{M}_B^m$) and the other by the seller ($\mathcal{M}_S^m$), where m indexes different benchmarks.[17] For any m, bargaining weights α_B^m and α_S^m describe where an agent's expected surplus under $\mathcal{M}_{RW}$ lies between her expected surplus under $\mathcal{M}_B^m$ and $\mathcal{M}_S^m$:

$$U_B(\mathcal{M}_{RW}) = \alpha_B^m U_B(\mathcal{M}_B^m) + (1-\alpha_B^m) U_B(\mathcal{M}_S^m), \quad U_S(\mathcal{M}_{RW}) = \alpha_S^m U_S(\mathcal{M}_S^m) + (1-\alpha_S^m) U_S(\mathcal{M}_B^m)$$

[16] The term *real-world mechanism* is not synonymous with the game's *protocol* — a secret reserve ascending auction followed by alternating offers. Within that protocol are infinitely many equilibria, and a real-world mechanism $\mathcal{M}_{RW}$ refers to the direct mechanism corresponding to a specific equilibrium.

[17] Görlach and Motz (2024) proposed an axiomatic definition of bargaining power, measured for a given agent as the change in outcomes that would arise if all agents were to adopt the preferences of that agent.

$$\Rightarrow \alpha_B^m = \frac{U_B(\mathcal{M}_{RW}) - U_B(\mathcal{M}_S^m)}{U_B(\mathcal{M}_B^m) - U_B(\mathcal{M}_S^m)}, \qquad \Rightarrow \alpha_S^m = \frac{U_S(\mathcal{M}_{RW}) - U_S(\mathcal{M}_B^m)}{U_S(\mathcal{M}_S^m) - U_S(\mathcal{M}_B^m)} \tag{1}$$

We consider several candidate benchmarks. The first involves *first-best*, ex-post efficient trade: trade occurs whenever $V_{B^{(1)}} \geq V_S$, at a price between the two. Among such mechanisms, sellers prefer a price equal to $V_{B^{(1)}}$, giving sellers all of the surplus and buyers zero; we denote this by $\mathcal{M}_S^{1^{st}}$. Similarly, $\mathcal{M}_B^{1^{st}}$ denotes buyers receiving all the first-best surplus.

We also consider a benchmark that constitutes a specific split of the first-best surplus subject to what we will refer to as the *competition constraint*: our setting involves multiple buyers, leading the final transaction price to be weakly greater than the highest auction bid, denoted P^A. Let $\mathcal{M}_B^{1^{st},con}$ denote first-best trade *constrained by competition*: trade occurs whenever $V_{B^{(1)}} \geq V_S$, at a price equal to $\max\{P^A, V_S\}$.[18] For the seller, $\mathcal{M}_B^{1^{st},con}$ yields the same outcome she would achieve in a mechanism consisting of an auction followed by the seller accepting or rejecting P^A (but never countering). Benchmarking real-world bargaining outcomes against $\mathcal{M}_B^{1^{st}}$ and separately against $\mathcal{M}_B^{1^{st},con}$ allows us to quantify how the overall division of surplus between the high bidder and the seller is driven by the direct effect of competition. This speaks to the distinction in antitrust discussions between *bargaining leverage* — the direct effect of competition on bargaining power — and the additional advantage an agent possesses beyond that competitive effect.[19]

We also consider *second-best* versions of these benchmarks. Under incomplete information, Myerson and Satterthwaite (1983) showed that first-best trade is generally infeasible and derived the second-best, ex-ante efficient mechanism maximizing the gains from trade subject to information constraints.[20] Williams (1987) extended this analysis, deriving the full ex-ante efficient Pareto frontier. Let $M_S^{2^{nd}}$ and $M_B^{2^{nd}}$ denote second-best mechanisms favoring sellers and buyers, respectively.[21] Under $A1$, these mechanisms are simple to characterize:

[18]This is the competition-constrained first-best mechanism preferred by the buyer. We could similarly define $\mathcal{M}_S^{1^{st},con}$, but this will mechanically equal $\mathcal{M}_S^{1^{st}}$, as $\mathcal{M}_S^{1^{st}}$ involves a price equal to the high bidder's value, $V_{B^{(1)}}$, and therefore satisfies the competition constraint.

[19]See, for example, Peters (2014), Asil et al. (2024), and https://www.justice.gov/atr/speech/mergers-increase-bargaining-leverage. We thank Tom Wollmann for this point.

[20]Specifically, if the supports of buyer and seller values overlap, no mechanism achieves first-best trade while satisfying incentive compatibility, individual rationality, and ex-ante budget balance.

[21]The term *bargaining power* is used similarly in other settings involving a combination of auctions and bargaining, such as Bulow and Klemperer (1996) and Menezes and Ryan (2005), who described an agent with all of the bargaining power as one who can implement her preferred second-best mechanism. In Loertscher and Marx (2022), who extended the idea of Williams (1987) to cases with multiple agents on one or both

(A1) (i) Agents are risk-neutral with independent, private values (IPV). (ii) Buyers are symmetric. (iii) Agents have weakly increasing virtual values.

$A1$ enumerates the conditions imposed by Myerson (1981), Myerson and Satterthwaite (1983), and Williams (1987).[22] Under $A1$, $M_S^{2^{nd}}$ can be implemented by an auction with an optimal *public* reserve price (Myerson 1981).[23] If the highest-value bidder were to negotiate with the seller *absent* competition from other buyers, $\mathcal{M}_B^{2^{nd}}$ would correspond to a take-it-or-leave-it offer (TIOLIO). Similar to $\mathcal{M}_B^{1^{st},con}$, we let $\mathcal{M}_B^{2^{nd},con}$ denote the second-best mechanism favoring the buyer, subject to the competition constraint: a TIOLIO from the high bidder to the seller where the price weakly exceeds P^A.[24] As in the first-best case, the comparison between $\mathcal{M}_B^{2^{nd}}$ and $\mathcal{M}_B^{2^{nd},con}$ allows us to speak to the direct effect of competition (*bargaining leverage*) and parties' residual bargaining power beyond this effect.

Thus, the four sets of benchmarks we consider are $\mathbb{M} \equiv \{$"1^{st}"; "$1^{st}, con$"; "2^{nd}"; "$2^{nd}, con$"$\}$. Figure 1 shows a geometric interpretation of these bargaining power measures, focusing for simplicity on the unconstrained first- and second-best benchmarks ($m =$ "1^{st}" or "2^{nd}").[25] The figure shows the expected surplus of buyers on the vertical axis and sellers on the horizontal. The two green points represent expected payoffs under the second-best buyer- and seller-optimal mechanisms, $\mathcal{M}_B^{2^{nd}}$ and $\mathcal{M}_S^{2^{nd}}$. The green line illustrates the second-best

sides, bargaining power is defined as the weight an agent receives in a weighted welfare maximization problem.

[22]For seller values distributed according to F_S (with density f_S), virtual values are $\psi_S(v_S) \equiv v_S + \frac{F_S(v_S)}{f_S(v_S)}$. Virtual values of buyers, with CDF F_B and density f_B, are $\psi_B(v_B) \equiv v_B - \frac{1-F_B(v_B)}{f_B(v_B)}$.

[23]The seller-optimal, second-best outcome can equivalently be achieved by an auction with no reserve price followed by a TIOLIO from the seller to the high bidder. We use the term *public* to refer to an auction with a reserve price that the seller commits to before the auction (and that the bidders are aware of), where trade occurs if the highest value exceeds the public reserve price. In contrast, in the secret-reserve auction of used-car markets, trade occurs if the *second-highest value* exceeds the secret reserve price, and otherwise bargaining ensues. While an optimal public reserve auction is the optimal mechanism for a seller (Myerson 1981), sellers cannot necessarily implement this mechanism in practice because it is the auction house that chooses the mechanism. Elyakime et al. (1994) showed that an optimal public reserve price for a seller (who has some value V_S of keeping the car herself) will not necessarily be the optimal mechanism *for an auction house* that only receives payment when trade occurs.

[24]Under $A1$, $\mathcal{M}_B^{2^{nd},con}$ is optimal for the buyer among mechanisms with transfers weakly above P^A. As in the first-best case, we could also define $\mathcal{M}_S^{2^{nd},con} \equiv \mathcal{M}_S^{2^{nd}}$, as $\mathcal{M}_S^{2^{nd}}$ automatically involves trading at a price weakly higher than P^A. Note that, if $A1$ is not satisfied, bargaining power measured relative to a buyer TIOLIO (or competition-constrained TIOLIO) and a public reserve auction would still be interpretable as describing where an agent's expected surplus lies between two benchmark mechanisms, one that the agent prefers over the other and one that the opponent prefers over the other.

[25]The competition-constrained benchmark cases look similar, with payoffs under buyer-preferred mechanisms ($\mathcal{M}_B^{1^{st},con}$ and $\mathcal{M}_B^{2^{nd},con}$) shifted to the right, as competition ensures the seller a higher payoff.

Figure 1: Illustration of Bargaining Power Measure

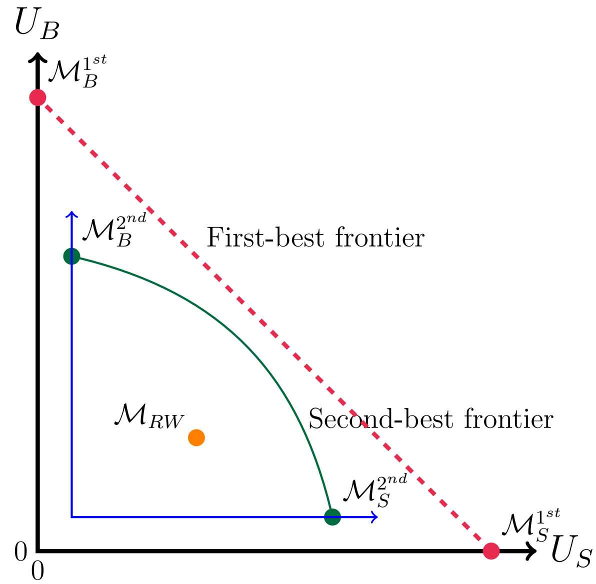

Notes: Visualization of bargaining power. Expected buyer surplus is on the vertical axis and expected seller surplus on the horizontal. The two green points represent expected payoffs under the second-best buyer- and seller-optimal mechanisms, $\mathcal{M}_B^{2^{nd}}$ and $\mathcal{M}_S^{2^{nd}}$. The green line illustrates the second-best, ex-ante efficient Pareto frontier. The red dashed line illustrates the first-best, ex-post efficient frontier. The orange point shows the location of a hypothetical real-world mechanism.

frontier derived by Williams (1987).[26] The red dashed line illustrates the first-best, ex-post efficient frontier — the maximum expected surplus if information were complete or if information constraints were relaxed. The orange point shows a hypothetical $\mathcal{M}_{RW}$. Under complete information, the red and green lines would coincide, as would blue and black.

In Figure 1, the buyer's expected surplus under $\mathcal{M}_{RW}$ lies about one-fourth of the way between the buyer's expected surplus under $\mathcal{M}_S^{1^{st}}$ and $\mathcal{M}_B^{1^{st}}$, while the seller's lies about one-third of the way between what she would expect under $\mathcal{M}_B^{1^{st}}$ and $\mathcal{M}_S^{1^{st}}$, together implying $(\alpha_B^{1^{st}}, \alpha_S^{1^{st}}) = (0.25, 0.33)$. Relative to the second-best benchmarks ($\mathcal{M}_B^{2^{nd}}$ and $\mathcal{M}_S^{2^{nd}}$), the buyer's and seller's expected surplus under $\mathcal{M}_{RW}$ implies $(\alpha_B^{2^{nd}}, \alpha_S^{2^{nd}}) = (0.33, 0.50)$.

These bargaining power metrics have several noteworthy properties. First, for each $m \in \mathbb{M}$, α_S^m and α_B^m are extensions of Nash bargaining weights to incomplete-information bargaining. In a transferable utility model like ours, if agents were to have complete information and

[26]The Pareto frontier is the highest combinations of buyer and seller expected surplus achievable by incentive-compatible, individually rational, budget-balanced mechanisms. It consists of all mechanisms $\mathcal{M}_\eta^{2^{nd}}$ that maximize the weighted sum of welfare, $\eta U_S(\mathcal{M}_\eta^{2^{nd}}) + (1-\eta) U_B(\mathcal{M}_\eta^{2^{nd}})$ for $\eta \in [0,1]$. This welfare weight, η, might itself be thought of as one notion of bargaining power *among second-best mechanisms*, but this notion would not suffice for our purposes; we seek a notion of bargaining power applicable to real-world bargaining situations, which are not guaranteed to achieve payoffs on the frontier.

use Nash bargaining, we could define bargaining power in two equivalent ways: an agent's bargaining power describes (i) the fraction of the total first-best surplus she receives or (ii) where her payoff lies between what she would achieve under her preferred mechanism and under her opponent's preferred mechanism. With incomplete information, the two definitions are not equivalent due to inefficiency: the sum of buyer and seller surplus can fall short of total first-best surplus (Myerson and Satterthwaite 1983). Furthermore, a trade-off generally exists between efficiency and rent-extraction: as one agent gets a larger surplus share, total surplus can shrink (i.e., incomplete-information bargaining is not zero-sum). The sum $\alpha_S^m + \alpha_B^m$ is informative about the extent of this inefficiency: for $m =$ "1^{st}" or $m =$"$1^{st}, con$", $\alpha_S^m + \alpha_B^m < 1$ implies a deadweight loss, as suggested by the Myerson-Satterthwaite Theorem. Nash bargaining, in contrast, requires $\alpha_B^m + \alpha_S^m = 1$, imposing efficiency a priori. For second-best mechanisms, $m =$ "2^{nd}" or $m =$ "$2^{nd}, con$", any convex combination of $\mathcal{M}_B^m$ and $\mathcal{M}_S^m$ corresponds to a mechanism that randomly selects between the buyer's or seller's preferred mechanisms. In Figure 1, these mechanisms lie along a straight line from $\mathcal{M}_B^{2^{nd}}$ to $\mathcal{M}_S^{2^{nd}}$. Thus, for second-best mechanisms m, examining whether $\alpha_B^m + \alpha_S^m$ is less than or greater than 1 sheds light on how efficient the real-world mechanism is relative to this alternative *random-proposer* game.

Second, for some benchmarks, α_i^m for some agent i may be negative because, among infinitely many equilibria, some are inefficient. An $\alpha_B^{2^{nd}} < 0$, for example, can reflect an inefficient equilibrium in which the buyer's expected payoff is even lower than under $\mathcal{M}_S^{2^{nd}}$. A negative α_i^m does not necessarily reflect a negative *utility* for i, but rather that i's real-world utility lies in the gap between the black and blue axes in Figure 1. This gap exists because second-best mechanisms provide an *information rent*: the buyer gets a strictly positive expected payoff even under $\mathcal{M}_S^{2^{nd}}$ (and vice versa).

Similarly, for some benchmarks, α_i^m can exceed 1 if, at some point in the game, i's opponent does not fully exploit her information rent, behaving non-strategically or more generously to i than would be implied by incentive constraints.[27] Saran (2011) provided an incomplete-information bargaining model in which some agents (labeled *naive*) ask for or offer their true value (rather than shading), which can lead to outcomes beyond the information-constrained

[27] In contrast, because first-best yields the maximum surplus, $\alpha_i^{1^{st}} \leq 1$ and $\alpha_i^{1^{st},con} \leq 1$ for $i \in \{S, B\}$.

second-best frontier. Valley et al. (2002) implemented incomplete-information bargaining in a lab experiment, finding that communication among agents can lead to outcomes that lie between the second- and first-best frontiers, a situation that can generate $\alpha_i^m > 1$. Valley et al. (2002) attributed their findings to agents' preferences for fairness or altruism; Keniston et al. (2024) documented evidence consistent with such preferences in a variety of negotiation settings including wholesale used-car auctions. If agent $-i$ (i's opponent) gives up even more than her information rent, this could also lead to $\alpha_{-i}^m < 0$.

Third, as the number of buyers/bidders in the mechanism increases, $\alpha_S^{2^{nd},con}$ and $\alpha_B^{2^{nd},con}$ become more difficult (and less useful) to measure because the gap between an agent's expected surplus under $\mathcal{M}_B^{2^{nd},con}$ approaches her surplus under $\mathcal{M}_S^{2^{nd}}$. For example, solving for $\alpha_S^{2^{nd},con}$ using (1) involves dividing by $U_S(\mathcal{M}_S^{2^{nd}}) - U_S(\mathcal{M}_B^{2^{nd},con})$ and thus, when this number is close to zero, $\alpha_S^{2^{nd},con}$ grows large in magnitude.

Finally, variation in our measures of bargaining power can arise for similar reasons as in empirical work with Nash bargaining. Consider a complete-information world, where $V_{B^{(1)}}$ and V_S are *known* and $V_{B^{(1)}} > V_S$. Let Z (with $V_S < Z < V_{B^{(1)}}$) be a price at which agents trade, where the seller receives Z if trade occurs and V_S otherwise. The buyer receives $V_{B^{(1)}} - Z$ if trade occurs and 0 otherwise. Consider two datasets such that the analyst infers a higher Nash weight for sellers in the first dataset than in the second. A number of channels could lead to this inference. For example, the seller's disagreement payoff (V_S) may be lower in the first dataset: fixing Z and $V_{B^{(1)}}$, a lower V_S would imply that the seller receives a larger share of the surplus, and hence has more bargaining power.[28] Similarly, holding fixed agents' values, a higher Z in the first dataset would imply more seller bargaining power. Unmodeled factors such as agents' patience, bargaining costs, or negotiation expertise could underlie such price differences, and, consequently, differences in inferred power. In empirical work, Nash weights do not correspond to microfoundations of these factors, but instead serve as a residual that allows the analyst to describe surplus split.[29]

Returning to incomplete information, the preceding statements about Nash weights still apply: higher inferred seller bargaining power can be driven by lower seller values or higher

[28] All else equal, a lower $V_{B^{(1)}}$ would lead an empirical analyst to infer more seller bargaining power.

[29] Binmore et al. (1986) showed that patience offers a partial microfoundation: as agents become infinitely patient, the equilibrium of a complete-information alternating-offer game corresponds to the Nash outcome.

prices, and higher prices can arise from unmodeled features such as patience, bargaining costs, or skill. Additionally, in incomplete-information bargaining, differing equilibria across two datasets, and the possibility of equilibrium disagreement, can lead to differences in inferred bargaining power. For example, if two datasets involve the same value distributions and equilibria that are similar in most respects but the first involves more trade for low-value sellers, this could affect inference about seller bargaining power in the two datasets. Bargaining power can also be affected by the negotiation protocol, such as who can make offers and how many they can make; Nash bargaining abstracts away from protocol.[30]

4 Identifying Negotiators' Private Value Distributions

Seller- and buyer-preferred benchmarks require several key inputs: distributions of buyer values F_B and seller values F_S and the real-world mechanism $\mathcal{M}_{RW}$. These objects are the primary goals of the identification and estimation results we now address in a model of the wholesale used-car market game. The game involves 1 seller and N bidders, where N is random and varies across instances of the game.[31] We assume the following:

(A2) *$N \geq 2$ risk-neutral, symmetric bidders participate in an ascending button auction. For $i = 1, ..., N$, each buyer i has a value $Y'\beta + W + V_{B_i}$, with $V_{B_i} \sim F_B$ and $W \sim F_W$ (both with bounded support), and with $\{Y, W, N, \{V_{B_i}\}_{i=1}^N\}$ mutually independent.*

(A3) *A risk-neutral seller has a value $Y'\beta + W + V_S$, with $V_S \sim F_S$ (with bounded support) and with V_S independent of $\{Y, W, N, \{V_{B_i}\}_{i=1}^N\}$.*

(A4) *The auction follows a button format, followed by a finite number of bargaining periods,*

[30] In empirical work, Nash weights and our measures are both inferred based on how observed prices divide the available surplus, but a key distinction is that, under incomplete information, there is no standard definition of the "available surplus": it could be defined as the total surplus achieved in a given equilibrium, the total potential surplus that could be achieved by a second-best mechanism preferred by one party or the other, or the total first-best surplus. Note also that, from a theoretical perspective, in a Nash bargaining model, *power* is a primitive; it is not in our model. However, a researcher could use our measure in *counterfactual* exercises in a similar fashion to how estimated Nash weights are used. For example, in a counterfactual exercise in which the distribution of buyer values is altered, payoffs under benchmark mechanisms $\mathcal{M}_B^{2^{nd}}$ and $\mathcal{M}_S^{2^{nd}}$ would change, and the corresponding surplus of the buyer and seller under the *counterfactual* bargaining can be constructed holding $\alpha_B^{2^{nd}}$ and $\alpha_S^{2^{nd}}$ constant (just as Nash bargaining weights are held constant in counterfactuals in practice; e.g. Crawford and Yurukoglu 2012).

