Intermediaries in Bargaining: Evidence from Business-to-Business Used-Car Inventory Negotiations*

Bradley J. Larsen, Carol Hengheng Lu, and Anthony Lee Zhang†

November 3, 2025

Abstract

We analyze 75,000 alternating-offer, business-to-business negotiations in the wholesale used-car market, each mediated (over the phone) by a third-party. We find that it matters who intermediates: high-performing mediators are 18.07 percentage points more likely to close a deal than low performers, and these differences are not due to statistical noise. Better mediators achieve trades quickly, and mediator skill improves with experience. We estimate buyer and seller value distributions using a structural model of two-sided incomplete-information in which mediators of different skill levels correspond to different equilibria. High-skilled mediators improve efficiency and overcome some of the inefficiency inherent in incomplete-information settings.

JEL Codes: C7, D8, L1, L81

Keywords: Bargaining, intermediaries, mediation, incomplete information, alternating offers, inefficiency

*We thank Rebecca Li, Jessie Liao, Wendy Yin, Zichen Zhao, and Tianyi Zhang for outstanding research assistance, as well as Jack Fanning, Matt Grennan, Nagore Iriberri, Pranav Jindal, Jingxi Li, Rachael Meager, Bruno Pellegrino, Joris Pinske, Hrishikesh Relekar, Yijing Ren, Renan Yoshida, Xingtan Zhang, and seminar and conference participants at the University of California Berkeley Haas, University of Toronto, Shanghai Jiao Tong University, University of Chicago, Queens University, Penn State University, Stanford-Berkeley IO Fest, the North America Winter Meetings of the Econometric Society, the European Winter Meetings of the Econometric Society, the International IO Conference, BEET (Bargaining: Experiments, Empirics, and Theory), the International Association of Conflict Management, and the NBER Market Design Working Group Meeting. We acknowledge support from the National Science Foundation through grants SES-1629060 and SES-1530632.

†Larsen: Washington University in St. Louis - Olin Business School and NBER; blarsen@wustl.edu. Lu: Amazon; chhlu@amazon.com. Zhang: University of Chicago Booth School of Business; anthony.zhang@chicagobooth.edu

1 Introduction

Many real-world bargaining situations — among nations, businesses, investors, consumers, or litigants — involve a third-party *intermediary* or *mediator*. These mediators are often at the center of massive transactions and are highly paid for their services: consider investment banks handling firm acquisitions or lawyers mediating pre-trial settlements. A 2011 survey of Fortune 1000 companies found that 98% of respondents used a mediator at least once in the previous three years.[1] To date, however, there is little real-world evidence on whether intermediaries make a difference for negotiation outcomes and welfare.

We analyze 75,000 business-to-business negotiations from the wholesale used-car market. The industry consists of hundreds of auction house locations nationwide that facilitate trade between manufacturers, fleet companies, banks, and used and new car dealerships, with more than $110 billion worth of cars traded annually.[2] Each car is auctioned individually in a rapid process. When the auction price (the highest bid at the auction) fails to reach the seller's secret reserve price, it frequently occurs that a *mediator* facilitates bargaining between the highest bidder and seller over the phone. The buyers, sellers, and mediators in this market are professionals; they frequently engage in negotiations, as each auction house location sells hundreds to thousands of cars on a fixed day each week. Stakes are high, especially for smaller buyers or sellers, where each transaction can make the difference for achieving preferred inventory levels or profit margins. The primary goal of the auction house company throughout this process is to maximize the trade probability (see discussions of the industry in Treece 2013, Lacetera et al. 2016).

The dataset we study is rich, containing information from six auction house locations owned by one company. Collectively, these locations sold hundreds of thousands of cars from 2006–2010. For each attempt to sell a car, the data record the auction price, the seller's secret reserve price, every action taken by each party in the negotiation process, and — critically for our study — the mediator's identity, which varies across transactions. The data also contain detailed vehicle characteristics, the timing and location of each transaction, and identities of buyers and sellers. Such data are rare in the literature — only a handful of existing studies analyze information on

[1] See Stipanowich and Lamare (2014), who also showed that the use of mediators by major companies has grown since 1997 in every category of dispute: commercial contracts, employment, consumer issues, corporate finance, environmental issues, intellectual property, personal injury, product liability, real estate, and construction.

[2] https://www.autoremarketing.com/ar/how-vehicle-volumes-sales-rates-performed-at-naaa-member-auctions-last-year/.

offers and counteroffers within real-world bargaining, and we know of only one other study containing information on both counteroffers and mediation.[3] We view this market as a rare opportunity to study bargaining intermediaries in the field.

In the data, negotiations end in agreement 62% of the time. In the remaining 38%, the buyer and seller fail to trade. Such breakdown is consistent with the presence of *two-sided incomplete information*, where the buyer and seller both have private values (i.e., unknown to the opposing party) for the car, and inconsistent with complete information.[4] Two-sided incomplete information is a primary rationale for the wholesale used-car market being an interesting place to study mediator heterogeneity and its effects on welfare. This is for two reasons. First, in such an environment, the number of equilibria of a sequential bargaining game is infinite (Ausubel et al. 2002), with equilibria known to have vastly different efficiency outcomes, ranging from no trade (Perry 1986, Ausubel and Deneckere 1992) to relatively efficient trade (Cramton 1992, Ausubel and Deneckere 1993). Thus, if one mediator can implement a different equilibrium than another, there is scope for her to increase trade probabilities and realized gains from trade. Second, with incomplete information, there is an inherent trade-off between efficiency and either party's rent extraction, and thus bargaining is not zero-sum: mediators who increase agreement can potentially also increase the size of the pie.