[31] We generally follow the convention that, for random variables, upper case denotes a random variable and lower case a realization or a specific observation in the data.

ending at a final price weakly higher than P^A.

(A5) Any buyer incurs a common cost η_B each time making an offer in bargaining.[32]

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 $\{Y, W, N, V_{B_i}\}$ are mutually independent for all i in a given instance of the game. Thus, agents' overall values are correlated through game-level heterogeneity terms Y and W, but, conditional on these variables, values are independent.[33] To simplify exposition, we condition on $N = n$ and return to this in Section 4.2. Likewise, our discussion largely conditions on $W = w$ and $Y = y$, omitting this dependence until Section 4.3. Note that, When we order bidders' values, we let $V_{B(i)}$ represent the i^{th}-highest value.

The game begins with the seller choosing a secret reserve price, R. Bidders then participate in the auction. If the auction price, P^A, exceeds R, the high bidder (the *buyer*) and seller trade at price P^A. If $P^A < R$, the buyer may choose to opt out of bargaining. If he does not opt out, he incurs a cost η_B, and it is the seller's turn to accept, counter, or quit (ending negotiations) in response to P^A. If the seller counters, it is the buyer's turn. Turns alternate until one party accepts or quits. If the game ends with the seller of type v_S and a buyer type v_{B_i} trading at some price p, the buyer receives a payoff of $v_{B_i} - p$ and the seller receives a payoff of p. If trade fails, the seller receives a payoff of v_S (her value for keeping the car) and the buyer receives a payoff of 0.[34] The auction house enforces that the final price weakly exceeds P^A, the competition constraint (Assumption $A4$).[35]

We refer to V_{B_i} and V_S as *values*, but we could equivalently refer to them as "net values" or "willingness to pay" and "willingness to sell." Suppose bidder i gets $\bar{V}_{B_i}$ when trade occurs and μ_{B_i} when it fails, where μ_{B_i} is a discounted continuation payoff of a broader game in which the bidder can re-enter the market or give up. What we refer to as the bidder's value

[32]It is unnecessary to (and convenient not to) model seller bargaining costs. See Appendix A.

[33]Our model is thus one of IPV with unobserved and observed game-level heterogeneity. The IPV assumption in this market can be motivated by our discussions with industry participants, who tell us that heterogeneity in willingness to pay across bidders arises from differences in dealers' geographic locations, inventory needs, and local demand. $A2$ and $A3$ nest $A1.i$–$A1.ii$.

[34]The buyer's payoff would also include any bargaining costs she has incurred.

[35]This is more of an industry norm than an explicit requirement, but the auction house informs us that violations are rare (and not looked on kindly). Larsen (2021) drops 9 observations violating this norm.

is $V_{B_i} \equiv \bar{V}_{B_i} - \mu_{B_i}$. Similarly, V_S is the seller's discounted continuation payoff in a broader dynamic game.[36] For our research question, we do not need to model this broader game. Our aim is to quantify average bargaining power *given* agents' net private values. See also Freyberger and Larsen (2025).

We introduce some additional notation to define a pure-strategy Bayes Nash Equilibrium (PSBNE) of this game. Let ζ_i be the price at which bidder i drops out of the auction. Let $D_t^B \in \{a, c, q\}$ represent the buyer's decision to accept, counter, or quit in odd periods t. At $t = 1$, let $D_1^B = q$ represent the buyer choosing to opt-out of bargaining when informed that $P^A < R$ (and $D_1^B \neq q$ represents not opting out). Let P_t^B represent the buyer's counteroffer (if the buyer counters) in period t; at $t = 1$, $P_1^B = P^A$. 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, including all offers in bargaining. The strategy of a bidder of type v_{B_i} is a history-contingent set of actions $\sigma_i^B(v_{B_i}) = \{\zeta_i, \{D_t^B|H_t\}_{todd}\}, \{P_t^B|H_t\}_{todd}\}$. The strategy of a seller of type v_S is $\sigma^S(v_S) = \{\rho, \{D_t^S|H_t\}_{t>1,even}, \{P_t^S|H_t\}_{t>1,even}\}$, where $\rho(v_S) = r$ is the seller's reserve price strategy. These constitute a PSBNE 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 reached by some types.[37]

A wide array of behavior can be sustained in equilibrium. Below we offer several examples of behavior, beginning at the $t = 2$, that could be sustained in equilibrium:

Example 1*: At $t = 2$, sellers counter at the optimal TIOLIO Z_S^* facing $V_{B^{(1)}} \geq P^A$. At $t = 3$, buyers accept (if $V_{B^{(1)}} \geq Z_S^*$) or quit. At $t = 4$, (only reached off path) sellers quit.*

Example 2*: At $t = 2$, sellers accept (if $V_S \leq P^A$) or counter at $\bar{b}$ (denoting the upper bound of buyer types, an uninformative offer). At $t = 3$, buyers quit if facing any (off-path) offer*

[36]For example, when trade fails, the seller can attempt to sell the car on a future date (through the same auction house or through a different outlet). In a broader continuation game, F_S and F_B would not be primitives. Also, our exercise of computing benchmarks should not be considered a counterfactual in which the benchmark is adopted in *all future interactions*, as this would alter F_B and F_S in broader continuation game. Instead, the exercise quantifies where surplus lies relative to a *one-time switch* to a benchmark.

[37]We do not impose refinements, such as Perfect Bayes Equilibrium (PBE). Gul and Sonnenschein (1988) and Ausubel et al. (2002) showed these refinements do little or nothing to restrict the set of equilibria in sequential incomplete-information bargaining because, in a BNE, behavior can be sustained by specifying actions that, though only occurring off the equilibrium path, incentivize agents to take certain actions on the path to avoid unappealing off-equilibrium outcomes. Similarly, in a PBE, 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.

other than $\bar{b}$; otherwise, they counter at the optimal TIOLIO Z_B^* facing $V_S > P^A$. At $t = 4$, sellers accept (if $V_S \leq Z_B^*$) or quit. At $t = 5$ (off path), buyers quit.

***Example 3**: At $t = 2$, sellers accept if $V_S \leq P^A$ and quit otherwise. At $t = 3$ (off path), buyers quit.*

***Example 4**: At $t = 2$, sellers either accept P^A, quit, or counter at a price of γP^A (with $\gamma > 1$ fixed and known to all parties), whichever yields a higher expected payoff. At $t = 3$, buyers accept if $V_{B^{(1)}} \geq \gamma P^A$ and quit otherwise. At $t = 3$, buyers quit if facing any (off-path) price other than γP^A. At $t = 4$ (off path), sellers quit.*

In Example 1, seller behavior at $t = 4$ disciplines buyers to only accept or quit (never counter) at $t = 3$. The $t = 5$ behavior of buyers in Example 2 has a similar role. Example 1 is similar to what occurs in $\mathcal{M}_S^{2^{nd}}$, Example 2 is similar to $\mathcal{M}_B^{2^{nd},con}$, and Example 3 is similar (in terms of the seller payoff) to $\mathcal{M}_B^{1^{st},con}$. Example 4 gives the seller a payoff between what she receives in Examples 1 and 3. Example 4 involves strategies conditioning on P^A: both the buyer and seller observe P^A, and any seller counteroffer at $t = 2$ other than γP^A prompts the buyer to quit at $t = 3$.[38] These are only four of infinitely many possibilities.

The following proposition includes several properties used in Larsen (2021) and Larsen et al. (2024), relying on a technical assumption that aids in proving monotonicity of $\rho(\cdot)$:

(A6) The seller's expected payoff conditional on bargaining occurring is continuous in P^A.

Proposition 1. *(i) In any PSBNE, holding fixed any post-auction path of play reached by some types on the equilibrium path and the reserve price strategy of the seller, the following is a weak best response for each bidder: bid truthfully in the auction and enter post-auction bargaining only doing so yields a non-negative expected payoff. Moreover, under A2–A5, in any PSBNE in which bidders follow the strategy in (i) and in which A6 holds, (ii) $\rho(\cdot)$ is strictly increasing and (iii) P^A, R, and final prices are additively separable in (and the trade probability is invariant to) game-level heterogeneity $Y'\beta + W$.*

We assume bidders follow the strategy in (i). The intuition for (i) is that remaining in

[38] Our definition of equilibrium allows agents to condition strategies on P^A but not on *payoff-irrelevant* variables, such as N or losing bids below P^A. For example, our equilibrium definition rules out behavior such as agents coordinating on the Example 1 equilibrium when the third-highest-value buyer has a bid of δ and coordinating on a distinct equilibrium when that bid is $\delta' \neq \delta$. As we show in Proposition 1, the game is location invariant, so W and Y are also not payoff relevant.

the bidding beyond one's value cannot yield a positive payoff given that the final price is constrained to exceed P^A, and dropping out below one's value is unnecessary, even to avoid facing bargaining costs, because these costs are only incurred if the buyer enters bargaining ($A5$). Buyers opting out also drives (ii): too high of a reserve price is unappealing to the seller as it triggers the buyer opting out more often. Part (iii) is a common result in empirical auction models (e.g., Haile et al. 2003).

PSBNE behavior is sufficient but not necessary for Proposition 1 to hold.[39] For example, suppose agents' actions align with $\mathcal{M}_B^{1^{st},con}$: sellers choose $\rho(V_S) = V_S$ and, whenever $P^A < V_S$, counter at a price of V_S. Such behavior, while not a BNE, would still satisfy properties (i)–(iii) in Proposition 1.[40] We will use the term *equilibrium* throughout the paper to denote a PSBNE satisfying $A6$. At some points, for clarity of exposition, we will use the term *behavioral equilibrium* to denote a strategy set involving some agents taking actions that appear to be suboptimal at some nodes of the game reached with positive probability by some types (which should not occur in a BNE).

As highlighted in Section 3, the key objects to identify to evaluate bargaining power are F_S, F_B, and $\mathcal{M}_{RW}$. With these objects, we can compute the expected payoff for each player under benchmark mechanisms 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, $\mathcal{M}_{RW}$, and game-level heterogeneity rely largely on Proposition 1 or prior results from the literature. Sections 4.2–4.4 discuss these. In summary, our identification results in Sections 4.1–4.4 rely on joint variation among four key variables: an indicator for whether trade occurs, and noisy measures (i.e., those including W) of auction, secret reserve, and final prices.

4.1 Equilibrium Menus. Rather than working with the full, complex set of equilibrium strategies, we follow the mechanism design literature and analyze the game as a direct mechanism. In a direct mechanism $\mathcal{M}$, each agent, $i \in \{S, B_1, ..., B_n\}$, reports (or potentially misreports) her private value to a designer, who assigns allocations according to a function

[39] Mixed strategy equilibria also exist, but limiting to pure strategies is sufficient and useful. Strictly speaking, our identification results hold under some degree of mixing: a seller of type v_S can mix over reserve prices r in some set $\mathcal{R}(v_S)$ as long as $\mathcal{R}(v_S)$ and $\mathcal{R}(v'_S)$ are disjoint for any $v_S \neq v'_S$. Larsen and Zhang (2018) derive results for the case where mixing can lead to different v_S choosing the same r.

[40] This behavior is not a BNE because offering V_S is not incentive compatible for sellers.

$x_i\left(v_S \ldots v_{B_n}\right)$. For a mechanism corresponding to a pure strategy equilibrium, the allocation is equal to 1 if the agent is allocated the car and zero otherwise. The designer allocates net payments made by i according to the transfer function $t_i\left(v_S \ldots v_{B_n}\right)$.

An agent's expected outcome is described by a *menu* of probability-transfer pairs. If player i behaves as if she is a type v_i' (potentially misreporting her type), she attains

$$P_i\left(v_i'\right) \equiv E\left[x_i\left(v_i', V_{-i}\right)\right], \; T_i\left(v_i'\right) \equiv E\left[t_i\left(v_i', V_{-i}\right)\right].$$

$P_i\left(v_i'\right)$ and $T_i\left(v_i'\right)$ are, respectively, the expectation of i's allocation $x_i\left(v_i', V_{-i}\right)$ and transfer $t_i\left(v_i', V_{-i}\right)$, 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 $v_i P_i\left(v_i'\right) - T_i\left(v_i'\right)$. In any incentive compatible (IC) mechanism, v_i maximizes this payoff:

$$v_i P_i\left(v_i\right) - T_i\left(v_i\right) \geq v_i P_i\left(v_i'\right) - T_i\left(v_i'\right) \forall v_i'. \tag{2}$$

This implies bounds on agent i's value, the proof of which follows immediately from (2):

Theorem 1. *Under A2–A3, for any agent i, v_i must satisfy*

$$v_i \geq \frac{T_i\left(v_i\right) - T_i\left(v_i'\right)}{P_i\left(v_i\right) - P_i\left(v_i'\right)} \; \forall v_i' : \; P_i\left(v_i'\right) < P_i\left(v_i\right)$$

$$v_i \leq \frac{T_i\left(v_i'\right) - T_i\left(v_i\right)}{P_i\left(v_i'\right) - P_i\left(v_i\right)} \; \forall v_i' : \; P_i\left(v_i'\right) > P_i\left(v_i\right).$$

In any setting 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, the seller's secret reserve price is such an action. We combine Theorem 1 with this property to obtain the following corollary specific to the seller's value; as such, we state it only for $i = S$:

Corollary 1. *Under A2–A6, for a seller v_S picking r, v_S satisfies (i) $v_S \geq \frac{T_S(r) - T_S(r')}{P_S(r) - P_S(r')}$ for all r' with $P_S\left(r'\right) < P_S\left(r\right)$ and (ii) $v_S \leq \frac{T_S(r') - T_S(r)}{P_S(r') - P_S(r)}$ for all r' with $P_S\left(r'\right) > P_S\left(r\right)$.*

Corollary 1 modifies notation slightly, with 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 $v_S = \rho^{-1}(r)$. We maintain this change moving forward. We also write the seller's allocation and transfer functions as $x_S\left(r, V_{-S}\right)$ and $t_S\left(r, V_{-S}\right)$.

Figure 2: Hypothetical Equilibrium Menu

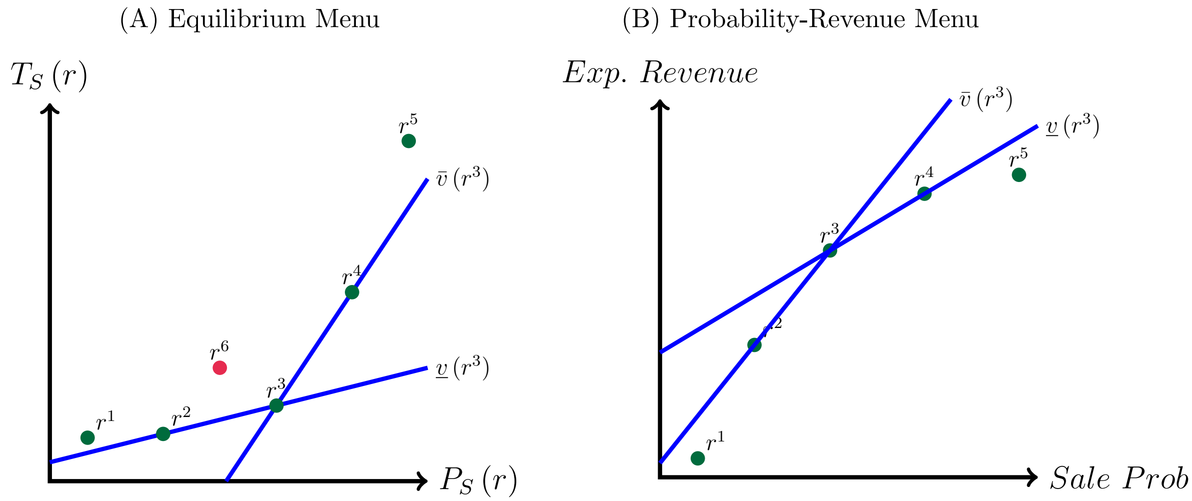

Notes: Panel A shows a hypothetical equilibrium menu. The slopes of the blue lines are upper and lower bounds for the value of an agent choosing action r^3. Panel B shows the menu in terms of the probability of sale, $1 - P_S(r)$, and expected revenue, $-T_S(r)$. Again, the slopes of the blue lines bound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

Figure 2.A illustrates a hypothetical equilibrium menu facing the seller. We consider a case where the seller's possible choices of secret reserve prices are $r' \in \{r^1, ..., r^5\}$. 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.[41] 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 probability of keeping the good. If the seller prefers 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, Corollary 1 implies 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.A.

The equilibrium menu can alternatively be thought of in terms of the sale probability $(1 - P_S(r))$ and expected revenue $(-T_S(r))$, shown in Figure 2.B. Sellers choose a probability-revenue pair from this menu: when sellers set lower r, they sell more often and get higher

[41] Here 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.

expected revenues, but *marginal* revenue from selling with higher probability is lower with lower r. If a seller chooses r^3, the secant lines at r^3 bound her value for the car.

This leads to our main (and most powerful) result: for a sufficiently smooth and continuous game, the Theorem 1 bounds yield point identification. $A7$ states these conditions:[42]

(A7) (i) $P_S(r)$ is continuous and strictly monotone with derivative bounded away from 0; (ii) $P_S(r), T_S(r)$ are both continuously differentiable; and (iii) $\rho(v_S)$ is continuous and strictly monotone in v_S.

Corollary 2. *Under A3 and A7, $v_S(r)$, the inverse of $\rho(v_S)$, satisfies $v_S(r) = \frac{T'_S(r)}{P'_S(r)}$ for all reserve prices r played by some type v_S in equilibrium.*

Corollary 2 is the smooth, continuous-action analog of what is illustrated in Figure 2. The required smoothness conditions are that different types of sellers choose different reserve prices, different reserve prices lead to different probabilities of the seller keeping the car, 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.[43] This can be considered a generalization of GPV.[44] The GPV argument is specific to a first-price auction (FPA). In contrast, our approach applies to any game where the econometrician observes whether trade is successful, transfers, and some variable that is one-to-one with agents' types. This feature may be advantageous in other bargaining settings as well, where equilibria may be difficult to characterize, and game rules differ from well-studied auction formats.

[42] $A7$ is assumed directly on equilibrium objects (the menu and $\rho(v_S)$) rather than on primitives.

[43] Corollary 2 applies only to strictly separating equilibria in the sense that no two distinct v_S choose the same r, and each r leads to a different $P_S(r)$. In some bargaining games, equilibria feature partial-pooling regions where, for example, all sellers with 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.