To analyze whether different mediators differentially affect welfare, we introduce a structural model that allows for different mediators to correspond to distinct pure-strategy Perfect Bayesian Equilibria (PBE) of a two-sided incomplete-information game. As a precursor for this structural analysis, we first document descriptive evidence on the assignment of mediators to specific negotiations and on the heterogeneity in observable outcomes across mediators. Our conversations with sales managers suggest that mediators are largely randomly assigned, although empirically we find some evidence that assignment is not entirely random and can depend on features determined before bargaining starts, such as car characteristics, buyer and seller identities, and outcomes of the pre-bargaining stages of the game (the auction). Our detailed data allow us to control for these

[3] In contemporaneous work, Dindaroglu and Ertac (2024) studied a little over 1,000 manually collected negotiations in Turkey's sheep market leading up to the Festival of Sacrifice. 7 percent of these involved a third-party mediator. Mediated negotiations had a higher trade probability and showed no difference in prices, consistent with our findings.

[4] In a complete-information model, such as the Rubinstein (1982) noncooperative model or cooperative Nash bargaining, parties only begin negotiating if agreement is known a priori to be efficient. The same is true in many incomplete-information models where the incomplete information is about features other than values, such as discount rates (as in Rubinstein 1985), or in models of reputational bargaining (Abreu and Gul 2000), where all rational types eventually trade and trade is always the efficient outcome. In contrast, with two-sided incomplete information about values, and with overlapping support for buyer and seller values, *disagreement* will sometimes be efficient, and impasse will generally occur with positive probability even when gains from trade exist (Myerson and Satterthwaite 1983).

features, and we show that, conditional on controls, variation in mediator assignment is indeed largely random. We further show that our key findings are similar with or without these controls.

We document heterogeneity in mediators' average trade probability through various regressions of an agreement indicator on mediator fixed effects, following the analysis of human auctioneer heterogeneity in Lacetera et al. (2016). We observe significant dispersion: the 75th percentile mediator trades 32.9 percentage points more often than the 25th percentile mediator. Large differences persist as we add rigorous controls for features of the negotiation. In our most saturated specification, the 75th–25th-percentile gap is 18.07 percentage points. Some variation in outcomes would occur in any finite data due to statistical noise even if underlying population probabilities were equal across mediators. Through a bootstrap simulation exercise, we show that mediator heterogeneity is a real phenomenon, wider than can be explained by statistical error. While mediators have large effects on trade probability, their effects on prices are indistinguishable from zero.[5] In this light, we refer to mediators with a higher trade probability as having higher *skill* or higher *performance*.

Having established heterogeneity in mediators' trade probability, we then examine how mediators differ in *when* they come to agreement and how outcomes differ by mediators' experience. Negotiations handled by effective mediators have a higher likelihood of ending in the first bargaining round, rather than after substantial back-and-forth action between buyers and sellers. Effective mediators stand out in their ability to achieve trade in difficult negotiations, where the auction price is not already close to the seller's secret reserve price. More experienced mediators — those who have been employed by the company longer — are more effective on average at reaching agreement, and within-mediator performance improves to some extent over time. Importantly, however, there is substantial residual variation in mediator ability that is not explained by experience.

These results are consistent with any number of possible channels through which mediators might influence outcomes. The lens through which we view this for our structural exercise is that different mediators implement different equilibria, some with higher probabilities of trade. The model is closely related to that of Larsen (2021). We first identify the buyer value distribution using an order statistics argument that relates buyer values to auction prices. We bound the seller value distribution using revealed preference arguments from the seller's decision to accept or reject the auction price. Finally, for different mediators, we identify the function determining whether

[5] As highlighted above, a lack of price effects is consistent with statements of industry participants, who suggest that the primary aim of the auction house is to facilitate trade independent of price. In our bootstrap simulation exercise, we confirm that any differences in average prices across mediators can be explained by statistical noise.

a given buyer and seller pair trade. This final step relies on a revelation-principle argument: any equilibrium has a corresponding *direct mechanism*. The direct mechanism for a given mediator skill level can be identified from trade probabilities in the data conditional on auction prices and reserve prices. Our identification and estimation also incorporate unobserved game-level heterogeneity — information that agents observe but that is unobservable to the econometrician.

With these estimated direct mechanisms and distributions of buyer and seller values, we compute surplus (the gains from trade) realized in the trade-maximizing (first-best efficient) mechanism and the second-best efficient mechanism, as well as the real-world surplus achieved by mediators of different skill levels.[6] We find that mediators who achieve a higher trade probability also achieve greater gains from trade. This implies that higher-skilled mediators are not only capturing additional low-surplus trades (i.e., cases where the buyer values the car only slightly more than the seller). Rather, these higher trade-volume mediators increase the total realized *gains from trade*, creating significant value. As highlighted above, this reflects the feature that, under two-sided incomplete information about values, the game is not zero-sum.

2 Related Literature

Our study examines mediator heterogeneity. A related theoretical and empirical literature compares mediated bargaining to unmediated bargaining. On the theory side, Jarque et al. (2003) (JPS) and Fanning (2021) studied continuous-time, war-of-attrition models where mediators can improve outcomes by speaking privately to agents and only revealing the opposing parties' offers once they agree. The models differ in how they arrive at this result: JPS studied a double auction with two-sided incomplete information about values. Fanning (2021) considered reputation, with two-sided incomplete information about whether an agent is rational or not. In the latter model, mediators introduce randomness in when they reveal offers compatible with agreement.[7] Our model differs from Fanning (2021) in that we consider incomplete information about values, as in JPS.[8] We differ

[6] The first-best mechanism maximizes trade gains as well as probability. With incentive constraints, the second-best mechanism that maximizes trade gains does not necessarily maximize probability. We choose to evaluate both the first-best and surplus-maximizing second-best.