[44] To see this, consider an n-bidder FPA in a symmetric IPV environment. Let the distribution of bids be $G(\cdot)$, with density $g(\cdot)$. In an FPA, the expected probability of winning, (P_i), and expected transfer (T_i) for bidder i with bid b_i are known transformations of G, given by $P_i(b_i) = G(b_i)^{n-1}$ and $T_i(b_i) = b_i G(b_i)^{n-1}$. Differentiating 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 term equals GPV's, but is derived without exploiting any knowledge that the game is a FPA.

Corollary 2 only requires optimality (or equivalently, IC) of the seller's secret reserve choice; it does not require optimality of buyers' actions or of the seller's own actions at later stages of the game (e.g., in bargaining). Thus, maintaining the assumption of optimality of seller's secret reserve choices, Corollary 2 provides identification of seller values even if other parts of agents' strategies belong to a behavioral equilibrium. Optimality of the seller's secret reserve choice implies two additional restrictions on the equilibrium menus. First, equilibrium menus must satisfy *individual rationality* (IR), implying that the seller's expected payoff under the optimal r must exceed v_S, her value of keeping the car herself:

$$\max_{r'} v_S P_S(r') - T_S(r') \geq v_S \quad \forall v_S \tag{3}$$

Second, equilibrium menus must be convex:

Theorem 2. *Under A3 and A7, 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. Points interior to the menu's convex hull are dominated: no type would find it optimal to play such actions. We impose convexity and IR on our estimated menu.

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 model. Data on failed attempts to trade that are essential for rejecting a complete-information environment.

4.2 Identification of F_B. Identification of F_B is relatively standard. By Proposition 1.i, the auction price equals the second-highest bidder's value. 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 $F_{P^A}(\cdot)$. For any y, the following holds:

$$F_{P^A}(y) = \sum_n \Pr(N = n)\left[nF_B(y)^{n-1} - (n-1)F_B(y)^n\right]. \tag{4}$$

The right-hand side of (4) is monotonic in y, and thus F_B is identified.[45] The distribution

[45]This approach requires N being independent of bidders' values ($A2$). Appendix B.4 shows evidence

of the highest-value bidder (the maximum order statistic, which we denote $V_{B(1)}$) is then

$$F_{B^{(1)}}(y) = \sum_n \Pr(N = n) F_B(y)^n \tag{5}$$

Appendix B.3 discusses how we specify $\Pr(N = n)$ and shows that the estimate of $F_{B^{(1)}}$ is not sensitive to this choice.

4.3 Game-Level Heterogeneity. We now incorporate game-level heterogeneity $Y\beta + W$, where the econometrician observes Y but not W. Proposition 1.iii implies *location invariance*: auction, reserve, and final prices shift additively in $Y'\beta + W$, and trade probability is invariant. This applies to W and $Y'\beta$, but the two forms of heterogeneity require different identification arguments. Let $R^{raw} \equiv Y'\beta + \tilde{R}$ and $P^{A,raw} \equiv Y'\beta + \tilde{P}^A$ represent reserve and auction prices with both heterogeneity terms included, where $\tilde{R} \equiv R + W$ and $\tilde{P}^A \equiv P^A + W$ (with joint density $f_{\tilde{R},\tilde{P}^A}$). Proposition 1 implies mutual independence of $\{Y, W, R, P^A\}$, and thus an estimate of β can be obtained from a linear regression of observations of R^{raw} and $P^{A,raw}$ on Y, with residuals corresponding to estimates of $\tilde{R}$ and $\tilde{P}^A$.

Evdokimov and White (2012), Lemma 1, yields the following.[46] For $\omega \in \{R, P^A, W\}$, let f_ω be a density and ϕ_ω a characteristic function (with derivative ϕ'_ω).

(A8) (i) ϕ_R and ϕ_W have only isolated real zeros; (ii) the real zeros of ϕ_{P^A} and ϕ'_{P^A} are disjoint; (iii) E[W]=0.

Proposition 2. *Under A2–A8, f_W, f_{P^A}, and f_R are identified from $f_{\tilde{R},\tilde{P}^A}$.*

The intuition behind Proposition 2 is that, due to independence, correlation between $\tilde{R}$ and $\tilde{P}^A$ is driven by W. Means of R, P^A, and W are not identified without a location normalization; we impose $E[W] = 0$. $f_{\tilde{P}^A,\tilde{R}}$ can be written as a convolution of marginals:

$$f_{\tilde{R},\tilde{P}^A}(\tilde{r}, \tilde{p}^A) = \int f_{P^A}(\tilde{p}^A - w) f_R(\tilde{r} - w) f_W(w) dw \tag{6}$$

Proposition 2 yields identification of F_{P^A}, and then we apply (4) and (5) to identify F_B and $F_{B^{(1)}}$. We now describe identification of the (P_S, T_S) menu under game-level heterogeneity. Consider a realization of game-level heterogeneity $Y\beta + W = \tau$. We re-

consistent with this assumption. Appendix B.4 also analyzes the correlation of N with other variables.

[46] By Evdokimov and White (2012), standard distributions satisfy $A8$, which is weaker than the conditions of Krasnokutskaya (2011), who first applied a similar result to unobserved game-level heterogeneity.

quire that agents' strategies constitute an equilibrium conditional on any such realization. Define the seller's expected allocation and transfer when playing reserve price $r + \tau$ in equilibrium as $P_S^\tau(r+\tau) \equiv E[x_S(r+\tau, V_{-S}+\tau) \mid Y'\beta + W = \tau]$ and $T_S^\tau(r+\tau) \equiv E[t_S(r+\tau, V_{-S}+\tau) \mid Y'\beta + W = \tau]$, with expectations taken over other agents' values (and hence their equilibrium actions). By Proposition 1.iii, adding τ to values shifts prices by τ:

$$price^\tau(r+\tau) = price^0(r) + \tau \;\forall \tau, r \tag{7}$$

where $price^\tau(r+\tau)$ is the expected final price conditional on trade when game-level heterogeneity equals τ. The expected transfer can be written $T_S^\tau(r+\tau) = price^\tau(r)\left(P_S^\tau(r+\tau) - 1\right)$, which, combined with (7) and with $P_S^0(r) = P_S^\tau(r+\tau)$ (another implication of Proposition 1.iii), becomes

$$T_S^\tau(r+\tau) = \left(price^0(r) + \tau\right)\left(P_S^0(r) - 1\right) = T_S^0(r) + \tau\left(P_S^0(r) - 1\right) \;\forall \tau, r \tag{8}$$

Thus, relative to a case where game-level heterogeneity is zero, a case where $\tau \neq 0$ shifts the argument inside the probability function by τ, and shifts and rotates the expected transfer function: (8) implies that $\frac{dT_S^\tau}{dP_S^\tau}\Big|_{r+\tau} = \frac{dT_S^0}{dP_S^0}\Big|_r + \tau$. This also implies that equilibrium menus are fully characterized by probabilities and transfers at $\tau = 0$, $P_S^0(r), T_S^0(r)$.

(P_S^0, T_S^0) 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.

Theorem 3. *Under A2–A8, $P_S^0(\cdot), T_S^0(\cdot)$ are identified (on the support of R that occur in equilibrium for at least some realizations of V_S) from $\tilde{P}_S(\tilde{r}), \tilde{T}_S(\tilde{r})$, f_R, and f_W.*

The proof of Theorem 3 demonstrates that $P_S^0(\cdot)$ and $T_S^0(r)$, the underlying expected allocation and transfer functions purged of unobserved heterogeneity, solve

$$\tilde{P}_S(\tilde{r}) = \frac{\int P_S^0(r) f_R(r) f_W(\tilde{r} - r)\, dr}{\int f_R(r) f_W(\tilde{r} - r)\, dr} \tag{9}$$

$$\tilde{T}_S(\tilde{r}) - E[W \Delta P_S \mid \tilde{r}] = \frac{\int T_S^0(r) f_R(r) f_W(\tilde{r} - r)\, dr}{\int f_R(r) f_W(\tilde{r} - r)\, dr} \tag{10}$$

$$E[W \Delta P_S \mid \tilde{r}] \equiv \frac{\int (\tilde{r} - r)\left(P_S^0(r) - 1\right) f_R(r) f_W(\tilde{r} - r)\, dr}{\int f_R(r) f_W(\tilde{r} - r)\, dr} \tag{11}$$

We describe in Section 5 how we exploit (9)–(11) to estimate $P^0_S(\cdot)$ and $T^0_S(\cdot)$.

4.4 Allocation Function for $\mathcal{M}_{RW}$. We assume the following:

(A9) Conditional on $P^A = p^A$ and $V_S = v_S$, trade occurs for a buyer $V_{B^{(1)}} = v_{B^{(1)}}$ if $v_{B^{(1)}}$ is greater than a function $g_S(v_S, p^A)$.

We do not believe $A9$ is overly strong for our setting: $A9$ holds in many bargaining models (e.g., Myerson and Satterthwaite 1983, Williams 1987), and Theorem 1 of Storms (2015) implies that it holds in any PSBNE of our game when $V_{B^{(1)}}$ and V_S have finite type spaces, including arbitrarily fine type spaces.[47] Behavior in Examples 1–4 does not violate $A9$, and non-BNE behavior can also satisfy $A9$ (e.g., $\mathcal{M}_B^{2^{nd},con}$).

Under $A9$, the allocation function for $\mathcal{M}_{RW}$ is $1\{v_B^{(1)} \geq g_S(v_S, p^A)\}$. Because $\rho(v_S)$ is strictly increasing, we can re-write this as a function of r rather than v_S. Let $g_R(\rho(v_S), p^A) \equiv g_S(v_S, p^A)$. The function g_R is related to the trade probability at a realization of $\tilde{P}^A$ and $\tilde{R}$:

$$\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A) = \int \frac{F_{B^{(1)}|P^A}\left(g_R(\tilde{r} - w, \tilde{p}^A - w) \mid \tilde{p}^A - w\right) I_g(\tilde{r}, \tilde{p}^A, w)}{\int I_g(\tilde{r}, \tilde{p}^A, z)dz} dw \qquad (12)$$

where $I_g(\tilde{r}, \tilde{p}^A, w) \equiv f_R(\tilde{r} - w) f_{P^A}(\tilde{p}^A - w) f_W(w)$ is the joint density of (R, P^A, W), and $\mathcal{A}$ is the event that trade occurs. $F_{B^{(1)}|P^A}\left(v_{B^{(1)}} \mid p^A\right) \equiv \frac{1 - F_B(v_{B^{(1)}})}{1 - F_B(p^A)}$ is the CDF of $V_{B^{(1)}}$ conditional on $V_{B^{(1)}} \geq P^A$. F_B and the densities in (12) are identified, and thus $g_R(\cdot)$ is as well, by a similar convolution theorem argument as in the proof of Theorem 3.

5 Estimation of Menus and Private Values

5.1 Estimation Details. Estimation follows identification closely. Let j denote an observation, consisting of the allocation (an indicator for the seller keeping the car), the transfer (the final price if the car sells and zero otherwise), the secret reserve price, and the auction price. Let y_j be a vector of controls. Estimation requires two additional assumptions:

(A10) Random variables $\{V_S, \{V_{B_i}\}_{i=1}^N, W, N\}$ are independently and identically distributed across instances of the game.

[47] Ausubel and Deneckere (1993) referred to $A9$ as the "northwestern criterion," as it implies a function of seller types such that, in a plot with sellers' types on the horizontal axis and buyers' on the vertical, all buyer and seller types northwest of the function will trade. In our model, the property is most straightforward to prove for finite type spaces, as this restriction permits conditioning on elements of finite sets rather than positive probability subsets of infinite sets.

(A11) Observations in the data arise from the same equilibrium.

$A10$ and $A11$ are common in structural work and allow us to apply the model's properties to the data. They can be relaxed somewhat by performing estimation separately in subsamples of the data, which we do below. We describe each estimation step in turn.

Step 1) Observable Heterogeneity. Let $r_j^{raw} \equiv y_j'\beta + \tilde{r}_j$ and $p_j^{A,raw} \equiv y_j'\beta + \tilde{p}_j^A$ be the reserve and auction prices for observation j before removing game-level heterogeneity, where $\tilde{r}_j \equiv r_j + w_j$, $\tilde{p}_j^A \equiv p_j^A + w_j$. We estimate β with the following stacked regression:

$$\begin{bmatrix} r_j^{raw} \\ p_j^{A,raw} \end{bmatrix} = \begin{bmatrix} y_j'\beta \\ y_j'\beta \end{bmatrix} + \begin{bmatrix} \tilde{r}_j \\ \tilde{p}_j^A \end{bmatrix}, \tag{13}$$

This is the bid *homogenization* approach of Haile et al. (2003). 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.[48] To aid in estimating this large number of categorical dummies, we use an augmented dataset that includes our main sample plus 39,700 observations in which cars that failed to sell were later re-auctioned and 80,213 observations for which we observe only a reserve or auction price but not both. The adjusted R^2 is 0.95, suggesting that most of the variation in these prices arises from observable differences across cars.[49] We refer to the predicted value $y_j'\hat{\beta}$ as the *market value* of the car. The average estimated market value is \$10,255. Let $\widetilde{price}_j$ be given by $price_j^{raw} - y_j'\hat{\beta}$ when trade occurs and 0 otherwise.[50]

Step 2) Unobserved Heterogeneity. We estimate f_W, f_R, and f_{P^A} and their corresponding CDFs via maximum likelihood, where (6), evaluated at the Step 1 residuals, is the

[48] Other controls are fifth-order polynomials in the auction house's book value 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 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).

[49] Estimating (13) with outcomes in logs instead of levels (consistent with a model of multiplicative rather than additive separability) yields a lower adjusted R^2 (0.88) and more residual variance across the range of predicted values of the regression (i.e., greater heteroskedasticity).

[50] Appendix B.1 offers more discussion of the augmented dataset and Appendix B.2 illustrates what variation in these residuals aids in the menu identification approach.

contribution of j to the likelihood. We model each $\omega \in \{P^A, R, W\}$ as $N(\mu_\omega, \sigma_\omega)$.[51]

Step 3) Buyer Value CDF. We estimate F_B by solving (4) on a grid. This requires an estimate of F_{P^A}, which comes from Step 2, and an estimate of $\Pr(N = n)$, which we obtain from the data subsample with bid logs, from which an auction-by-auction lower bound on the number of bidders, $\underline{N}$, can be imputed.[52] We replace N in (4) and (5) with $\underline{N}$. We construct the value distribution for the winning bidder, $F_{B^{(1)}}$, using (5).

Step 4) $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$. We estimate the noisy expected allocation function $\tilde{P}_S(\tilde{r})$ through a local linear regression of an indicator for the seller keeping the car on $\tilde{r}_j$. We estimate the noisy expected transfer $\tilde{T}_S(\tilde{r})$ analogously through a local linear regression of $-\widetilde{price}_j$ on $\tilde{r}_j$. These regressions use a Gaussian kernel and \$500 bandwidth. For comparison, the mean and standard deviation of reserve prices in the data are \$10,368 and \$5,929.

Step 5) Estimating $P_S^0(r)$ via Splines. We parameterize the function $P_S^0(r)$ as a quadratic I-spline (Ramsay 1988) with 5 knots, constrained to be nondecreasing in r.[53] Denote this $P_S^0(r;\theta_P)$. We estimate the spline coefficients θ_P as the solution to

$$\min_{\theta_P} \int \left[\left(\frac{\int P_S^0(r;\theta_P)\, \hat{f}_R(r)\, \hat{f}_W(\tilde{r}-r)\, dr}{\hat{h}(\tilde{r})} \right) - \hat{\tilde{P}}_S(\tilde{r}) \right]^2 \hat{h}(\tilde{r}) d\tilde{r} \tag{14}$$

In words, (14) numerically solves (9) by minimizing the distance between the estimated function $\hat{\tilde{P}}_S(\tilde{r})$ 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 $\hat{h}(\tilde{r}) \equiv \int \hat{f}_R(r)\, \hat{f}_W(\tilde{r}-r)\, dr$, the estimated density function of $\tilde{r}$.[54] For brevity, we sometimes denote $P_S^0\left(r;\hat{\theta}_P\right)$ by $\hat{P}_S^0(r)$. The square root of the value of (14) at the optimum constitutes

[51] A likelihood ratio test comparing this model to one with fifth-degree Hermite polynomials for densities (as in Gallant and Nychka 1987) fails to reject the more parsimonious model.

[52] This lower bound is the number of unique bidders who placed a bid online (bidder identities are observed for online bidders) plus 1 if 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 attributes all physical bids to a single bidder unless there are two such bids in a row, motivated by the intuition that no bidder bids against herself and so two consecutive physical bids correspond to an auction with at least two physical bidders. Appendix B.3 discusses alternative choices for $\Pr(N = n)$.

[53] These knots are placed at evenly spaced quantiles of the reserve price distribution. The choice of 5 as the number of knots for P_S^0 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. See Appendix B.5 for additional discussion of estimates from Steps 5–6. Appendix B.6 shows how our estimates differ if we ignore unobserved heterogeneity.

[54] Weighting 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 $\hat{\tilde{P}}_S(\tilde{r})$ may be less accurately estimated.

one measure of fit — a root weighted mean squared error (RMSE). Because we estimate a probability in this step, the RMSE lies in [0,1]. We find a RMSE of 0.0021, suggesting that the convolution differs from the local linear fit by 0.21 percentage points on average.

Step 6) Estimating $T_S^0(r)$ via Splines. Using the estimated probability $P_S^0\left(r;\hat{\theta}_P\right)$ 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 $\breve{T}_S^0\left(P_S^0\left(r;\hat{\theta}_P\right);\theta_T\right)$.[55] This spline approximation allows us to directly impose convexity of the transfer-probability menu. We estimate

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

subject to the constraint that $\breve{T}_S^0(p)$ is weakly convex. This requires an estimate of $E\left[W\Delta P_S\mid\tilde{r}\right]$, which we construct using (11). We denote the value of $\breve{T}_S^0\left(P_S^0\left(r;\hat{\theta}_P\right);\hat{\theta}_T\right)$ 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 (\$4.08) relative to prices in this market, indicating a good fit.

Step 7) IR constraint, $\rho(\cdot)$, and F_S. We construct, for a grid of values for r, an estimate of the type v_S that would choose each r, equal to $\frac{d\breve{T}_S^0}{dp}$, the derivative of the outer spline function $\breve{T}_S^0(p;\theta_T)$ with respect to p, evaluated at the point $\hat{P}_S^0(r)$. This derivative has a closed form given our parametrization of $\breve{T}_S^0(p;\theta_T)$. We then impose IR as follows: For any implied v_S choosing r such that (3) does not hold, we replace v_S with the marginal seller type receiving zero utility and, holding $\hat{P}_S^0(r)$ fixed, set $\hat{T}_S^0(r)$ to the value making (3) hold with equality.[56] This step affects only 2.44% of sellers. From the v_S corresponding to each r we obtain an estimate of the mapping $\rho(v_S)$. F_S is then given by $F_S(v_S)=F_R(\rho(v_S))$ for any v_S. Simulating independent draws from our estimates of $Y'\beta$, F_W, and F_S, we find that 97.37% of sellers are estimated to have total values $(Y'\beta+W+V_S)$ that are positive.

Step 8) Trade Cutoff Function, $g_R(\cdot)$. We minimize the distance between the left- and right-hand sides of (12), plugging in F_B from Step 3 and f_R, f_{P^A}, and f_W from Step 2, as

[55] For $\breve{T}_S^0$, we use 4 knots uniformly spaced from 0 to 5% above the highest predicted value of $P_S^0\left(r;\hat{\theta}_P\right)$ on our grid of r-values. We choose 4 knots for this step to ensure some undersmoothing in the earlier step.

[56] This treats these sellers as though trading with probability $\hat{P}_S^0(r)$ but having zero expected surplus, making them exactly indifferent to participating.

well as an estimate of $\Pr(\mathcal{A}|\tilde{R}=\tilde{r},\tilde{P}^A=\tilde{p}_A)$.[57] We parameterize $h_g(r,p^A)\equiv\frac{1-F_B(g(r,p^A))}{1-F_B(p^A)}$ using a flexible bilinear spline and then invert the estimate of $h_g(\cdot)$ to obtain $\hat{g}_R(r,p^A)=\hat{F}_B^{-1}(1-(1-\hat{F}_B(p^A))\hat{h}_g(r,p^A))$.[58]

5.2 Estimated Menu and CDFs. Figure 3.A shows the menu, with the estimate of $P_S^0(r)$ on the horizontal axis and $T_S^0(r)$ on the vertical. Each point on the menu corresponds to the seller's expected payoff from choosing a given secret reserve price. The units for the vertical axis are \$1,000. We make several remarks before discussing the estimates. First, our formulation for payoffs, $v_SP_S^0(r)-T_S^0(r)$, means $P_S^0(r)$ is the probability the seller *keeps* the good, so $1-P_S^0(r)$ is the sale probability, and $-T_S^0(r)$ is the expected payment received by the seller. Second, recall that seller (and buyer) values are additive in game-level heterogeneity, with a seller's total value for the car given by $V_S+Y'\beta+W$. Therefore, for a car with $Y'\beta+W=$ \$10,000, $V_S=$ -\$1,000 (a negative value) would imply the seller's total value for the car is \$9,000, and $V_S=$ \$1,000 would mean her total value is \$11,000. $V_S<0$ are types whose optimal action lies on the *downward-sloping* portion of the menu in Figure 3.A.

With these remarks in mind, we compare several points along the menu, $r_A,...,r_D$, 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 car 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-sloping portion of the menu. Choice r_C yields a lower probability

[57] We estimate this conditional probability by regressing an indicator for trade occurring on a tensor product of cubic b-spline functions with fifteen uniformly spaced knots in each dimension.

[58] Parameterizing and estimating $h_g(r,p^A)$ rather than $g_R(\cdot)$ is useful because the former is bounded on $[0,1]$. Our parameterization of $h_g(\cdot)$ uses 25 knots in each dimension, uniformly spaced between the 0.001 and 0.999 quantiles of $\tilde{R}$ and $\tilde{P}^A$. We impose several constraints in estimation: (i) $h_g(r,p^A)\in[0,1]$; (ii) $h_g(r,p^A)$ decreasing in r; (iii) $h_g(r,p^A)=1$ if $r\le p^A$; (iv) $g_R(r,p^A)\ge\underline{g}(r,p^A)\equiv\max\{p^A,\rho^{-1}(r)\}\Rightarrow h_g\le\frac{1-F_B(\underline{g}(r,p^A))}{1-F_B(p^A)}$; (v) $E_{R,P^A}[h_g(r,p^A)\underline{g}(r,p^A)]\le E[\widetilde{price}|\mathcal{A}]\Pr(\mathcal{A})$. Condition (iv) ensures that the estimated g does not allow trade when $v_S>v_B$. Condition (v) follows from the fact that the observed average price must exceed the lowest possible price, which is $\max\{p^A,\rho^{-1}(r)\}$.

Figure 3: Estimated Menu

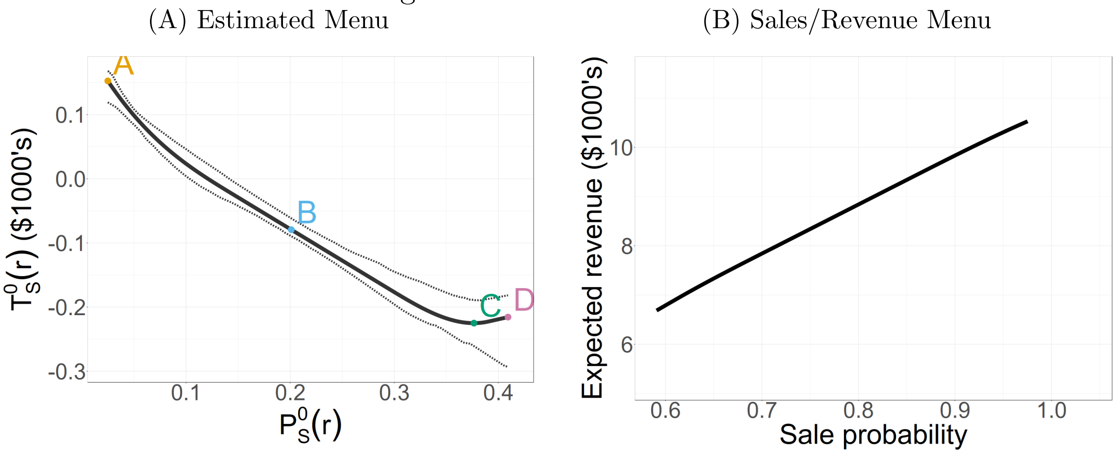

Notes: Panel A displays the final estimated menu, $(\hat{P}^0_S, \hat{T}^0_S)$. Dashed lines show pointwise 95% confidence bands a nonparametric percentile bootstrap with 200 replications. Units on the vertical axis are $1,000 relative to the market value estimate. The points marked A–D are discussed in the body of the paper. Panel B shows an alternative version of the menu, as in Figure 2, with the mean game-level heterogeneity term added to all values and prices: $\tau = E[Y'\hat{\beta} + W] = E[Y'\hat{\beta}] =$ \$10,255. Panel B thus shows sale probabilities, $1 - P^\tau_S(r + \tau)$, and expected revenues, $(-T^\tau_S(r + \tau))$, computed as described in (7)–(8).

of keeping the good but a higher expected transfer to the seller (i.e. a more negative T^0_S). Therefore, a seller choosing r_D wants to keep the good more (i.e. has a higher v_S) than a seller choosing r_C, manifest by a derivative that is more positive at D than at C. Figure 3.B offers an alternative version of this menu in terms of *expected revenues* and *sale probabilities*, analogous to Figure 2.B. To improve readability, we adjust expected revenues by adding an amount $\tau = E[Y'\hat{\beta} + W] = E[Y'\hat{\beta}] =$ \$10,255 to all agents' values and, consequently (by Proposition 1.iii), to all prices, shifting upward and rotating the menu as in (7)–(8).

From the menu derivatives, we obtain $v_S = \rho^{-1}(r)$. 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 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). For a seller who chooses a reserve price equal to the heterogeneity value of the car $(Y\hat{\beta} + W)$, represented by 0 on horizontal axis, Figure 4.A implies that the seller's value v_S was substantially below the car's market value. Consider a car with market value equal to the mean, $\tau =$ \$10,255. A seller who chooses a reserve price of \$10,000 for this car has a value of \$8,533.2; one who chooses \$11,000 has a value of \$9,765.7; and one who chooses \$12,000 has a value of \$9,966.9. Taking into account the distribution of reserve prices, the average difference between reserve prices and seller values is \$1,569.4.

Figure 4: Value Mapping and Value CDFs

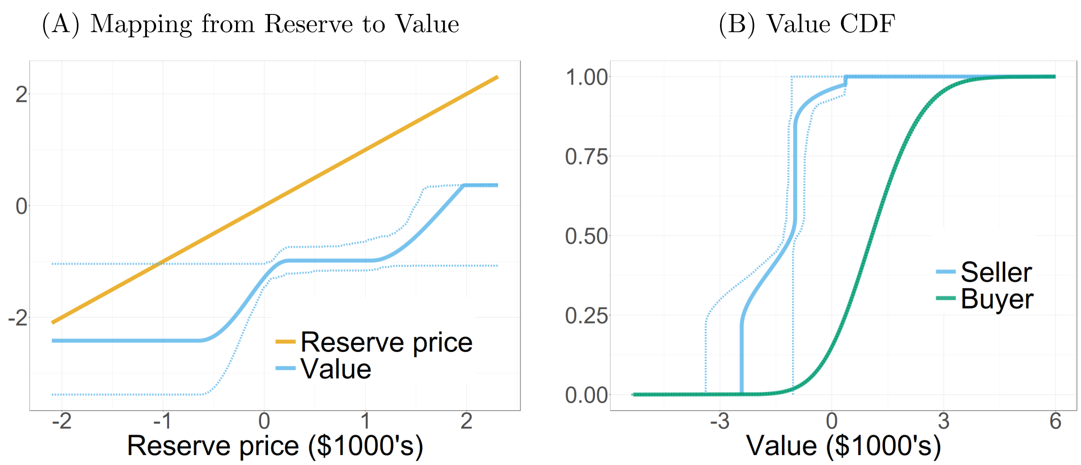

Notes: Left panel shows the estimated mapping from reserve prices to values (blue line) and the reserve price (yellow line, $y = x$). Right panel displays estimated CDF of seller values, F_S (blue line), and estimated CDF of maximum order statistic of buyer values, $F_{B(1)}$ (green line). Dashed lines 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.

The estimated $\rho^{-1}(r)$ mapping, combined with F_R, gives us an estimate of F_S, which we plot in Figure 4.B.[59] We also plot the estimate of $F_{B(1)}$, which lies to the right, suggesting gains from trade exist 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. Figure 4.B also reveals overlap between F_S and $F_{B(1)}$, which, by the Myerson-Satterthwaite Theorem, can lead to inefficiency, with some trades failing even when $V_{B(1)} > V_S$.

6 Quantifying Bargaining Power

Computing α_S^m and α_B^m using (1) requires computing buyer and seller surplus under a given benchmark mechanism $m \in \mathbb{M}$ and under the real-world mechanism, $\mathcal{M}_{RW}$.[60] In this section

[59] The convexity constraint imposed on the menu binds over a region of secret reserve prices, implying an atom in F_S where 28.90% of sellers have the same value.

[60] Game-level heterogeneity is irrelevant for calculating surplus because buyer and seller surplus are invariant to it. Specifically, by Proposition 1.iii, a shift in game-level heterogeneity $Y'\beta + W = \tau$ has no effect on the trade probability and shifts buyer values, seller values, and prices by the same amount, τ, leaving surplus unchanged. For example, consider two cases: case (i) τ = $10,000, the seller's total value (i.e., including τ) is $9,000, the buyer's is $12,000, and the price is $11,000; and case (ii) τ = $5,000, the seller's total value is $4,000, the buyer's is $7,000, and the price is $6,000. In both cases, the realized surplus is $2,000 for sellers and $1,000 for buyers; the level of τ is irrelevant as it appears in seller values, buyer values, and prices, and thus it differences out. This can also be seen mathematically. By Proposition 1.iii, if a seller of value v_S would have chosen a reserve price of r when $Y'\beta + W = 0$, then, when $Y'\beta + W$ is equal to some $\tau \neq 0$, her value and reserve price shift additively by τ. Then, by (7)–(8), her surplus is $(v_S + \tau)P_S^\tau(r + \tau) + price^\tau(r + \tau)(1 - P_S^\tau(r + \tau)) = v_S P_S^0(r) - T_S^0(r)$, where τ falls out of the expression.

we first describe how we compute these expected surplus measures using the objects estimated in Section 5. We report and discuss bargaining power estimates relative to first-best benchmarks, then second-best benchmarks, then separately for OEM vs. non-OEM sellers.

6.1 Computing Expected Surplus and Bargaining Power. Total expected surplus under any mechanism can be described by an allocation and transfer. For mechanism $\mathcal{M}$ and any realizations of $(V_S, V_{B^{(1)}}, P^A)$, let $x^{\mathcal{M}}(v_S, v_{B^{(1)}}, p^A)$ be a dummy denoting whether trade occurs. Total expected surplus under $\mathcal{M}$ is

$$\int \left[\int \int \left(v_{B^{(1)}} - v_S\right) x^{\mathcal{M}}(v_S, v_{B^{(1)}}, p^A) dF_S\left(v_S\right) dF_{B^{(1)}|P^A}\left(v_B \mid p^A\right) \right] dF_{p^A}\left(p^A\right) \tag{15}$$

$U_B(\mathcal{M})$ is given by replacing $v_{B^{(1)}} - v_S$ in (15) with $v_{B^{(1)}}$ and then subtracting the expected buyer payment under $\mathcal{M}$. $U_S(\mathcal{M})$ is given by replacing $v_{B^{(1)}} - v_S$ in (15) with v_S and then subtracting that quantity from the expected buyer payment.

We now describe the allocation and expected buyer payment for the real-world and benchmark mechanisms. Under $\mathcal{M}_{RW}$, $x^{\mathcal{M}_{RW}} = 1\left\{v_{B^{(1)}} \geq g\left(\rho(v_S), p^A\right)\right\}$, where $g_R(\cdot)$ is the function described in Section 4.4 and estimation Step 8. The expected payment is $\int -T_S^0\left(\rho\left(v_S\right)\right) dF_S(v_S)$.[61] For any first-best $\mathcal{M}$ ($\mathcal{M}_S^{1^{st}}$, $\mathcal{M}_B^{1^{st}}$, or $\mathcal{M}_B^{1^{st},con}$), $x^{\mathcal{M}} = 1\{v_{B^{(1)}} \geq v_S\}$. The expected payment is given by replacing $v_{B^{(1)}} - v_S$ in (15) with $v_{B^{(1)}}$ for $\mathcal{M}_S^{1^{st}}$; with v_S for $\mathcal{M}_B^{1^{st}}$; and with $\max\{p^A, v_S\}$ for $\mathcal{M}_B^{1^{st},con}$. $\mathcal{M}_S^{2^{nd}}$ involves an optimal public reserve price, given by $z_S^*\left(v_S\right) = \arg\max_{z_S} \int \max\left\{z_S, p^A\right\} dF_{P^A}\left(p^A\right) - v_S - F_{B^{(1)}}\left(z_S\right)\left(z_S - v_S\right)$.[62] The allocation function is $x^{\mathcal{M}_S^{2^{nd}}} = 1\{v_{B^{(1)}} \geq z_S^*\left(v_S\right)\}$ and the expected payment is given by replacing $v_{B^{(1)}} - v_S$ in (15) with $\max\{z_S^*\left(v_S\right), p^A\}$. $\mathcal{M}_B^{2^{nd}}$ involves a buyer TIOLIO $z_B^*\left(v_{B^{(1)}}\right) = \arg\max_{z_B}\left(v_{B^{(1)}} - z_B\right) F_S\left(z_B\right)$. The allocation function is $x^{\mathcal{M}_B^{2^{nd}}} = 1\{v_S \leq z_B^*\left(v_{B^{(1)}}\right)\}$ and the expected payment is given by replacing $v_{B^{(1)}} - v_S$ in (15) with $z_B^*\left(v_{B^{(1)}}\right)$. Finally, $\mathcal{M}_B^{2^{nd},con}$ involves a buyer TIOLIO $z_B^*\left(v_{B^{(1)}}, p_A\right)$ maximizing $\left(v_{B^{(1)}} - z_B\right) F_S\left(z_B\right)$ subject to $z_B\left(v_{B^{(1)}}, p_A\right) \geq p^A$. The allocation function is $x^{\mathcal{M}_B^{2^{nd},con}} = 1\{v_S \leq z_B^*\left(v_{B^{(1)}}, p_A\right)\}$ and the expected payment is given by replacing $v_{B^{(1)}} - v_S$ in (15) with $z_B^*\left(v_{B^{(1)}}, p_A\right)$.

6.2 Bargaining Power Relative to First-Best Benchmarks. We now discuss the estimated expected *levels* of surplus for sellers and buyers under the real-world and benchmark

[61] We do not incorporate bargaining costs into this welfare calculation; Larsen (2021) estimated that the total expected loss due to bargaining costs is below \$33.60 for buyers and below \$5.20 for sellers.

[62] This particular formulation for the optimal reserve price comes from Aradillas-López et al. (2013).

mechanisms, beginning with first-best benchmarks. Note that all statements in the discussion of our results are statements that hold *on average*. Table 2.A shows surplus estimates and Table 2.B shows the implied bargaining power weights. The real-world mechanism, $\mathcal{M}_{RW}$, gives an expected surplus of \$1,353 to the seller and \$1,068 to the buyer, with a total surplus of \$2,420. $\mathcal{M}_S^{1^{st}}$ gives all of this surplus to the seller and none to the buyer, while $\mathcal{M}_B^{1^{st}}$ does the opposite. The seller's bargaining power implied by comparing expected surplus under $\mathcal{M}_{RW}$ to these benchmarks is $\alpha_S^{1^{st}} = 0.525$, meaning that, on average, sellers achieve 52.5% of the first-best surplus. The corresponding number for buyers is 41.4%.

Table 2: Expected Surplus and Bargaining Power – Full Sample

	A. Surplus and Trade Probability				**B.** Bargaining Power		
	U_S (\$1,000s)	U_B (\$1,000s)	Total Surplus (\$1,000s)	Trade Prob		Seller	Buyer
$\mathcal{M}_{RW}$	1.353 (0.227)	1.068 (0.015)	2.420 (0.233)	0.841 (0.005)			
$\mathcal{M}_S^{1^{st}}$	2.576 (0.234)	0 –	2.576 (0.234)	0.977 (0.006)	$\alpha^{1^{st}}$	0.525 (0.045)	0.414 (0.038)
$\mathcal{M}_B^{1^{st}}$	0 –	2.576 (0.234)	2.576 (0.234)	0.977 (0.006)			
$\mathcal{M}_B^{1^{st},con}$	1.435 (0.224)	1.141 (0.019)	2.576 (0.234)	0.977 (0.006)	$\alpha^{1^{st},con}$	-0.072 (0.012)	0.936 (0.006)
$\mathcal{M}_S^{2^{nd}}$	1.618 (0.206)	0.776 (0.042)	2.394 (0.247)	0.792 (0.024)	$\alpha^{2^{nd}}$	0.755 (0.026)	0.280 (0.044)
$\mathcal{M}_B^{2^{nd}}$	0.535 (0.245)	1.819 (0.161)	2.354 (0.258)	0.818 (0.073)			
$\mathcal{M}_B^{2^{nd},con}$	1.470 (0.233)	1.055 (0.041)	2.525 (0.239)	0.926 (0.027)	$\alpha^{2^{nd},con}$	-0.787 (0.197)	1.044 (0.101)

Notes: Panel A shows estimated expected seller and buyer surplus, total surplus, and trade probability in the real-world and benchmark mechanisms. Panel B shows estimates of α_S^m and α_B^m for four combinations of benchmark mechanisms: first-best, first-best constrained, second-best, and second-best constrained. Standard errors from 200 bootstrap samples are in parentheses.

These first-best bargaining power weights, $\alpha_S^{1^{st}}$ and $\alpha_B^{1^{st}}$, suggest that sellers achieve a greater fraction of the first-best surplus than buyers. To assess how much of this bargaining power is driven by bidder competition, we now consider the benchmark $\mathcal{M}_B^{1^{st},con}$, in which the buyer receives the most surplus possible subject to the constraint that the price weakly exceeds the auction price — and in which the seller's surplus is entirely due to competition.

This mechanism gives expected surplus of \$1,141 to the buyer, implying a competition-constrained bargaining power for the buyer of $\alpha_B^{1^{st},con} = 0.936$. In other words, buyers receive 93.6% of the surplus they would under a first-best outcome constrained only by competition, not by any information constraints ($\mathcal{M}_B^{1^{st},con}$, like the other first-best benchmarks, is not subject to information constraints).