[7] The model of Copic and Ponsati (2008) is built on JPS, while Fanning (2023) is built on Fanning (2021) to derive the optimal mediation protocol in reputational bargaining. Basak (2024) also studied information in reputational bargaining. Fanning (2021) offered a thorough literature review that we build on.

[8] Gottardi and Mezzetti (2024) modeled two-sided incomplete information about values, where mediators know more about agents' values than do agents. Mediators restrict information flow to agents to allow gradual learning and improve efficiency. In our setting, mediators do not know more than agents about the car; Table 3 and Figure 7 of Lacetera et al. (2016) presented survey evidence consistent with this for auctioneers, whose level of knowledge is

from JPS by considering alternating offers and discrete time, realistic features of our setting.[9]

Theory also considers mediation in one-shot bargaining games. In Goltsman et al. (2009), a mediator can improve outcomes over unmediated bargaining by filtering information and adding noise, akin to Fanning (2021).[10] In Kydd (2003), mediators must be biased toward one party or the other in order to send credible messages.[11] Hörner et al. (2015) modeled international conflict, finding that mediation can be effective when mediators do not fully reveal when a country is weak. Li and Zhang (2024) demonstrated theoretically that more informed mediators can improve efficiency.[12]

In our study, we aim for a structural model that is patterned closely after the real-world game and that allows for the wide variety of outcomes we observe while still being empirically tractable. This motivates our choice of a model with the following traits: two-sided incomplete information about values, continuous value distributions, discrete time, and both parties being allowed to make multiple offers. None of the above models of mediation in bargaining fit these requirements. We therefore adapt the two-sided incomplete-information model of Larsen (2021) to allow for multiple mediator types.[13] Our model does not incorporate some potential roles of mediators from the theoretical models above, such as mediators having different information than agents or being able to elicit private information to selectively (and potentially randomly) disclose. Instead, in our model, mediator types differ in which of infinitely many equilibria they implement, leading to different agreement rates and trade gains.

A higher-skilled mediator in our framework can be viewed as one that helps agents play an equilibrium closer to the theoretical second-best described in Myerson and Satterthwaite (1983) and Williams (1987). For example, a good mediator may be better at withholding trades between players

similar to auction house mediators in this market.

[9]JPS highlighted that modeling their setting in continuous time greatly facilitates the theoretical analysis.

[10]Goltsman et al. (2009), as well as Fanning (2023), found that arbitration is generally more effective than mediation. Like mediation, arbitration is considered a form of alternative dispute resolution (ADR), but arbitrators can enforce (rather than only recommend) an outcome. Unlike mediation, arbitration usage by firms has declined in recent decades (Stipanowich and Lamare 2014). Kong et al. (2025) presented a structural analysis of arbitration.

[11]Kim (2024) showed theoretically that a mediator who maximizes trade (not revenue) can send more credible messages. The author finds supporting evidence in Korean real estate.

[12]Glode and Opp (2016) considered an intermediary who buys from a seller and resells to a buyer, finding that efficiency improves if the intermediary is more informed than one of the negotiators.

[13]The broader bargaining theory literature also tends to focus on special cases of the game we study, such as take-it-or-leave-it offers or one-sided incomplete information. Sequential bargaining with two-sided incomplete information about values is known to yield infinitely many, qualitatively different equilibria, with no complete characterization (Ausubel et al. 2002), posing challenges for theoretical and empirical work. Equilibrium multiplicity is driven by off-equilibrium beliefs, which can sustain a large range of on-path behavior. A feature that supports tractability in our setting is that negotiation occurs *after* an auction, so auction actions (e.g., auction prices and secret reserve prices) help us identify primitives while remaining agnostic about the nature of equilibria in the post-auction bargaining.

when necessary — even when she knows the buyer values the good more than the seller — to keep agents' reporting incentive compatible.[14] In theory work building on Myerson and Satterthwaite (1983), Eilat and Pauzner (2021) showed that, if mediators cannot commit to the trade-withholding necessary to maintain incentive compatibility, efficiency declines.[15] How close real-world bargaining gets to the Myerson-Satterthwaite-Williams frontier, and whether some mediators get closer than others, is an empirical question.

Empirical evidence on mediation effectiveness is mixed. Several laboratory experiments have studied mediation: in Bazerman et al. (1992), mediation decreases the trade probability and increases price, depending on mediators' incentives and knowledge; Casella et al. (2020) found that mediation does not increase agreement; Eisenkopf and Bachtiger (2013) found the opposite, but only when mediators can punish uncooperative agents. In international conflict, Dixon (1996) found that mediation increases the likelihood of peaceful resolution, while Fey and Ramsay (2010) found no effect.[16] In law, McEwen and Maiman (1981) found that mediation increases compliance and satisfaction rates in small claims suits, and Emery et al. (1991) showed positive effects in child custody cases in which parents were randomly assigned to mediated vs. unmediated negotiations.

An important distinction of our study from previous empirical work — as well as previous theoretical work — is that we study *differences across mediators* rather than comparing mediated to unmediated bargaining. To our knowledge, in the theoretical or empirical studies highlighted above, all mediation is considered equal — there is no scope for any one mediator to *do* something better or inherently *be* better than another, no notion of *mediator types* or *mediator skill*.[17] This contrasts starkly with the way mediators are discussed in popular writing and negotiation training, where discussions abound of the existence of high-quality vs. low-quality mediators.[18] Our study

[14]Myerson and Satterthwaite (1983) considered the second-best efficient bilateral trade mechanism, maximizing the gains from trade, with equal weights on buyer and seller surplus, under the constraints of individual rationality, incentive compatibility, and ex-ante budget balance. Williams (1987) extended this analysis to the full Pareto efficient frontier with unequal welfare weights.