The seller, in contrast, does worse in the real-world mechanism than under $\mathcal{M}_B^{1^{st},con}$. As noted in Section 3, the seller's expected surplus in $\mathcal{M}_B^{1^{st},con}$ equals what she would receive in an equilibrium in which she chooses $r = v_S$ and never counters (i.e., running an auction and then accepting or rejecting P^A). Our finding that $U_S(\mathcal{M}_B^{1^{st},con}) > U_S(\mathcal{M}_{RW})$ suggests the seller is giving up some surplus to the buyer. While such behavior would violate rationality for a seller who only cares about her own surplus, it is consistent with recent evidence from Yu (2024) (studying Amazon) and Rosaia (2025) (studying Uber/Lyft), who found that empirical behavior of these firms is captured well by a model in which firms place weight not just on their own profits but also on *consumer surplus*, potentially due to long-term, dynamic concerns.[63] Sellers in our data are large firms, such as Ford, Bank of America, or Hertz, and may have motivation to concede some surplus to buyers in hopes of incentivizing buyers to return for future sales.[64]

The story that emerges from the comparison of $\mathcal{M}_{RW}$ to first-best benchmarks is that, while sellers achieve a higher overall surplus level than buyers, this is entirely driven by sellers' bargaining leverage from competition among buyers. Sellers are doing no better than they would with only an auction plus an accept/reject stage, and buyers are achieving a surplus close to their most favorable outcome given the competition constraint. We thus find no evidence that sellers have more *residual* bargaining power (due to skill, patience, or lower costs) than buyers; they only enjoy the power conveyed by competition.

[63] As highlighted in Section 4, our identification of a seller's value relies on optimality of the seller's choice of secret reserve price (i.e., a choice of r maximizing her own surplus) and not on optimality of her choices at later stages of the game. If a seller is indeed yielding some surplus to buyers, our identification strategy for seller values treats the seller as choosing r optimally in a broader behavioral equilibrium in which she takes as *given* the behavior of her future self (i.e., her behavior at the post-auction stage).

[64] Consistent with this view, Steve Lang — who has participated in the wholesale used-car market as a buyer, seller, auctioneer, and part owner of an auction house — explained, in a personal communication, that large sellers tend to prioritize achieving a high probability of sale in the current auction in hopes of encouraging buyers to attend future sales of these sellers.

Table 2 also demonstrates that the total first-best surplus ($2,576) is higher than the real-world surplus, evidence of the inefficiency inherent in incomplete-information bargaining settings when type distributions have overlapping supports (Myerson and Satterthwaite 1983); $\alpha_S^{1^{st}} + \alpha_B^{1^{st}} < 1$ reflects this same point. Inefficiency can also be seen by the trade probability in the real-world (0.841) falling short of the first-best (0.977).

6.3 Bargaining Power Relative to Second-Best Benchmarks. The final three rows of Table 2.A show expected surplus estimates under second-best benchmarks — benchmarks that require IC, IR, and ex-ante budget balance. Under the seller's preferred second-best mechanism, $\mathcal{M}_S^{2^{nd}}$, the seller's expected surplus is $1,618 and the buyer's is $776. In contrast, under the buyer's preferred second-best mechanism without the competition constraint, $\mathcal{M}_B^{2^{nd}}$, the seller would achieve only $535 and the buyer $1,819. Imposing the competition constraint — that the transaction price must exceed P^A — $\mathcal{M}_B^{2^{nd,con}}$ yields $1,470 for the seller and $1,055 for the buyer. Under the first two of these mechanisms, the trade probability and total surplus is lower than in the first-best, reflective of the trade-off between efficiency and rent extraction inherent in mechanisms that must satisfy information constraints. The total surplus and trade probability are much higher under $\mathcal{M}_B^{2^{nd,con}}$ (but still lower than in the first-best), suggesting that competition in the auction goes a long way toward removing inefficiency.[65]

Table 2.B shows the bargaining power weights implied by comparing the $U_i(\mathcal{M}_{RW})$ to $U_i(\mathcal{M}_S^{2^{nd}})$ and $U_i(\mathcal{M}_B^{2^{nd}})$ for $i \in \{S, B\}$ — yielding our estimates of $(\alpha_S^{2^{nd}}, \alpha_B^{2^{nd}}) = (0.755, 0.280)$. Thus, the real-world mechanism yields a surplus to the seller that lies 75.5% of the way — and buyer's lies 28.0% of the way — between what she would achieve under $\mathcal{M}_S^{2^{nd}}$ and $\mathcal{M}_B^{2^{nd}}$. As with the bargaining weights constructed using first-best benchmarks, these weights imply that, ignoring the competition constraint, sellers are doing better than buyers relative to their preferred second-best outcomes. The final row of Table 2.B shows estimates of $(\alpha_S^{2^{nd},con}, \alpha_B^{2^{nd},con})$, which account for the competition constraint by replacing the buyer-preferred benchmark with $\mathcal{M}_B^{2^{nd,con}}$. Here we find a negative estimate of $\alpha_S^{2^{nd,con}}$. This reflects, in part, the same possibility of the seller conceding some surplus to the buyer, as discussed in

[65]This result is consistent with Williams (1999), who showed that, for a single seller facing N buyers, it is theoretically possible for outcomes to approach the first-best as N increases.

Section 6.2, but here the negative result can also be driven by inefficiency in the equilibrium played, which can lead to $U_S(\mathcal{M}_{RW}) < U_S(\mathcal{M}_B^{2^{nd},con})$.

The buyer, in contrast, has $\alpha_B^{2^{nd},con}$ above 1 (1.044), with the standard errors suggesting we cannot rule out $\alpha_B^{2^{nd},con} = 1$. As highlighted in Section 3, bargaining weights constructed with second-best benchmarks can exceed 1 if, in the real-world mechanism, some agents are not fully exploiting their information rents, consistent with the naive agents in the theoretical model of Saran (2011) or the laboratory participants in Valley et al. (2002). The finding that $\alpha_B^{2^{nd},con}$ is near 1 is closely intertwined with $\alpha_S^{2^{nd,con}} < 0$ and $\alpha_S^{1^{st},con} < 0$: sellers appear to be conceding some surplus to buyers to the extent that buyers achieve a payoff close to or potentially higher than $U_B(\alpha_B^{2^{nd},con})$.

6.4 OEM Bargaining Power. Bargaining power is of particular interest between dealers and original equipment manufacturers (OEMs). For decades, OEMs, such as Ford and GM, have been involved in the used-car market only in a *wholesale* role, without directly interacting with consumers: OEMs buy back leased vehicles at the end of the lease term and sell these to used-car dealers (through wholesale auction houses), who in turn sell to end consumers. In the past three years, OEMs have taken small steps in the U.S. toward *direct-to-consumer* (DTC) used-car sales. In February 2021, Ford launched a website, Ford Blue Advantage, that allows consumers to directly search and purchase from Ford's used-car stock.[66] GM followed in 2022 with its version, CarBravo.

This U.S. trend has been accelerated by changes in Europe, where OEMs are adopting what is referred to as an *agency model*, in which consumers search for and purchase vehicles directly from OEM websites. OEMs argue that an agency model retains a role for dealers (one of product support rather than sales) while avoiding downstream price competition between dealers, giving OEMs control over prices, and also giving OEMs market insights from customer data. The shifting European landscape applies to both new and used cars. Industry discussions highlight that the agency model is unlikely to be implemented in the U.S. *new-car market* because of state-level laws preventing DTC sales for new cars.[67] For

[66]See https://media.ford.com/content/fordmedia/fna/us/en/news/2021/02/11/ford-blue-advantage-used-vehicle-marketplace.html and https://investor.gm.com/news-releases/news-release-details/general-motors-introduces-carbravotm-new-way-shop-used-vehicles.

[67]See recent discussions in https://www.forbes.com/sites/michaeltaylor/2022/10/28/bmw-still-

used cars, however, no such restrictions exist.[68]

OEMs' motives to transition to used-car DTC sales in the U.S. may depend to some extent on how much of their potential surplus OEMs receive in the traditional wholesale model, a question to which our bargaining power metrics can speak. We divide our sample into cars sold by OEMs vs. non-OEMs and estimate bargaining power in the two samples.[69] Table 3 shows buyer and seller expected surplus and bargaining power for non-OEM sellers in panels A and B and for OEM sellers in panels C and D. We find that the total expected surplus in the real-world mechanism is higher when the seller is an OEM: \$3,602 as opposed to \$2,137. This higher total surplus for OEMs is driven by *sellers'* surplus being higher for OEMs (\$2,538 vs. \$1,101); buyers' real-world surplus is nearly equivalent in the two cases. The total first-best surplus is also higher for OEM sales than non-OEM sales. The first row of panels B and D shows the bargaining weights relative to first–best benchmarks constructed without enforcing the competition constraint: $(\alpha_S^{1^{st}}, \alpha_B^{1^{st}}) = (0.671, 0.281)$ in the OEM sample vs. (0.472,0.444) in the non-OEM sample. Thus, the potential surplus is higher in OEM sales, and OEM sellers appear to capture more of that surplus than buyers and more than non-OEM sellers.

We find similar results comparing to the unconstrained second-best benchmark, with $(\alpha_S^{2^{nd}}, \alpha_B^{2^{nd}}) = (0.918, 0.048)$ in the OEM sample vs. (0.631,0.318) in the non-OEM sample. This low $\alpha_B^{2^{nd}}$ (0.048) in the OEM sample is evidence that competition drives the overall result: it suggests that the buyer's real-world payoff (\$1,064) lies only 4.8% of the way between what he receives under $\mathcal{M}_S^{2^{nd}}$ (\$938) and the vastly larger expected payoff (\$3,574) if he could make a TIOLIO offer to the seller without facing any competition.

We now compare the real-world outcomes to benchmarks that account for the competition

---

wants-an-agency-model--but-not-for-the-us/ and https://www.am-online.com/news/market-insight/2023/08/21/dealers-troubled-by-agency-model-impact-on-used-car-supply. Lafontaine and Scott Morton (2010) discussed the economics of state-level laws governing the U.S. new-car market.

[68]One industry expert commented in September 2022, "The scary part is that if any manufacturer wanted to exert more force and control in the used-car market for late-model used cars, I'm not sure any state franchise laws governing new-car sales would hinder them. Again, the OEMs are the largest 'manufacturers' of used cars, financing and owning hundreds of thousands of off-lease, off-fleet and off-rental vehicles. At the end of the term, these are used cars, so any controls or retailing dictates are fair game, as I believe state franchise laws only focus on new-vehicle sales protections." See https://www.wardsauto.com/financials/will-oems-compete-with-dealers-for-used-car-sales-, which also describes Europe/Africa shift to DTC used-car sales, with the prime example being Spoticar, the DTC website of Stellantis, the parent of Chrysler.

[69]Appendix B.7 discusses model fit in these subsamples and the data variation driving our results.

Table 3: Expected Surplus and Bargaining Power – OEM vs. Non-OEM

A. Non-OEM Sample: Surplus and Trade Probability					**B.** Non-OEM Sample: Barg. Power		
	U_S (\$1,000s)	U_B (\$1,000s)	Total Surplus (\$1,000s)	Trade Prob		Seller	Buyer
$\mathcal{M}_{RW}$	1.101	1.036	2.137	0.828			
	(0.178)	(0.025)	(0.180)	(0.019)			
$\mathcal{M}_S^{1^{st}}$	2.333	0	2.333	0.965	$\alpha^{1^{st}}$	0.472	0.444
	(0.183)	–	(0.183)	(0.005)		(0.045)	(0.043)
$\mathcal{M}_B^{1^{st}}$	0	2.333	2.333	0.965			
	–	(0.183)	(0.183)	(0.005)			
$\mathcal{M}_B^{1^{st},con}$	1.201	1.132	2.333	0.965	$\alpha^{1^{st},con}$	-0.089	0.915
	(0.173)	(0.019)	(0.183)	(0.005)		(0.009)	(0.025)
$\mathcal{M}_S^{2^{nd}}$	1.415	0.719	2.134	0.741	$\alpha^{2^{nd}}$	0.631	0.318
	(0.156)	(0.037)	(0.193)	(0.021)		(0.026)	(0.038)
$\mathcal{M}_B^{2^{nd}}$	0.564	1.713	2.277	0.873			
	(0.181)	(0.086)	(0.169)	(0.049)			
$\mathcal{M}_B^{2^{nd},con}$	1.259	1.063	2.322	0.935	$\alpha^{2^{nd},con}$	-1.017	0.920
	(0.184)	(0.028)	(0.182)	(0.017)		(0.179)	(0.069)
C. OEM Sample: Surplus and Trade Probability					**D.** OEM Sample: Barg. Power		
$\mathcal{M}_{RW}$	2.538	1.064	3.602	0.912			
	(0.325)	(0.024)	(0.324)	(0.006)			
$\mathcal{M}_S^{1^{st}}$	3.784	0	3.784	0.990	$\alpha^{1^{st}}$	0.671	0.281
	(0.346)	–	(0.346)	(0.005)		(0.029)	(0.029)
$\mathcal{M}_B^{1^{st}}$	0	3.784	3.784	0.990			
	–	(0.346)	(0.346)	(0.005)			
$\mathcal{M}_B^{1^{st},con}$	2.702	1.082	3.784	0.990	$\alpha^{1^{st},con}$	-0.152	0.983
	(0.348)	(0.025)	(0.346)	(0.005)		(0.030)	(0.003)
$\mathcal{M}_S^{2^{nd}}$	2.763	0.938	3.700	0.934	$\alpha^{2^{nd}}$	0.918	0.048
	(0.331)	(0.036)	(0.364)	(0.018)		(0.013)	(0.020)
$\mathcal{M}_B^{2^{nd}}$	0.009	3.574	3.582	0.906			
	(0.035)	(0.361)	(0.354)	(0.049)			
$\mathcal{M}_B^{2^{nd},con}$	2.745	1.021	3.765	0.973	$\alpha^{2^{nd},con}$	-11.705	1.525
	(0.370)	(0.060)	(0.342)	(0.014)		(109.969)	(7.588)

Notes: Panels A and C show estimated expected seller and buyer surplus, total surplus, and trade probability in the real-world and benchmark mechanisms. Panels B and D show estimates of α_S^m and α_B^m for four combinations of benchmark mechanisms: first-best, first-best constrained, second-best, and second-best constrained. Panels A and B use non-OEM sales and panels C and D use OEM sales. Standard errors from 200 bootstrap samples are in parentheses.

constraint. Our finding is that bargaining power of OEMs — as well as that of non-OEMs — is driven purely by bargaining leverage from bidder competition: as in the full sample, both the non-OEM and OEM sales show an expected buyer surplus quite close to what buyers would achieve under the competition-constrained first-best. Indeed, this finding that leverage drives sellers' bargaining power is most salient in the OEM sample, where $\alpha_B^{1^{st},con} = 0.983$ (compared to 0.915 in the non-OEM sample). A similar story emerges from the seller's first-best constrained bargaining weight: $\alpha_S^{1^{st},con} = -0.152$ in the OEM sample and -0.089 in the non-OEM sample, implying that the seller's payoff is even farther below the competition-only payoff, $U_S(\mathcal{M}_B^{1^{st},con})$, for OEM sellers than for non-OEM sellers.

The final rows of panels B and D tell a similar qualitative story: sellers' real-world outcomes are worse than in the second-best buyer-preferred mechanism constrained by competition ($\mathcal{M}_B^{2^{nd},con}$), reflecting an inefficient equilibrium. Buyers' outcomes in the real-world for non-OEM sales are only slightly below their expected payoff under $\mathcal{M}_B^{2^{nd},con}$. For OEM sales, buyers in the real world do even better than under $\mathcal{M}_B^{2^{nd},con}$, with an $\alpha_B^{2^{nd},con} > 1$ (with standard errors suggesting we cannot rule out $\alpha_B^{2^{nd},con} < 1$), reflecting a real-world outcome that potentially lies between the first- and second-best frontiers, consistent with OEM sellers not fully exploiting their information rent (or in some other fashion conceding some surplus to buyers). Note that, for the OEM sample, the point estimates of $\alpha_S^{2^{nd},con}$ and $\alpha_B^{2^{nd},con}$ are larger in magnitude and more noisily estimated. This reflects a point raised in Section 3: a bargaining weight constructed from benchmarks $\mathcal{M}_S^{2^{nd}}$ and $\mathcal{M}_B^{2^{nd},con}$ is less likely to be useful if the degree of competition is sufficiently high that there is little difference between the two benchmarks, as the construction of the bargaining weight involves dividing by a small number in such cases. Table 3 shows that this is indeed the case, where the point estimates of seller surplus under these two mechanisms are only \$18 apart in the OEM sample, whereas they are much farther apart in the full sample and non-OEM sample.

Together, these results suggest that, while OEMs achieve an overall surplus close to the seller-optimal level, this is due to dealer competition: OEMs, like non-OEMs, achieve little or no price increases *above* the auction price. This lack of OEM bargaining power beyond that granted by competition, along with the apparent inefficiency of the real-world bargaining equilibrium, is consistent with the idea that OEMs in the traditional wholesale market may indeed benefit from attempting a different sales model, such as DTC sales.[70]

7 Conclusion

This study analyzes bargaining power under incomplete information. We focus on negotiations between buyers (car dealers) and sellers (large institutions, such as manufacturers or fleet-owning companies) in the supply side of the U.S. used-car market. These negotiations are facilitated by wholesale used-car auction houses, which run a secret-reserve-price ascending auction and facilitate offers and counteroffers between the seller and highest bidder.

The private value distribution of the buyer can be estimated from data on auction prices. The distribution for 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. The derivative of this menu evaluated at the point chosen by the seller corresponds to the seller's privately known value.

As bargaining power is not a well-studied concept in incomplete-information bargaining, we propose a new definition: an agent's bargaining power describes where the agent's surplus lies between two benchmark mechanisms: one preferred by the seller and one by the buyer. This extends a traditional (complete-information, Nash bargaining) notion of bargaining power to an incomplete-information setting. We consider both first- and second-best benchmarks, as well as benchmarks that explicitly account for the bargaining leverage a seller achieves through competition among bidders.

We find that, on average, car dealers (buyers) achieve a lower level of surplus than sellers

[70]The success of this attempt may depend on whether OEMs can reach consumers whose willingness-to-pay distribution rivals that of $F_{V_{B(1)}|P^A}$, the distribution OEMs face at auctions. Bose and Deltas (2002) provided a model in which sellers prefer sales to *resellers* (analogous to dealers at our auctions) rather than DTC sales if dealers have access to a stronger buyer market downstream than what the original sellers can access. Thus, OEMs' steps toward DTC sales may be experimental attempts to learn whether they have the marketing and technological capability to reach consumers willing to pay as much as auction bidders. Other explanations for OEMs' movement toward DTC used-car sales include the COVID-19 pandemic (with its accompanying decrease of in-person interactions at dealerships) and supply-chain shortages (discussed in the industry articles cited herein). Note that if we had instead found that OEMs have strong bargaining power even after accounting for the competition effect, and simultaneously found that the equilibrium in the wholesale market is relatively efficient, we would instead infer that OEMs' shift to DTC sales is less likely to be motivated by a desire to overcome low bargaining power in the wholesale market.

(large companies), suggesting buyers have less bargaining power than buyers relative to first- or second-best benchmarks. However, we find that this higher bargaining power of sellers is driven by competition among buyers: sellers do not achieve more expected surplus than they would if all of their bargaining power were to derive only from buyer competition, and buyers achieve an expected surplus level that is only slightly below what they would achieve in the most favorable outcome for buyers in a first-best world subject to competition. Competition among buyers thus gives seller's substantial bargaining leverage but little residual bargaining power to raise prices above the auction price.

We examine the relationship between manufacturers (OEMs) and dealers in used-car markets. While OEMs achieve a surplus level close to their optimal payoff, this is again driven by dealer competition, just as with non-OEM sales. OEMs are unable to push prices much beyond the auction price. This lack of power may be one motivation for recent moves by major OEMs (Ford and GM) toward selling used-cars directly to consumers rather than to dealers at wholesale auto auctions. 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 incomplete information plays a role.

References

Agarwal, N., Li, P. Z., and Somaini, P. J. (2023). Identification using revealed preferences in linearly separable models. NBER WP 31868.

Aradillas-López, A., Gandhi, A., and Quint, D. (2013). Identification and inference in ascending auctions with correlated private values. *ECMA*, 81(2):489–534.

Asil, A., Ramos, P., Starc, A., and Wollmann, T. G. (2024). Painful bargaining: Evidence from anesthesia rollups. NBER WP 33217.

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

Ausubel, L. and Deneckere, R. (1992). Durable goods monopoly with incomplete information. *RESTUD*, 59(4):795–812.

Ausubel, L. and Deneckere, R. (1993). Efficient sequential bargaining. *RESTUD*, 60(2):435.

Binmore, K., Rubinstein, A., and Wolinsky, A. (1986). The Nash bargaining solution in economic modeling. *RAND*, pages 176–188.

Bose, S. and Deltas, G. (2002). Exclusive dealing through resellers in auctions with stochastic bidder participation. *SEJ*, 69(1):109–127.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *AER*, 86(1):180–194.

Capps, C., Dranove, D., and Satterthwaite, M. (2003). Competition and market power in option demand markets. *RAND*, pages 737–763.

Cramton, P. (1992). Strategic delay in bargaining with two-sided uncertainty. *RESTUD*, 59(1):205–225.

Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *AER*, 102(2):643–685.

Donna, J. D., Pereira, P., Pu, Y., Trindade, A., and Yoshida, R. C. (2024). Direct sales and bargaining. *RAND*, 55(4):749–787.

Elyakime, B., Laffont, J. J., Loisel, P., and Vuong, Q. (1994). First-price sealed-bid auctions with secret reservation prices. *Annales d'Economie et de Statistique*, 34:115–141.

Evdokimov, K. and White, H. (2012). Some extensions of a lemma of Kotlarski. *Econometric Theory*, 28(4):925–932.

Freyberger, J. and Larsen, B. J. (2025). How well does bargaining work in consumer markets? A robust bounds approach. *ECMA*, 93(1):161–194.

Gallant, A. and Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. *ECMA*, 55(2):363–390.

Görlach, J.-S. and Motz, N. (2024). A general measure of bargaining power for non-cooperative games. IZA Discussion Paper No. 16809.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *AER*, 105(1):172–203.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *AER*, 103(1):145–177.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first price auctions. *ECMA*, 68(3):525–574.

Gul, F. and Sonnenschein, H. (1988). On delay in bargaining with one-sided uncertainty. *ECMA*, 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. *AER*, 109(2):473–522.

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

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

Keniston, D., Larsen, B. J., Li, S., Prescott, J., Silveira, B. S., and Yu, C. (2024). Splitting the difference in incomplete-information bargaining: Theory and widespread evidence from the field. *IER*, 65(4):1911–1939.

Kline, B. (2023). Identification of incomplete information games in monotone equilibrium. WP, UT Austin.

Kong, Y., Silveira, B. S., and Tang, X. (2024). Risk and information in dispute resolution: An empirical study of arbitration. *JPE*, forthcoming.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *RESTUD*, 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. *AEJ-Micro*, 8(4):195–229.

Lafontaine, F. and Scott Morton, F. (2010). Markets: State franchise laws, dealer terminations, and the auto crisis. *JEP*, 24(3):233–50.

Larsen, B. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-

auto auctions. *RESTUD*, 88(2):851–882.
Larsen, B., Lu, C., and Zhang, A. L. (2024). Intermediaries in bargaining: Evidence from business-to-business used-car inventory negotiations. NBER WP 29159.
Larsen, B. and Zhang, A. (2018). A mechanism design approach to identification and estimation. NBER WP 24837.
Loertscher, S. and Marx, L. M. (2019). Merger review for markets with buyer power. *JPE*, 127(9).
Loertscher, S. and Marx, L. M. (2022). Incomplete information bargaining with applications to mergers, investment, and vertical integration. *AER*, 112(2):616–649.
Luo, Y., Perrigne, I., and Vuong, Q. (2018). Structural analysis of nonlinear pricing. *JPE*, 126(6):2523–2568.
Menezes, F. and Ryan, M. J. (2005). Reserve price commitments in auctions. *EL*, 87(1):35–39.
Meyer, M. C. (2008). Inference using shape-restricted regression splines. *Ann. App. Stat.*, 2(3):1013–1033.
Murry, C. and Zhou, Y. (2020). Consumer search and automobile dealer colocation. *Man. Sci.*, 66(5):1909–1934.
Myerson, R. (1979). Incentive compatibility and the bargaining problem. *ECMA*, pages 61–73.
Myerson, R. (1981). Optimal auction design. *Math. of OR*, 6(1):58–73.
Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *JET*, 29(2):265–281.
Perrigne, I. and Vuong, Q. (2011). Nonparametric identification of a contract model with adverse selection and moral hazard. *ECMA*, 79(5):1499–1539.
Perry, M. (1986). An example of price formation in bilateral situations: A bargaining model with incomplete information. *ECMA*, pages 313–321.
Peters, C. T. (2014). Bargaining power and the effects of joint negotiation: The "recapture effect". EAG Discussion Paper 14-3.
Pinkse, J. and Schurter, K. (2019). Estimation of auction models with shape restrictions. arXiv:1912.07466.
Ramsay, J. O. (1988). Monotone regression splines in action. *Stat. Sci.*, pages 425–441.
Rosaia, N. (2025). Competing platforms and transport equilibrium: Evidence from New York City. WP, Columbia.
Saran, R. (2011). Bilateral trading with naive traders. *GEB*, 72(2):544–557.
Storms, E. (2015). Monotonicity of trade in PSBNE of bargaining with alternating offers. Mimeo, Stanford.
Valley, K., Thompson, L., Gibbons, R., and Bazerman, M. H. (2002). How communication improves efficiency in bargaining games. *GEB*, 38(1):127–155.
Williams, S. R. (1987). Efficient performance in two agent bargaining. *JET*, 41(1):154–172.
Williams, S. R. (1999). A characterization of efficient, Bayesian incentive compatible mechanisms. *ET*, pages 155–180.
Yu, C. (2024). The welfare effects of sponsored product advertising. WP, UNC.

Appendix for Online Publication

A Model Discussion and Proofs

The models in Larsen (2021) and Larsen et al. (2024) incorporated an additive per-offer bargaining cost for sellers (in addition to buyers); we do not incorporate this feature here for several reasons. First, the properties of Proposition 1 do not rely on the existence of seller bargaining costs. A bargaining cost for buyers, on the other hand, is important in the model to rationalize why some buyers opt out of bargaining, and this opt-out possibility is important for the property that the seller's secret reserve price strategy is strictly increasing. Second, Larsen (2021) estimated an upper bound on the total expected seller surplus lost due to bargaining costs that is very small (less than \$5.2). Third, appropriately incorporating such seller bargaining costs requires modifying the menu approach, for example, to account for the seller's expected probability of trading *in each period* of the game, not just her expected probability of trading overall. While theoretically feasible, such a model would be cumbersome and add little benefit over the current model (given the first and second points highlighted in this paragraph).

Proof of Proposition 1. Part (i) of Proposition 1 follows from Proposition 1 of Larsen (2021). Part (ii) follows from Larsen (2021) Proposition 3. Part (iii) follows from Larsen et al. (2024) Proposition 5. □

Proof of Corollary 2. By $A7$, $\rho(v_S)$ is strictly increasing and $P_S(r)$ is strictly monotone. 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 $\left[\underline{b}, \overline{b}\right]$ denote the support of V_B. Choosing any secret reserve price r below $\underline{b}$ is a dominated action for the seller given that every buyer has a value of at least $\underline{b}$. Moreover, the seller is indifferent between a secret reserve price of $\overline{b}$ and any secret reserve price higher than this because no buyer would ever be willing to pay more than $\overline{b}$. Thus, the support of R can be treated as being bounded within $\left[\underline{b}, \overline{b}\right]$.

Consider some value v_S strictly in the interior of the interval $\left[\underline{b}, \overline{b}\right]$. First, we apply the upper bound from Corollary 1, 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,

$$v_S \leq \frac{T_S(\rho(v_S+\delta)) - T_S(\rho(v_S))}{P_S(\rho(v_S+\delta)) - P_S(\rho(v_S))} \tag{16}$$

where the right-hand side of (16) always exists, because by assumption both $\rho(\cdot)$ and $P_S(\cdot)$ are strictly monotone. Now, let $\delta \to 0$. Because ρ is strictly monotone and $T_S(\cdot)$ and $P_S(\cdot)$ are differentiable, we have

$$\lim_{\delta\to 0} \frac{T_S(\rho(v_S+\delta)) - T_S(\rho(v_S))}{P_S(\rho(v_S+\delta)) - P_S(\rho(v_S))} = \frac{T'_S(\rho(v_S))}{P'_S(\rho(v_S))} \tag{17}$$

The ratio of derivatives on the right-hand side of (17) always exists, because by assumption $T_S(\cdot)$ is differentiable, and $P_S(\cdot)$ is strictly monotone and differentiable, so $P'_S(\rho(v_S)) \neq 0$ for all v_S. Thus, the bound in (16) becomes:

$$v_S \leq \frac{T'_S(\rho(v_S))}{P'_S(\rho(v_S))} \tag{18}$$

Next, applying the lower bound from Corollary 1, we also have, for any δ,

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

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

$$v_S \geq \frac{T'_S(\rho(v_S))}{P'_S(\rho(v_S))} \tag{19}$$

Combining (18) and (19), along with the fact that $v_S(r)$ is the inverse of $\rho(v_S)$, yields the desired result for any v_S in the interior of $\left[\underline{b}, \bar{M}\right]$.

Now, for $v_S = \bar{M}$, we have:

$$v_S\left(\bar{M}\right) = \lim_{\epsilon\to 0} v_S\left(\bar{M}-\epsilon\right) = \lim_{\epsilon\to 0} \frac{T'_S\left(\rho\left(\bar{M}-\epsilon\right)\right)}{P'_S\left(\rho\left(\bar{M}-\epsilon\right)\right)} = \frac{T'_S\left(\rho\left(\bar{M}\right)\right)}{P'_S\left(\rho\left(\bar{M}\right)\right)}$$

where the last equality follows by continuous differentiability of T_S and P_S ($A7$). The proof that $v_S(\underline{b}) = \frac{T'_S(\underline{b})}{P'_S(\underline{b})}$ is analogous. □

Proof of Theorem 2. We prove the result by contradiction. Suppose 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 in equilibrium for some types, such that

$$\gamma P_S(r') + (1-\gamma) P_S(r'') = P_S(r) \tag{20}$$

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

$$T_S(r) > \gamma T_S(r') + (1-\gamma) T_S(r'') \tag{21}$$

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

$$v_S P_S(r) - T_S(r) \tag{22}$$

Playing r' with probability γ and r'' otherwise would yield instead

$$v_S [\gamma P_S(r') + (1-\gamma) P_S(r'')] - [\gamma T_S(r') + (1-\gamma) T_S(r'')] \tag{23}$$

Plugging (20) into (23) yields $v_S P_S(r) - [\gamma T_S(r') + (1-\gamma) T_S(r'')]$, which, by (21), is strictly greater than (22). Because r is optimal for type v_S, this yields a contradiction. □

Proof of Theorem 3. We describe identification separately for probabilities and transfers.

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

$$\begin{aligned}
\tilde{P}_S(\tilde{r}) &= E_{R,V_{-S},W}[x_S(R+W, V_{-S}+W) \mid R+W=\tilde{r}] && (24)\\
&= E_{R,V_{-S},W}\left[E_{V_{-S}}[x_S(R+W, V_{-S}+W) \mid W, R+W=\tilde{r}] \mid R+W=\tilde{r}\right] && (25)\\
&= E_{R,W}\left[P_S^W(R+W) \mid R+W=\tilde{r}\right] && (26)\\
&= E_{R,W}\left[P_S^0(R) \mid R=\tilde{r}-W\right] && (27)
\end{aligned}$$

Expression (25) follows from applying the law of iterated expectations to (24). (26) 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 (26) and (27) follows because of location invariance in Proposition 1.iii, which implies $P_S^w(r+w) = P_S^0(r)\ \forall w, r$. That is, the probability 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 (26) with $P_S^0(R)$ in (27). In integral form, expression (27) corresponds to (9) in the main text. In words, (9) 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. To show that $P_S^0(r)$ is identified, note that we can write (9) as:

$$\tilde{P}_S(\tilde{r}) \int f_R(r) f_W(\tilde{r}-r)\, dr = \int P_S^0(r) f_R(r) f_W(\tilde{r}-r)\, dr \tag{28}$$

The left-hand side of (28) involves $\tilde{P}_S(\tilde{r})$, $f_R(\cdot)$, and $f_W(\cdot)$, all of which are identified.

The right-hand side is the convolution of $P_S^0(r)\, f_R(r)$ against $f_W(r)$. $f_R(\cdot)$ and $f_W(\cdot)$ are known. Note that the distribution of W has bounded support and, as shown in the proof of Corollary 2, R will as well. Therefore, $P_S^0(r)\, f_R(r)$ is 0 for all r outside the support of R. Thus, both $P_S^0(r)\, f_R(r)$ and $f_W(r)$ are in L^1 and the convolution theorem applies, meaning that the convolution of $P_S^0(r)\, f_R(r)$ and $f_W(r)$ is invertible. Thus, $P_S^0(r)\, f_R(r)$ is identified. Therefore, at any realization of R that occurs in equilibrium for some seller type (and hence, where $f_R(r)$ is positive), $P_S^0(r)$ is identified.

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

$$\begin{aligned}
\tilde{T}_S(\tilde{r}) &= E_{R,V_{-S},W}\left[t_S(R+W, V_{-S}+W) \mid R+W=\tilde{r}\right] && (29)\\
&= E_{R,V_{-S},W}\left[E_{V_{-S}}\left[t_S(R+W, V_{-S}+W) \mid W, R+W=\tilde{r}\right] \mid R+W=\tilde{r}\right] && (30)\\
&= E_{R,W}\left[T_S^W(R+W) \mid R+W=\tilde{r}\right] && (31)\\
&= E_{R,W}\left[T_S^0(R)+(\tilde{r}-R)\left(P_S^0(R)-1\right) \mid W=\tilde{r}-R\right] && (32)
\end{aligned}$$

These equations are similar to expressions (24) to (27) above. (30) follows from applying the law of iterated expectations to (29). (31) 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, (32) follows from applying (8) to (31). In integral form, (32) corresponds to (10) from the main text. We can then write:

$$\left(\tilde{T}_S(\tilde{r}) - E\left[W\Delta P_S \mid \tilde{r}\right]\right)\left(\int f_R(r)\, f_W(\tilde{r}-r)\, dr\right) = \int T_S^0(r)\, f_R(r)\, f_W(\tilde{r}-r)\, dr, \quad (33)$$

In the left-hand side of (33), the term $\tilde{T}_S(\tilde{r})$ is identified in the data and the term $E\left[W\Delta P_S \mid \tilde{r}\right]$ can be calculated for any $\tilde{r}$ using (11) because $P_S^0(r)$ is identified. Also, $f_R(\cdot)$ and $f_W(\cdot)$ are identified. Thus, the left-hand side of (33) is identified. The right-hand side of (33) is the convolution of $T_S^0(r)\, f_R(r)$ against $f_W(r)$. By the convolution theorem, this convolution is invertible, and thus $T_S^0(r)\, f_R(r)$ is identified, hence $T_S^0(r)$ is identified at any realization of R that occurs in equilibrium for some seller type. □

B Additional Estimation Results and Discussion

B.1 Data Cleaning

We use the data of Larsen (2021), with the additional restriction that we limit to the first instance of a given car being offered for sale at a given auction house; here we refer to this as the *first run* restriction. Some of the data cleaning steps in Larsen (2021) include dropping observations with missing variables, incoherent bargaining sequences, extreme prices (those outside the 0.01 or 0.99 quantiles of a given price variable), or car types (a make-model-year-trim-age combination) that are sold less than ten times in the data. A total of 9 observations are dropped due to cases where the recorded final price lies below the recorded auction price. See Table A5 of Larsen (2021) for a full set of restrictions and the number of observations dropped from each.

As described in Section 5.1, our estimation Step 1 uses an augmented dataset to aid in controlling for as much observable variation as possible and, in particular, to help with estimation of our many category dummies (e.g., make-model-year-trim-age dummies). The augmented dataset includes 39,700 observations corresponding to repeat attempts to sell a given car, 15,034 observations in which a reserve price is recorded but not an auction price, and 65,179 observations in which an auction price is recorded but not a reserve price. A missing auction price corresponds to a case where a scheduled auction sale did not occur or where an auctioneer attempted to run the sale but there was insufficient activity from bidders. A missing reserve price corresponds to a case where the seller, rather than reporting a secret reserve price, requests that the auction house call her post-auction to let her accept the auction price or bargain further (effectively reporting a very high reserve price). The identification of our structural model requires observing both auction prices and secret reserve prices, and thus we cannot use these observations in our main estimation steps (Steps 2–8). Our results should thus be interpreted as quantifying bargaining power conditional on the types of cases that appear in our final sample.

To explore whether differences in these distinct types of observations in the augmented dataset drive any of our results, we estimate an alternative version of (13) — the stacked regression in estimation Step 1 — with a fully saturated set of dummy interactions for these distinct types of observations. Specifically, we include a dummy for whether a stacked observation corresponds a reserve vs. auction price, a dummy for whether the observation corresponds to an instance where only a reserve or auction price is observed for a given car, and a dummy for whether the observation is a repeat sales attempt, and all interactions of the preceding three dummies. Within our final sample, the predicted values from this alternative regression have a correlation of 0.999 with those from (13).

We now analyze differences between cases reserve prices are recorded vs. missing. For this exercise, we still impose the first run restriction on our sample. For three different outcomes — $\widetilde{price}_j$, $\tilde{p}_j^A$, and a trade indicator — we regress the outcome on an indicator for whether reserve prices are recorded. The results indicate that, in observations with missing secret reserve prices, the trade probability is lower by 1.1 percentage points, $\tilde{p}_j^A$ is higher by \$88, and $\widetilde{price}_j$ (conditional on trade) is higher by \$53. We also perform a similar regression comparing cases where auction prices are recorded vs. missing, regressing $\tilde{r}_j$ on a dummy for the auction price being missing. The estimated coefficient suggests that reserve prices are \$428 higher in observations with missing auction prices. These results are each statistically significant, but are not large relative to the means in Table 1 (average trade probability is 0.86, and average book value is \$11,030) or relative to the standard deviations of $\tilde{p}_j^A$, $\widetilde{price}_j$, or $\tilde{r}_j$ (which all exceed \$1,100).

B.2 Residuals From Observable Heterogeneity Regression

This section examines the residuals from the first-step regression. In Figure A.1.A 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 as indicating where the noisy reserve price lies relative to the market value estimate of the car. Figure A.1.B displays, on the vertical axis, $E[\tilde{P}^A|\tilde{R}]$ and $E[\widetilde{price}|\tilde{R}]$, again from local linear regressions.[71] Here we observe that higher reserve

[71]The regressions in each panel use observations with noisy reserve prices lying between the 0.01 and 0.99 quantiles of empirical noisy reserve prices.