[15]Recent theoretical work (Saran 2011) suggests that efficiency in the Myerson and Satterthwaite (1983) framework can also be affected — indeed, improved — by the presence of naive traders who do not strategically shade their actions relative to their true values.

[16]Bercovitch and Jackson (2001) found that mediation is preferred in international conflicts (over unmediated bargaining) in cases of outcome uncertainty or unequal bargaining power.

[17]Theory work by Kim (2017) sheds light on one reason why empirical analysis of mediator *heterogeneity* can be challenging: to avoid being revealed as a weak negotiator, both parties favor a mediator who is best for strong negotiators. This leads to complicated selection in settings where a mediator is chosen by the negotiators. In our setting, negotiators do not choose the mediator.

[18]See, for example, the training material offered by the Harvard Law School Program on Negotiation at https://www.pon.harvard.edu/freemium/mediation-secrets-for-better-business-negotiations-top-techniques-from-mediation-training-experts/. Evidence in Stipanowich and Lamare (2014) also pointed to the existence of vari-

attempts to fill this gap by analyzing *mediator heterogeneity*.

Our study also relates to structural empirical work studying bargaining under two-sided incomplete information (Keniston 2011, Li and Liu 2015, Larsen 2021, Larsen and Zhang 2026, Freyberger and Larsen 2025).[19] Our work contributes by analyzing the impact of intermediaries.[20] The data we use overlap to some extent with that of Larsen (2021) and Larsen and Zhang (2026), although those studies do not exploit mediator identities. Larsen (2021) provided a structural welfare analysis of the real-world average bargaining efficiency relative to theoretical benchmarks, and Larsen and Zhang (2026) analyzed the division of surplus in this context. Both studies leave unanswered the question of *what explains* variation in outcomes across negotiating pairs. Our paper takes a first step in this direction. Our model builds on Larsen (2021) to identify buyer and seller values, but differs in allowing the *direct mechanism* to depend on the mediator.

Our hope is that insights from this paper can inform several audiences. For practitioners, the results suggest that mediator heterogeneity can have sizeable effects on outcomes, much like takeaways from the teacher value-added literature, and hence interventions to improve mediator skill may have positive effects on trade and welfare.[21] For bargaining theory, we hope our findings can motivate models incorporating this heterogeneity, rather than equal treatment of all mediation. Finally, for empirical bargaining research, which typically models negotiated prices as arising from some form of Nash bargaining, where mediators can play no role, our finding that mediators play a large role may offer insights into the validity of such abstractions and the importance of incomplete information more broadly.

3 Institutional Background and Data

The wholesale used-car auction industry is the backbone of the supply-side of the used-car market, both in the U.S. and much of the world. Millions of used cars arrive each year to used-car dealerships as trade-in vehicles and then are never sold at those dealerships, but are instead brought by the

ation in quality: responding Fortune 1000 firms prefer mediators whom they know from previous experience or who come from private ADR firms, and only 37.7% of respondents rank mediators they work with as "very qualified."

[19] Empirical bargaining studies of one-sided incomplete information include Silveira (2017) and Ambrus et al. (2018).

[20] Biglaiser et al. (2020) offered a model and empirical evidence of how intermediaries can help reduce information asymmetries in retail used-car markets.

[21] Other empirical studies have documented heterogeneous outcomes in sales situations in the field, such as Barwick and Pathak (2015), Robles-Garcia (2020), and Gilbukh and Goldsmith-Pinkham (2024), studying real estate agents/brokers; Lacetera et al. (2016), studying auctioneers in wholesale used-car markets; Bruno et al. (2018), studying art auctions; and Jindal and Newberry (2022), studying heterogeneity across sellers in large-scale appliances. Relative to much of this literature, the heterogeneity we find is quantitatively large.

dealerships (referred to as dealer sellers) to a wholesale used-car auction house location, where the inventory is sold to other car dealerships. Inventory at used-car auction houses also comes from large sellers (referred to as fleet/lease sellers), including fleet companies (such as Wheels, selling old fleet cars), rental car companies (such as Hertz), banks (such as Bank of America, selling repossessed cars), or manufacturers (such as Ford or GM, selling off-lease or lease-buy-back vehicles). Total industry revenue is currently more than $110 billion annually.

Transactions involve a secret reserve price ascending auction potentially followed by post-auction bargaining. The role of the auction stage is to ensure that it is the highest-value bidder (rather than some other bidder) who trades with the seller through the auction or who faces the seller in a post-auction negotiation. For each car, an auctioneer runs a rapid (approximately 90-second) ascending auction. Prior to the auction, the seller chooses a reserve price, which is reported to the auction house but hidden from bidders. If the auction price (the highest bid) fails to reach the seller's secret reserve price, the buyer is given the option to walk away from the transaction. If the buyer chooses not to walk away, bargaining ensues between the high bidder (whom we will call the *buyer*) and the seller. For most sales, this bargaining stage simply involves the seller immediately accepting or rejecting the auction price in person.[22] About 26% of auctions end with the auction house facilitating negotiations over the phone through an auction house employee — a *mediator*.[23] The mediator first calls the seller and relays the auction price.[24] The seller can choose to accept this price, counter, or quit (ending the negotiation). If the seller counters, the mediator calls the buyer. This process continues until one party accepts or quits. The auction house company records all actions taken by either party, the identity of each party, and the mediator's identity.