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

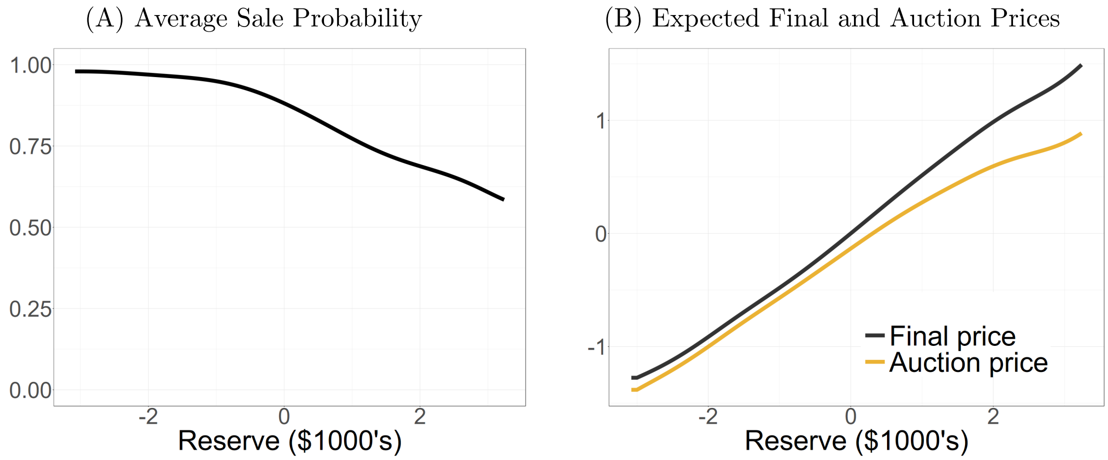

Notes: Panel A displays local linear regression estimates of an indicator for whether the car sold regressed on observations of $\tilde{R}$. Panel B contains similar local linear regression estimates where the outcome is observations of $\tilde{P}^A$ (the high bid from the auction) in yellow and observations of $\widetilde{price}$ (conditional on a sale occurring) in black. Units are in terms of $1,000, relative to the market value estimate. Uses main sample from the body of the paper.

prices are associated with expected final prices that represent a higher markup over the auction price.

The objects $\tilde{P}^A$ and $\tilde{R}$ represent the components of auction prices and reserve prices after controlling for observable heterogeneity. Figure A.1.B demonstrates that these objects are correlated with one another, underscoring the importance of accounting for *unobserved heterogeneity* in our setting. Sellers who choose higher reserve prices sell with lower probabilities (panel A) but attain higher prices conditional on sale (panel B). Importantly, the difference between the average final price conditional on sale and the average auction price is increasing in the reserve price. Note that the average auction price roughly measures the value of unobserved car-level heterogeneity conditional on the reserve price.[72] The fact that the *difference* between the final price and the auction price is *increasing* in the reserve price suggests that sellers who choose high reserve prices are forgoing some sale probability in

[72]To 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

$$E\left[\tilde{P}^A \mid \tilde{R}\right] = E\left[P^A + W \mid R + W\right] = E\left[P^A \mid R + W\right] + E\left[W \mid R + W\right] = E\left[P^A\right] + E\left[W \mid \tilde{R}\right]$$

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

Figure A.2: Results Relating to $\Pr(N = n)$

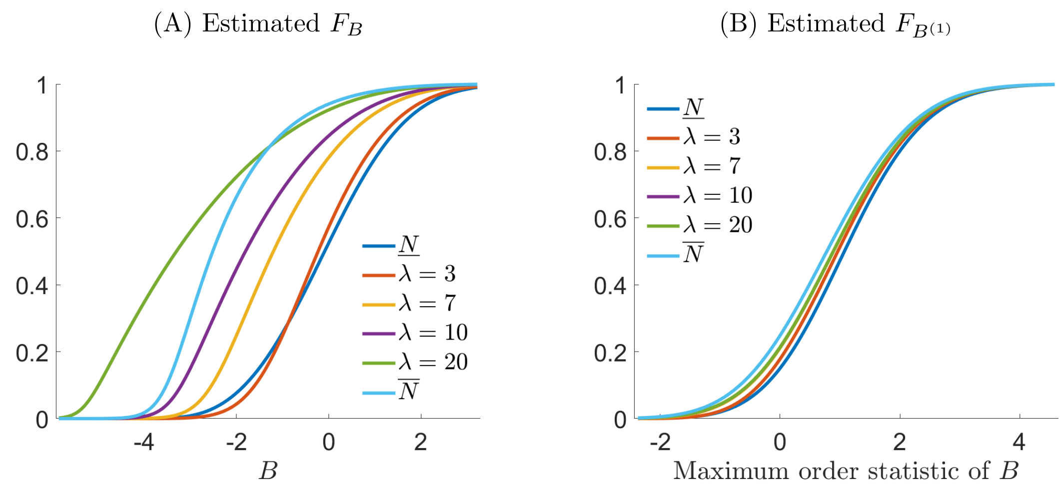

Notes: Panels A and B show the estimated F_B and $F_{B^{(1)}}$, respectively, obtained used different specifications for $\Pr(N = n)$: four Poisson specifications, with the Poisson parameter λ given by 3, 7, 10, or 20, and an auction-by-auction lower or upper bound on N, denoted $\underline{N}$ and $\overline{N}$, respectively. Similar figures appear as online appendix Figures A6.B and A6.D in Larsen (2021) for the samples used in that study. Units on horizontal axes are \$1,000.

order to obtain a higher sale price, as the menu approach requires.

B.3 Robustness to Specification of $\Pr(N = n)$

As shown in (4), identification and estimation of the underlying distribution of buyer values, F_B, relies on identification of F_{P^A} (the distribution of the second-highest value) and $\Pr(N = n)$, the probability mass function (PMF) of N. The construction of $F_{B^{(1)}}$, the maximum order statistic distribution, in turn relies on F_B and that same PMF, as shown in (5). It turns out that the implied $F_{B^{(1)}}$ from this procedure is not sensitive to the choice of the PMF for N; F_B is sensitive to this choice, but $F_{B^{(1)}}$ is not.

To see this, Figure A.2 shows estimates of F_B (in panel A) and $F_{B^{(1)}}$ (in panel B) under different specifications for $\Pr(N = n)$: four Poisson specifications, with the Poisson parameter given by 3, 7, 10, or 20, and an auction-by-auction lower or upper bound on N, denoted $\underline{N}$ and $\overline{N}$, respectively.[73] Panel A shows that the underlying estimate of F_B is naturally

[73] $\underline{N}$ is explained in Section 5.1. $\overline{N}$ is the sum of two objects: (i) the number of in-person bids and (ii) the number of bidders *registered* to bid online. Not all online bidders necessarily submit bids; by registering online, they have gained permission to access the online portal where video for a specific physical auction lane will be live-streamed. Thus, the number of bidders registered to bid online is an upper bound on the number of bidders actually participating online. Note that, anywhere in the paper where observations of $\underline{N}$

sensitive to the choice of $\Pr(N = n)$: treating the auction price distribution as though it represents a draw from a distribution with more bidders leads to the inference that the underlying F_B has more mass at low V_B. However, the estimate of $F_{B^{(1)}}$, constructed using the estimated F_B and $\Pr(N = n)$, varies little — even across widely different specifications for $\Pr(N = n)$ — as shown in panel B.

The insensitivity of the implied $F_{B^{(1)}}$ to the specification of $\Pr(N = n)$ is not just an empirical artifact: Larsen (2021) Proposition 10 demonstrated that, for any Poisson specification for $\Pr(N = n)$, the implied CDF $F_{B^{(1)}}(y)$ is completely insensitive to the Poisson parameter. Specifically, consider the following procedure: for a given Poisson specification of $\Pr(N = n)$ with mean λ (i) use $\Pr(N = n)$ in the process of inverting F_{P^A} to obtain F_B by solving (4) and then (ii) use F_B and $\Pr(N = n)$ to construct $F_{B^{(1)}}$ using (5). Proposition 10 of Larsen (2021) showed that the derivative of the implied $F_{B^{(1)}}(y)$ at any point y with respect to λ is identically zero. Panel B shows that this robustness holds even for the lower and upper bound PMFs, which are not necessarily Poisson.

This lack of sensitivity is important for our analysis because the inputs for our bargaining power metrics all rely on $F_{B^{(1)}}$, not F_B directly. This can be seen in the expressions for surplus in Section 6.1. Note that in some of the surplus expressions in Section 6.1, the integral against the maximum order statistic of buyer values takes the form of the density of the maximum order statistic *conditional* on the second-highest, $dF_{B^{(1)}|P^A}$, which is then integrated against the density of the auction price, dF_{P^A}, yielding again the density of the maximum order statistic.

This relates to a general property of button-like auctions with symmetric IPV: the density of a higher order statistic of bids conditional on a lower order statistic does not depend on N. Song (2004) proves this result generally. In our case, it takes the form of $F_{B^{(1)}|P^A}\left(v_{B^{(1)}} \mid p^A\right) \equiv \frac{1-F_B(v_{B^{(1)}})}{1-F_B(p^A)}$, presented in Section 4.4, showing that the conditional density of the highest bidder's value conditional on the auction price is independent of the number of bidders. This is also related to the well-known point (e.g., Riley and Samuelson 1981) that, with symmetric IPV bidders, the optimal reserve price is independent of N.

is used, we limit to cases where $\underline{N} \geq 2$, and anytime $\overline{N}$ is used, we limit to cases where $\overline{N} \geq 2$.

B.4 Analysis of Assumptions Involving N Independence

B.4.1 Test of Bidders' Values Independent of N. Our estimation of the underlying distribution of buyer values F_B relies on the assumption that bidders' values are independent of N ($A2$). This assumption is common in empirical auction studies, and the assumption has several names: Athey and Haile (2007) referred to the assumption as one of *exogenous participation*, while Aradillas-López et al. (2016) (AGQ) referred to it as *valuations being independent of* N. We adopt the latter terminology.

AGQ demonstrated the following (their Proposition 2): For any $k \leq n$ and any v, let $\psi_{k:n}(v) \equiv \frac{n!}{(n-k)!(k-1)!} \int_0^v t^{k-1}(1-t)^{n-k}dt$. In settings where bidders have private value $\tilde{B}_i = W + B_i$, where B_i are IPV and W is unobserved game-level heterogeneity, then, for any v and any $n < n'$, values being independent of N implies

$$\psi^{-1}_{n-1:n}(F_{\tilde{B}^{n-1:n}}(v)) \leq \psi^{-1}_{n'-1:n'}(F_{\tilde{B}^{n'-1:n'}}(v)) \tag{34}$$

where, for any n, $F_{\tilde{B}^{n-1:n}}$ represents the distribution of the auction price *including the unobservable heterogeneity* when n bidders are present. If values are independent of N, (34) would hold with equality if there were no unobserved heterogeneity (i.e., $W = 0$); in that case, then both the left- and right-hand sides of (34) would yield the underlying distribution of buyer values, F_B. With unobserved heterogeneity ($W \neq 0$), the inequality in (34) holds. In words, the AGQ result shows that, if valuations are independent of N, using noisy auction prices ($\tilde{P}^A$) and applying the second-order statistic inversion to a sample of n-bidder auctions will lead to an implied underlying distribution that stochastically dominates the distribution implied by repeating the exercise with data from n'-bidder auctions (for $n < n'$). Importantly, in addition to deriving this result, AGQ demonstrated that the canonical models of endogenous entry analyzed in Samuelson (1985) and Levin and Smith (1994) lead to violations of (34), and thus the condition is testable against these models of endogenous N.

We apply this result by evaluating the inequality in (34), where we perform the auction price distribution inversion at different values of $\underline{N}$ observed in our bid log subsample. For this exercise, we treat realizations of $\underline{N}$ as realizations n of the number of bidders. As described in Section 5, the lower bound $\underline{N}$ varies auction-by-auction; for the sake of this exercise, it captures variation in the degree of competition in different auctions.

Figure A.3: Test of Independence between N and Buyer Values

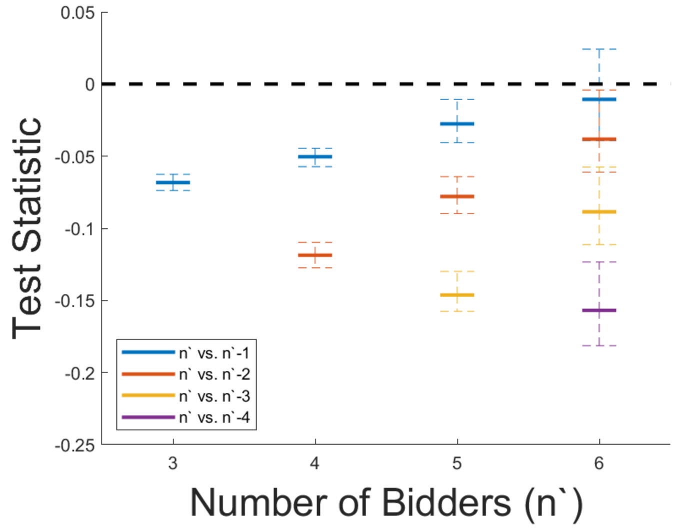

Notes: Figure shows the value of test statistic $\Psi(n, n')$ from (35) at different values of n and n', with two-sided 90% confidence intervals constructed from 200 bootstrap replications shown with dashed lines.

We construct the following statistic that summarizes violations of (34) for $n < n'$:

$$\Psi(n, n') \equiv \int [\psi^{-1}_{n-1:n}(F_{\tilde{B}^{n-1:n}}(v)) - \psi^{-1}_{n'-1:n'}(F_{\tilde{B}^{n'-1:n'}}(v))] dF_{\tilde{B}^{n-1:n}}(v) dv \tag{35}$$

(35) computes the amount by which (34) is violated and then integrates this amount against the density of the second order statistic of bids to incorporate the fact that estimates of the underlying value distribution will be more precise where we have more data.

We compute this statistic using all values of $n \geq 2$ for which we have at least 100 auctions, which corresponds to auctions with 2 to 6 bidders. Figure A.3 displays the estimate of Ψ for all of these values of n' and for various gaps between n and n'. In blue, the statistic is constructed by comparing values of Ψ for n' and for $n = n' - 1$; in red, $n = n' - 2$; in yellow, $n = n' - 3$; and, in purple, $n = n' - 4$. The solid point shows the test statistic at each value of n'. The dashed lines show two-sided 90% confidence intervals constructed from 200 bootstrap replications. Our test of interest is a one-sided test — the null hypothesis is that $\Psi(n, n') < 0$ — and thus the lower confidence bound is of particular interest. The lower 90% confidence bound from a two-sided confidence interval corresponds to the 95% confidence bound from a one-sided test of the null $\Psi(n, n')$.

Figure A.3 demonstrates that, at all values of n and n', the point estimate and the lower confidence bound of the statistic are negative. We thus reject the null that the AGQ inequality is violated, a finding consistent with values being independent of N and inconsistent with traditional models of endogenous N.[74]

B.4.2 Relationship Between N and Other Variables. Assumptions $A2$–$A3$ impose that N is independent of (Y, W, V_S). Maintaining the other assumptions of the model, the correlation between N and $Y\hat{\beta}$ should therefore be zero if N is independent of Y. To test this, we use the lower bound on the number of bidders, $\underline{N}$, as our proxy for N. We regress observations of $Y\hat{\beta}$ on dummies for $\underline{N}$, with the omitted dummy being the one corresponding to $\underline{N} = 3$, which is the realization of our lower bound in 84% of observations.[75] Figure A.4.A shows the estimated coefficients, with 95% confidence intervals surrounding each estimate.[76] We find that all of the estimates are statistically different from the omitted category (auctions with $\underline{N} = 3$). Interestingly, the average $Y\hat{\beta}$ at a given value of $\underline{N}$ is non-monotonic in $\underline{N}$: at both $\underline{N} = 2$ and $\underline{N} = 4$ the point estimate is greater than at $\underline{N} = 3$, suggesting that it is not the case that an increase in $\underline{N}$ always corresponds to an increase in the predicted market value $(Y\hat{\beta})$ of a car.

We repeat this exercise with the regression outcome being the estimate of the noisy reserve price $\tilde{R}$. Independence of $\underline{N}$ and $\tilde{R}$ should hold if $\underline{N}$ is independent of (V_S, W), maintaining the other assumptions of the model. The estimated coefficients on the dummies for values of $\underline{N}$ are shown in Figure A.4.B. We find that some (although not all) values of $\underline{N}$ have an average estimated $\tilde{R}$ that is statistically significantly different from that of $\underline{N} = 3$ (the omitted category). However, the relationship is quite scattered, with some point estimates above zero and some below, and the differences in point estimates are smaller in magnitude than in Figure A.4.A.

The results in Figure A.4 suggest that the assumption of N being independent of (V_S, W) may not be overly strong in our data, but independence of Y and N may be. We maintain the assumption despite these findings for two reasons. First, our results in Figure A.3 lead us to fail to reject N being independent of buyer values, and our results in Appendix B.3

[74] A similar (but abbreviated) analysis appears in Appendix D.3.2 of Larsen (2021).

[75] We omit from this analysis a realization of $\underline{N}$ that only occurs once in the data ($\underline{N} = 12$).

[76] These confidence intervals are constructed using homoskedastic standard errors, which is consistent with the independence assumptions of the model.

Figure A.4: Correlation Between N and Other Variables

(A) Correlation of $\underline{N}$ and $Y\hat{\beta}$ (B) Correlation of $\underline{N}$ and $\tilde{R}$

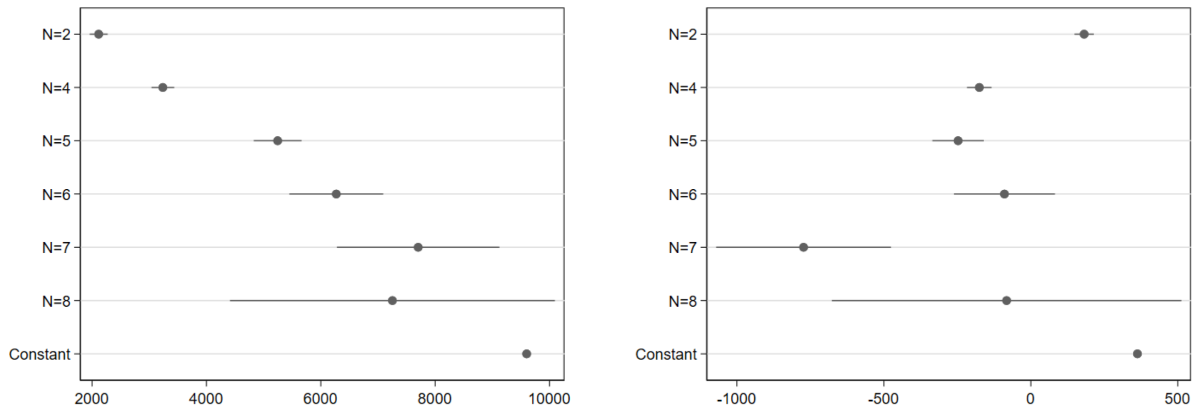

Notes: Panel A shows regression estimates from a regression of $Y\hat{\beta}$ (the predicted market value from Step 1 of our estimation) on dummies for $\underline{N}$ (the lower bound on the number of bidders at the auction). On the vertical axis, N represents $\underline{N}$. $N = 3$ is the omitted category. Confidence intervals constructed under homoskedasticity surround each estimate. Panel B repeats this exercise with the outcome being $\tilde{R}$. Units on the horizontal axis are dollars.

offer reassurance that variation in N has only minimal effects on our estimates of $F_{B^{(1)}}$, the only channel through which N enters our estimates of bargaining power directly. The second reason is one of practicality: allowing for correlation between N and $Y\hat{\beta}$ would require conditioning nearly all of our estimation steps on this index, which would be computationally quite challenging.