The mediated negotiation data consist of several hundred thousand realizations of bargaining sequences at six different auction house locations (in distinct geographic markets) run by the same parent company from 2006–2010. We will frequently use the term *thread* to refer to a bargaining

[22]In some cases, if the seller is not present and the auctioneer observes a sufficiently wide gap between the auction price and the reserve price, the auctioneer may decide to reject the auction price on behalf of the seller.

[23]This 26% is not conditional on the auction failing, but rather represents the percentage of observations for which we see a mediator identity recorded among all attempts to sell the vehicle in our full auction-plus-bargaining dataset, described in Section 5.5. Observations *without* a mediator recorded are those that fall into one of the following cases: (i) the car sells through the auction price exceeding the reserve price, (ii) the auction price is below the reserve price and is accepted or rejected by the seller in person (meaning a mediator is never involved), or (iii) the auction price is so far below the reserve price that the auctioneer opts against contacting the seller. The 26% of transactions with mediated negotiation — which we will refer to as our *mediated negotiation data* — constitute the primary dataset we use in this and the following sections to study mediator heterogeneity. In Section 5, we bring in observations corresponding to cases (i)–(iii) as well, as these observations are key for identifying buyer and seller value distributions.

[24]The mediator's identity is not known to the buyer or seller before the negotiation begins.

sequence. The data storage system creates a new record for each action taken during a given bargaining thread, allowing us to see that some threads involve multiple mediators facilitating different stages of the negotiation. Among all bargaining threads, 67.06% are handled by one mediator per thread, 30.34% are handled by two mediators, and the remaining by more than two. We limit our analysis to threads involving single mediators to facilitate our attribution of outcomes to mediators.[25] Because we want a sufficient number of observations per mediator to estimate mediator-specific outcomes, we restrict our sample to mediators whom we observe in at least 50 separate bargaining threads. Appendix B describes additional data cleaning. In the end, we have 114 mediators and 75,090 threads with 6,226 distinct sellers and 6,258 distinct buyers.

The industry's primary performance indicator is the probability with which trade occurs; the auction company's goal is to operate a liquid two-sided market that can attract both buyers and sellers (Treece 2013, Lacetera et al. 2016).[26] Each mediator also has this same goal. We will largely evaluate differences in mediator performance by this trade probability metric.[27] We also look at heterogeneity in the final price of successful trades. To make prices comparable across cars, for much of our analysis, we normalize them by the auction company's estimated market value of the car (which we will refer to as the *book price*).

Table 1 shows thread-level statistics for our main sample.[28] Agreement occurs in 62.3% of threads. *Final price* refers to final negotiated price when trade occurs. This price may coincide with the auction price if the seller accepts that price in the negotiation (the auction price is the first offer in the negotiation stage). The average final price for a successful trade is $5,605, lower than the average book price of $6,958. The average final price is also between the average auction price of $5,572 and average secret reserve price of $7,347, though it is much closer to the former.[29] A bargaining thread ends after 1.39 offers on average, and within six hours. The average car is 6.35

[25] Appendix B compares the single-mediator threads to those with multiple mediators and presents evidence suggesting that dropping these observations does not drastically affect our results.

[26] In contrast, price maximization (a common auction design objective) attracts sellers but not buyers.

[27] Mediators are not paid a commission based on individual trades they facilitate; they are paid a salary. However, our conversations with individual mediators, and our in-person observations of the mediation process, make it clear that mediators' primary objective is a high trade probability.

[28] Appendix Table B.1 shows means from Table 1 separately by the six auction house locations. Locations vary in the number of negotiations they handle, ranging from 6,766 to 19,293, with smaller locations having fewer mediators (this ranges from 10 to 31 across locations). Locations also differ in the fraction of cars sold by fleet/lease sellers, ranging from 0.396 to 0.536. See Appendix B for more discussion. In our analysis of mediator heterogeneity below, we control for these differences in a number of ways, both through location fixed effects and through controls varying at the thread-level within a location.

[29] Recall that any transaction that reaches the negotiation stage is one in which the reserve price exceeds the auction price, and hence trade failed to occur at the auction stage.

Table 1: Descriptive Statistics at Bargaining Thread Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.623	0.485	0	1
Final price ($)	5,605	4,940	850	12,700
Book price ($)	6,958	5,264	1,550	14,475
Auction price ($)	5,572	4,938	800	12,700
Reserve Price ($)	7,347	5,407	1,900	15,000
No reserve	0.255	0.436	0	1
# Offers in a thread	1.393	0.704	1	2
Length of a thread (hours)	5.769	14.934	0.337	15.6
Fleet/lease car	0.474	0.499	0	1
Car age (years)	6.353	3.605	2	12
Mileage	93,341	51,099	30,161	158,011
Engine displacement (liters)	3.598	1.525	2	5.7
No. Threads	75,090			

Notes: Table presents descriptive statistics at the thread level. Final price is conditional on trade occurring.

years old and has 93,341 miles on it.[30] Roughly half of threads correspond to fleet/lease sellers.

Table 2 shows descriptive statistics at the mediator level. Some of these outcomes are thread-level characteristics averaged within threads handled by a given mediator. Among the 114 mediators, 45.8% are female.[31] The average mediator has worked at her current auction house for four years.[32] In our data, she handles 659 bargaining threads and successfully facilitates 60.4% of them. While our sample restrictions enforce that each mediator handle at least 50 negotiations, Table 2 shows that even the 10th percentile mediator handles more than twice this amount. The achieved final price is, on average, 79% of the reserve price and 2.2% higher than the auction price. Dispersion in average trade probability is quite large across mediators, with a standard deviation of 0.196 (while the dispersion in prices is much smaller). The following section analyzes this dispersion.

[30]Table 1 shows that, for 25.5% of observations, there is no reserve price reported. This typically implies that (i) the seller plans to be present at the auction to signal to the auctioneer whether to let the car sell at the current bid (effectively acting as a shill bidder) or (ii) the seller wants an auction house mediator to call her regardless of the auction price to let her accept or reject this price over the phone. The probability of agreement is similar in cases where the reserve price is missing vs. not (see Online Appendix D.1.1 of Larsen 2021).

[31]We find no difference in trade probabilities by mediator gender (see Appendix Table B.5).

[32]We observe each mediator's employment start date at an auction house. The reported years of employment in Table 2 are computed by first calculating the mediator's employment length up through the date of each bargaining thread and then taking the average of this quantity across all of her threads in our data.

Table 2: Descriptive Statistics at Mediator Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.604	0.196	0.374	0.899
Final price/book price	0.834	0.0751	0.754	0.932
Final price/reserve price	0.790	0.0626	0.725	0.861
Final price/auction price	1.022	0.0359	1	1.04
Female	0.458	0.501	0	1
# Threads mediated	659	686	102	1,473
Years of employment	4.150	5.308	0.379	9.96
No. Mediators	114			

Notes: Table presents descriptive statistics at the mediator level.

4 Heterogeneity in Outcomes Across Mediators

In this section, we provide evidence on mediators' assignment and their heterogeneous performance as measured by their trade probability. We show that this heterogeneity is not due to sampling error. We also examine differences in their final negotiated prices. Finally, we examine how mediators differ in *when* they achieve trade, how experience correlates with performance, and how different mediators perform when faced with greater disparity in buyer and seller values.

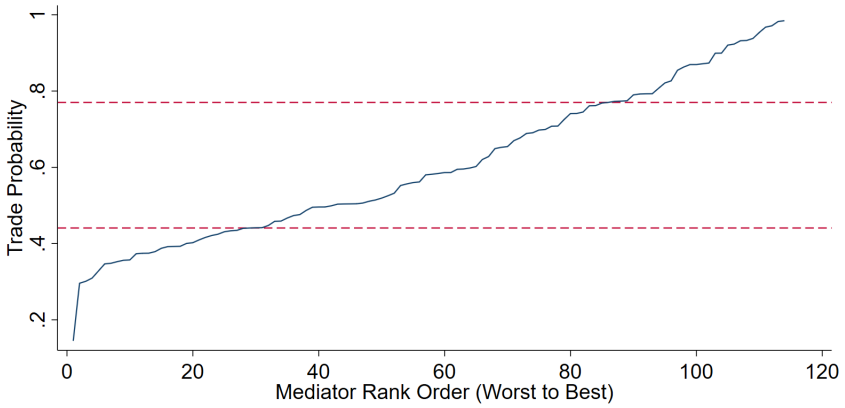
4.1 Mediator Effects on Agreement. We first report raw differences in trade probability across the 114 mediators. Let $Agree_i$ be an indicator variable representing whether trade occurs for thread i, which is mediated by intermediary $k(i)$. Let $\mathcal{J}_k$ be the set of all threads mediated by k. The average of $Agree_i$ over the subset of k's threads is

$$\overline{Agree}_k \equiv \frac{1}{|\mathcal{J}_k|} \sum_{i \in \mathcal{J}_k} Agree_i. \tag{1}$$

Figure 1 shows these probabilities, sorted from smallest to largest. A core stylized fact of our study is that different mediators have very different success probabilities. The best to worst mediators span a difference of 0.840. The 90–10 percentile spread is 0.525. The 25th-percentile of mediator trade probabilities is 0.441 and the 75th-percentile is 0.770 (shown with horizontal dashed lines in Figure 1). We will refer to higher-trade-probability mediators as *higher-skilled*.

To better understand the industry, we spent time observing the bargaining/mediation process and interviewing mediators, buyers, and sellers. Through this effort, we learned about potential

Figure 1: Mediator Differences in Trade Probability



Notes: Figure shows average trade probability for each mediator, ranking them from low to high. Dashed lines represent the trade probability achieved by the 25th and 75th percentile mediators.

drivers of mediator heterogeneity and the mediator assignment process. These conversations revealed that each seller typically has an auction house employee assigned to manage the relationship between that seller and the auction house, and, where possible, this employee mediates negotiations involving that seller. A seller's default mediator is often unavailable, however, and in such cases alternative mediators handle negotiations. At one location, we were told that negotiations involving fleet/lease sellers, rather than dealer sellers, are more often handled by a pre-assigned mediator, but even for fleet/lease sellers this is not always the case. At this location, we were able to observe that, for several exciting hours each week, the room in which the phone calls occur for dealer sellers is alive with activity, akin to a miniature stock market trading floor, with no obvious protocol determining which mediator handles a given thread.[33]

Mediator assignment for a given thread thus involves an element of randomness, but is unlikely to be completely random. Nonrandom assignment is potentially problematic for our goal of measuring mediator heterogeneity if good mediators are systematically assigned to threads with characteristics associated with higher trade probabilities; for example, if effective mediators are systematically assigned to less-expensive cars, cars with lower reserve prices or higher auction prices, or sellers or buyers who have a higher propensity to agree. We evaluate the extent of this non-random

[33]Throughout the body of the paper, we use a sample of negotiations that includes both dealer sellers and fleet/lease sellers. In Appendix D we replicate all tables and figures from the body of our paper on separate samples for dealers vs. fleet/lease sellers and find that our results are largely similar across the two seller types.

assignment and identify features of the negotiation that help control for non-random assignment.