B.5 Local Linear Regressions and Spline Steps

We next examine the local linear regression and spline estimation steps. In Figure A.5, 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 of $P^0_S(r)$ and $T^0_S(r)$. 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 $P^0_S(r)$ function; thus, correcting for unobserved heterogeneity yields an estimate of $P^0_S(r)$ that is steeper than $\tilde{P}_S(\tilde{r})$. This can be seen in panel A by comparing the $P^0_S(r)$ 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 the estimate of $\tilde{T}_S(\tilde{r})$ the estimate of $E[W\Delta P_S \mid \tilde{r}]$, which represents the expected value of the unobserved

Figure A.5: Removing Unobserved Heterogeneity from Allocation/Transfer Functions

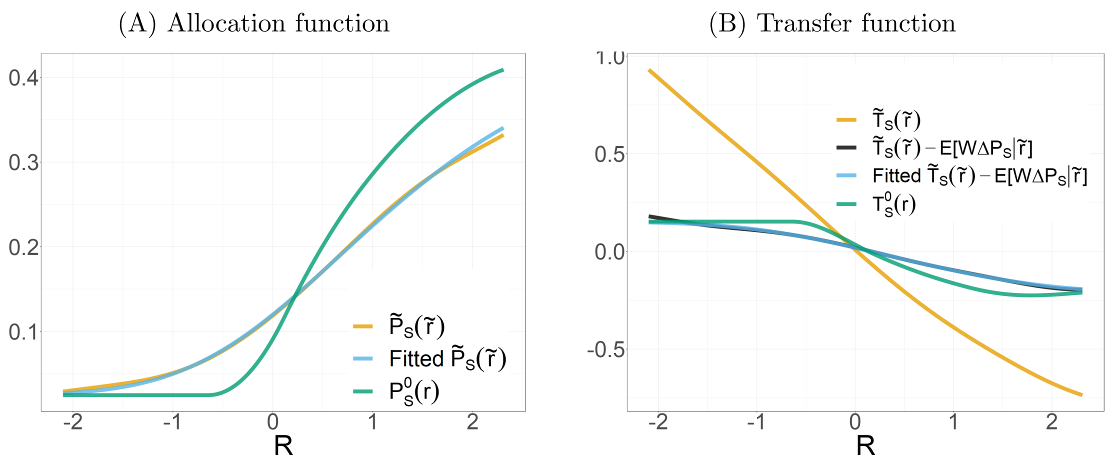

Notes: Figure displays heterogeneity correction for allocation function (Panel A) and transfer function (Panel B). Yellow lines display the original uncorrected estimates of $\tilde{P}_S(\tilde{r})$ and $\tilde{T}_S(\tilde{r})$ from local linear regressions, and green lines display final, corrected estimates, $\hat{P}^0_S(r)$ and $\hat{T}^0_S(r)$. 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, $\tilde{T}_S(\tilde{r}) - E[W\Delta P_S \mid \tilde{r}]$. 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.

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 $T^0_S(r)$ being higher. In panel B, comparing the $\tilde{T}_S(\tilde{r})$ line to the $\tilde{T}_S(\tilde{r}) - E[W\Delta P_S \mid \tilde{r}]$ line shows that this correction makes the slope of the expected transfer function significantly less negative. Secondly, the estimate of $T^0_S(r)$ is essentially a noise-corrected version of the estimated $\tilde{T}_S(\tilde{r}) - E[W\Delta P_S \mid \tilde{r}]$, and thus the slope and convexity of the estimated $T^0_S(r)$ are both larger in absolute value than the noisy version. The net effect is that the estimated $T^0_S(r)$ is much less negatively sloped—and somewhat more convex—than the original nonparametric estimate of $\tilde{T}_S(\tilde{r})$. In each panel, the blue line displays the fitted estimates, constructed by the convolution of the estimated allocation or transfer function against F_W; in each case, the estimate aligns closely with the local linear estimates. Quantitatively, the RMSE of the fitted $\tilde{P}_S(\tilde{r})$ function is 0.0021, and the RMSE of the $\tilde{T}_S(\tilde{r}) - E[W\Delta P_S \mid \tilde{r}]$ function is \$4.08.

B.6 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 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 ignoring unobserved heterogeneity, treating the estimates of the expected allocation and transfer functions, $\tilde{P}_S(\tilde{r}), \tilde{T}_S(\tilde{r})$, 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 (without enforcing IR). Figure A.6 shows results analogous to those from Figures 3.A, 4.A, and 4.B, comparing estimates accounting for and ignoring unobserved heterogeneity.

Figure A.6.A shows how the unobserved heterogeneity correction affects the seller menu. The menu is much steeper without the unobserved heterogeneity correction. The intuition behind this result is as follows. If unobserved heterogeneity is present, a positive relationship between sale prices and reserve prices is driven by both R and W, but without incorporating the unobserved heterogeneity correction, our menu approach misattributes all of this relationship to R, implying that sellers could achieve higher prices by setting higher reserve prices. Through the lens of the no-heterogeneity-correction model, cases where sellers do not choose higher reserve prices are then interpreted as being cases of sellers necessarily having very low values, equivalent to a steeper menu.

Panel B of Figure A.6 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. Without the unobserved heterogeneity correction, we estimate that a nontrivial fraction of sellers have values from -\$30,000 to -\$10,000 below the estimated market value of the car (unrealistically far below the market value). Thus, the unobserved heterogeneity correction appears to be important to account for in order to obtain reasonable estimates of sellers' values.

Figure A.6: Estimation With and Without Unobserved Heterogeneity Correction

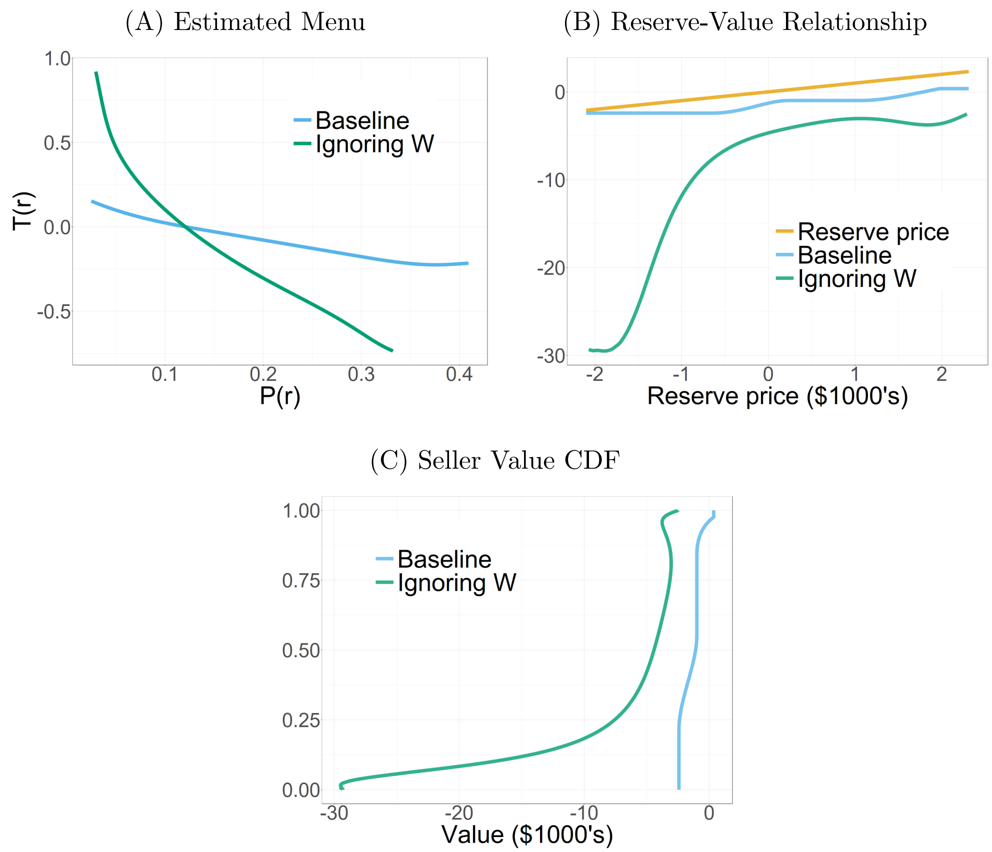

Notes: Estimates of the menu, reserve-value mapping, and F_S, with and without the unobserved heterogeneity correction. Panel A shows the baseline menu estimate (blue) and the menu estimate using the local polynomial estimates of allocation and transfer functions, $\tilde{P}_S(\tilde{r}), \tilde{T}_S(\tilde{r})$,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

B.7 OEM vs. Non-OEM Sales

In this section we include additional analysis of OEM vs. non-OEM sales. In our estimation using these subsamples, we find a strong fit in terms of RMSE, just as in the full sample. The RMSE from the spline estimation for the probability function (estimation Step 5) is 0.0009 for OEM sales and 0.0013 for non-OEM sales. The RMSE for the spline estimation of the transfer function (Step 6) is \$14.71 for OEM sales and \$17.61 for non-OEM sales. The convexity constraint binds more in these subsamples than in the overall sample — especially for OEM sellers — implying that, through the lens of our menu model, sellers behave as though 90.11% of them have the same value in the OEM sample; the corresponding number is 45.04% in the non-OEM sample and 28.90% in the full sample (as reported in Section 5.2). IR constraints are binding for only a small fraction of sellers in the OEM sample (3.54%) and for zero sellers in the non-OEM sample.

We now explore what variation in the data leads to our conclusions in Section 6.4 regarding bargaining power in transactions involving OEMs vs. non-OEMs. Recall from Section 4 that all of our key identification arguments come from the joint distribution of four objects that are observable in our data: an indicator for the whether trade occurs ($\mathcal{A}$ below) and three noisy price variables ("noisy" meaning they are contaminated with unobserved heterogeneity): auction prices ($\tilde{p}^A$), secret reserve prices ($\tilde{r}$), and final prices ($\widetilde{price}$). Therefore, any differences we find between OEM and non-OEM bargaining power must arise from differences in these variables between the two subsamples.

In Table A.1, we first consider the variation in the data that leads to the finding that OEMs have more unconstrained bargaining power than non-OEMs. Columns 1–4 of panel A show results from separate regressions of each of these four variables on an indicator for whether the seller was an OEM.[77] Columns 1 and 2 show that OEMs achieve a higher trade probability than non-OEMs (by 9.5 percentage points), with no significant reduction in final prices: an increase in sale probability would typically be associated with a decrease in sale price, but here we observe OEMs receiving a price that is \$40.86 higher on average. Together, this leads to higher revenue for OEMs, pushing their surplus closer to the revenue

[77] For any regression involving $\widetilde{price}$ as an outcome, we use only observations in which trade occurs; for all others, we use the full sample.

Table A.1: Variation in Data Underlying OEM Bargaining Power

A.	(1)	(2)	(3)	(4)	(5)	(6)
	$\mathcal{A}$	$\widetilde{price}$	$\tilde{r}$	$\tilde{p}^A$	$\underline{N}$	$\overline{N}$
OEM	0.0954***	40.86***	-80.58***	150.9***	0.0109***	16.18***
	(0.00255)	(8.809)	(8.647)	(8.765)	(0.00371)	(0.103)
Constant	0.829***	0.184	136.3***	-146.6***	3.000***	23.08***
	(0.00136)	(4.877)	(4.607)	(4.670)	(0.00197)	(0.0556)

B.	(1)	(2)	(3)	(4)	(5)	(6)
	$\mathcal{A}$	$\widetilde{price}$	$\tilde{r}$	$\widetilde{price} - \tilde{p}^A$	$\widetilde{price} - \tilde{p}^A$	$\widetilde{price}$
OEM	0.0716***	-8.201***	-134.4***	-8.495***	-8.201***	
	(0.00234)	(0.615)	(8.075)	(0.620)	(0.615)	
$\tilde{p}^A$	0.000102***	0.990***	0.357***		-0.00952***	0.993***
	(0.000000943)	(0.000275)	(0.00304)		(0.000275)	(0.000249)
$\tilde{r}$	-0.000106***	0.00612***			0.00612***	
	(0.000000956)	(0.000274)			(0.000274)	
Constant	0.859***	8.467***	188.6***	8.495***	8.467***	5.939***
	(0.00126)	(0.340)	(4.318)	(0.343)	(0.340)	(0.285)

Notes: Results from regressions of various outcomes (shown above each column) on a dummy for the seller being an OEM and other controls. $\underline{N}$ is the lower bound on the number of bidders and $\overline{N}$ is the upper bound (both of these bounds vary at the auction level, and are described in Appendix B.3). Homoskedastic standard errors (consistent with the model's independence assumptions) are in parentheses. Prices are in units of dollars.

under $\mathcal{M}_S^{1^{st}}$ or $\mathcal{M}_S^{2^{nd}}$.

Table A.1 also reveals the underlying sources for these findings: OEMs set lower reserve prices (column 3), consistent with them having lower private values for the cars they sell. OEMs having lower private values and yet receiving similar final prices that are higher than non-OEMs — and a higher trade probability — implies that the seller surplus for OEM sellers is larger (and Table 3 shows that they are garnering more of that surplus). The higher trade probability is also driven by OEMs facing higher auction prices (column 4).[78] Consistent

[78]Note that our empirical analysis the OEM and non-OEM subsamples separately can be considered a relaxation of the assumptions ($A2$–$A3$) that other random variables in the model are independent of Y, the observable heterogeneity vector, because this vector includes fixed effects for large sellers such as OEMs. It can also be considered a relaxation of the assumption that all data are generated by the same equilibrium ($A11$). The OEM status of a seller can be considered a variable that is correlated with a seller's private value, V_S. Under this interpretation, our finding in Table A.1.A that noisy reserve prices differ between OEMs and non-OEMs is consistent with these independence and single-equilibrium assumptions holding in the full

with auction prices being higher for OEM sellers, column 6 shows that, for OEM sellers, the upper bound on the number of bidders ($\overline{N}$) is higher by 16.18 bidders and the lower bound ($\underline{N}$, column 5) is higher by 0.011 bidders. The combination of these features, through the lens of our model, implies higher bargaining power for OEMs absent the competition constraint.

We now consider the variation in the data that leads us to infer that, after accounting for the competition constraint, OEM bargaining power is low, and is not better than that of non-OEMs. The key result in this regard is found in column 4 of panel B, which shows that OEMs are no better than non-OEMs at pushing the final price above the auction price, and they may even be worse, as the point estimate is negative (albeit quite small, suggesting a difference of only \$8.50 for OEM vs. non-OEM sellers). Column 5 shows that this result holds even after controlling for the (noisy) reserve and auction price. These data features drive our finding that OEMs have little or no competition-constrained bargaining power.

Other results in panel B confirm points from panel A. The higher trade probability experienced by OEMs is evident even after controlling for noisy auction prices and reserve prices (column 1). After including these controls, we find that prices are slightly lower for OEMs than non-OEMs (by \$8.20). Column 3 shows that OEMs choose lower reserve prices even conditional on the (noisy) auction price; as explained in Appendix B.2, the noisy auction price can be considered a proxy for unobserved heterogeneity. Other relationships in panel B move in sensible directions: higher auction prices and lower reserve prices increase trade (column 1); noisy auction prices and noisy reserve prices are correlated, as would be expected in the presence of unobserved heterogeneity (column 3); and final prices are higher when auction prices increase, as implied by Assumption $A4$ (column 2). This last result is also shown in column 6 of panel A, unconditional on OEM status.

sample, as is the finding that agreements rates are higher, as this can follow from lower reserve prices, all else equal. However, the finding that noisy auction prices and the number of bidders differs across the two subsamples suggests that the distribution of buyer values or $\Pr(N = n)$ may differ in OEM vs. non-OEM sales, and hence the assumption of independence of Y and buyer values, or the assumption of independence of Y and N, is at best an approximation in the full sample.

Figure A.7: Dealer Sellers: Sale Prob., Auction Price, and Final Price by Reserve Price

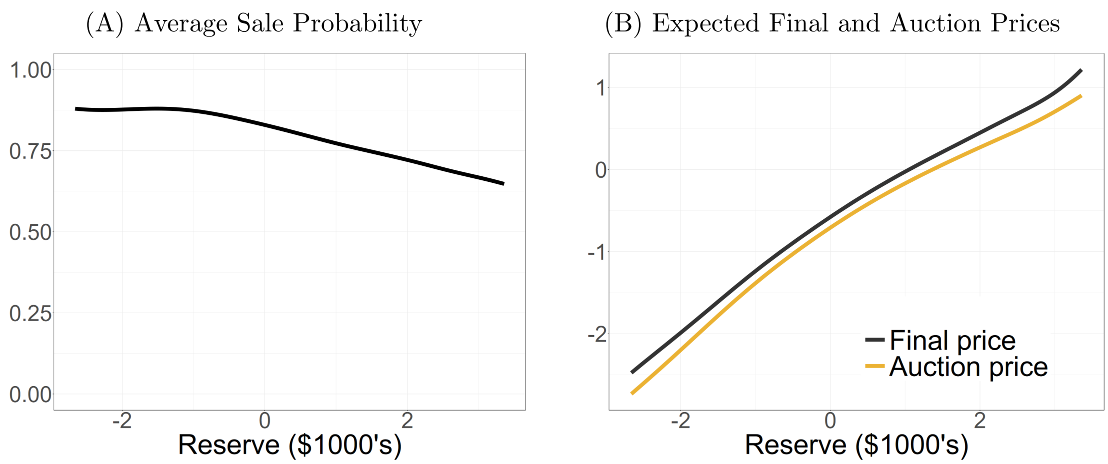

Notes: Figure shows results as in Figure A.1 but using observations in which sellers are used-car dealers.

B.8 Dealers vs. Fleet/Lease Sellers

Our study focuses on cars sold by large fleet/lease sellers because we are particularly interested in bargaining power between dealers and OEMs or other large institutional sellers. In Figure A.7, we replicate Figure A.1 but using a sample of cars sold by dealer sellers rather than fleet/lease sellers. Recall that Figure A.1 provided some descriptive evidence that, in our main sample, the menu approach is indeed appropriate in this setting, i.e., different reserve prices yield different payoffs for sellers, and hence can serve to help separate seller types as our method requires. In Figure A.7, however, we observe a relatively flat trade probability and an expected final price that represents roughly a constant markup over the auction price, regardless of the seller's reserve price. These results suggest that it would be challenging to use our menu approach to identify seller values in this dealer sellers sample, as the approach requires that sellers face a trade-off between trade probabilities and transfers at different reserve prices — a trade-off that does not jump out from Figure A.7, unlike Figure A.1.

Consistent with this evidence, when we estimate our model on the dealer sellers sample, we find that, prior to enforcing the IR constraint, the constraint is violated for 56.36% of observations, unlike in the fleet/lease sample, where only a small fraction of observations (2.44%) require the IR-enforcement step of our estimation. This suggests that it may be challenging to infer seller values based on an assumption of optimally chosen secret reserve prices in the dealer sellers sample. Larsen (2021) took a different approach, 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 large fleet/lease sellers in this paper.

Appendix References

Aradillas-López, A., Gandhi, A., and Quint, D. (2016). A simple test for moment inequality models with an application to English auctions. *J. Econometrics*, 194:96–115.

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

Levin, D. and Smith, J. L. (1994). Competitive bidding with entry costs. *AER*, 84(3):585–599.

Riley, J. G. and Samuelson, W. F. (1981). Optimal auctions. *AER*, 71(3):381–392.

Samuelson, W. F. (1985). Competitive bidding with entry costs. *EL*, 17(1–2):53–57.

Song, U. (2004). Nonparametric estimation of an eBay auction model with an unknown number of bidders. *WP, UBC*.