We first construct a leave-one-out estimate of a mediator's agreement rate, given by $\overline{Agree}_{-i,k(i)} \equiv \frac{1}{|\mathcal{J}_k|-1} \sum_{\{j:j\neq i, j\in\mathcal{J}_k\}} Agree_j$. This is the average trade probability for all threads *other than* i that are mediated by k. We construct a similar leave-one-out average of the trade probability for the seller and buyer of thread i.[34] We then estimate the following regression:

$$Agree_i = Z_i'\alpha + \epsilon_i, \tag{2}$$

and an alternative version of this regression where the outcome variable is instead $\overline{Agree}_{-i,k(i)}$.

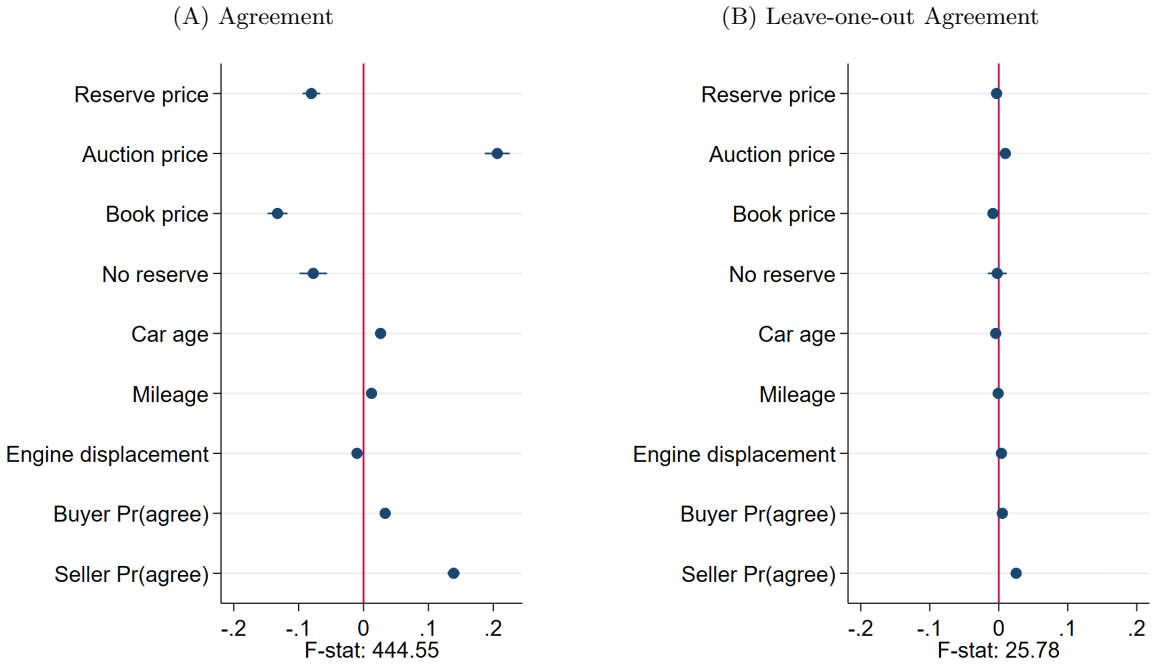
Z_i contains characteristics of thread i that are determined before the negotiation begins (and hence before any mediator influence). The auction price and reserve price are conceptually quite important to include in Z_i. All observations in our mediated negotiation data are cases where the auction price is below the reserve price (the condition that leads to auction failure and entering negotiations). However, observations in which the auction price is closer to the reserve price are more likely cases where gains from trade exist and agreement is feasible. This is because a higher auction price should reflect a higher willingness to pay of the buyer, and a lower reserve price a lower value of the seller. We formalize this in our structural analysis in Section 5, where auction prices and reserve prices play key roles in identifying buyer and seller values.

Z_i also includes other factors that potentially relate to mediator assignment and to trade probabilities: fixed effects for the auction house location and negotiation date; vehicle mileage, age, engine displacement, and book price; and the leave-one-out agreement rates of the seller and buyer. The first version of (2), with $Agree_i$ as the outcome, allows us to explore whether thread characteristics correlate with the likelihood that the thread ends in agreement. The second version of (2), with the leave-one-out agreement as the outcome, speaks to whether these characteristics correlate with a mediator's average agreement rate on the mediator's *other* threads. Comparing estimates from these two regressions allows us to examine whether mediators with higher agreement rates are systematically matched to cars, buyers, or sellers that are more likely to result in an agreement.

Figure 2.A shows estimates from the first regression, with 95% confidence intervals about each

[34] The leave-one-out agreement rate for seller $s(i)$ is the empirical frequency with which $s(i)$ comes to agreement based on threads other than i that involved s. Using similar notation to the mediator leave-one-out rate, this can be written $\overline{Agree}^{seller}_{-i,s(i)} \equiv \frac{1}{|\mathcal{J}_s|-1} \sum_{\{j:j\neq i, j\in\mathcal{J}_s\}} Agree_j$, where $\mathcal{J}_s$ is the set of all threads with s as the seller. The buyer leave-one-out agreement rate can be written $\overline{Agree}^{buyer}_{-i,b(i)} \equiv \frac{1}{|\mathcal{J}_b|-1} \sum_{\{j:j\neq i, j\in\mathcal{J}_b\}} Agree_j$, where $\mathcal{J}_b$ is the set of all threads with b as the buyer.

Figure 2: Mediator Assignment Test



Notes: Figure shows coefficients from (2). Panel A shows coefficients when the outcome is $Agree_i$ and panel B where the outcome is the leave-one-out agreement rate $\overline{Agree}_{-i,k(i)}$. In addition to the variables listed in the figure, regressions include fixed effects for auction house locations and negotiation date. Covariates are standardized to have standard deviations equal to 1. 95% confidence intervals, computed by clustering at the mediator level, surround each point estimate. Buyer Pr(agree) and Seller Pr(agree) represent buyer and seller leave-one-out agreement.

estimate. Other than engine displacement, each feature is significantly related to the likelihood that a thread ends in agreement. In particular, threads are more likely to end in agreement if they have a lower reserve price or higher auction price (as foreshadowed in our discussion above), if the car is a less-expensive car (lower book price), or if the buyer or seller is a priori a more agreeable agent (i.e., more likely to agree as measured by the buyer's or seller's leave-one-out agreement rate).

Figure 2.B shows the coefficients from the second regression, where the outcome is $\overline{Agree}_{-i,k(i)}$, the leave-one-out agreement rate. If good mediators are not systematically assigned to threads that are a priori more likely to end in agreement, we would expect coefficients from this regression to be much closer to zero. Figure 2.B confirms this. The results still attest that assignment is not completely random: better mediators — those with higher leave-one-out average trade probabilities — tend to be assigned to threads with slightly higher auction prices and with sellers who are a priori slightly more likely to agree. However, these effects are small relative to the magnitudes in

panel A. Indeed, the F-statistic decreases from 445 in the first regression to 26 in the second, and the coefficients are about five times smaller. Thus, while thread-level features do predict realized agreement, they have much less predictive power for leave-one-out agreement rates. These results suggest that, while some of the raw heterogeneity in trade probability across mediators arises from other thread-level characteristics, these characteristics themselves are not strongly correlated with the assigned mediator's overall performance in terms of trade probability. We will control for this non-mediator-related heterogeneity in our analysis below.[35]

To measure mediator performance incorporating these controls, we expand (1) to estimate an effect for each mediator conditional on thread-level characteristics, as follows:

$$Agree_i = \beta_{k(i)} + X_i'\phi + \epsilon_i. \tag{3}$$

In (3), $\beta_{k(i)}$ is the effect for mediator k and X_i is a vector varying across specifications.[36] Following Lacetera et al. (2016), we reposition estimated mediator effects to be mean zero:

$$\hat{\beta}_{norm,k} = \begin{cases} \hat{\beta}_k - \frac{1}{M}\sum_{j=2}^{M}\hat{\beta}_j & for\ k \neq 1 \\ 0 - \frac{1}{M}\sum_{j=2}^{M}\hat{\beta}_j & for\ k = 1, \end{cases}$$

where 1 denotes the omitted mediator in regression (3).

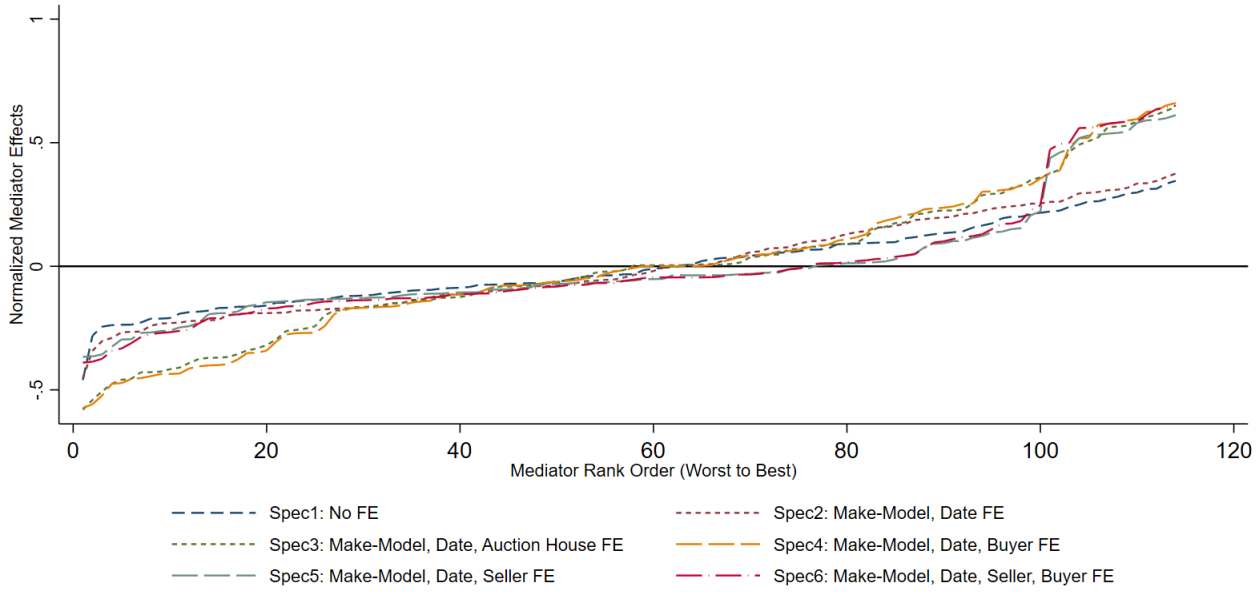
Figure 3 displays estimates of mediator fixed effects from increasingly stringent specifications for X_i. In the baseline case, which we denote *spec 1*, X_i includes various thread-level features: the car's age, book price, engine displacement and mileage; a dummy for whether the car is sold by a fleet/lease seller; and reserve and auction prices. In spec 2, we add fixed effects for the car's make-model combination and the negotiation date. Spec 3 adds auction house location fixed effects and spec 4 adds buyer fixed effects.[37] Spec 5 replaces buyer with seller fixed effects. The most-saturated

[35] Appendix B offers an alternative test of mediator assignment that suggests that mediators are largely randomly assigned to car types (make and model) and to buyers, and that mediator assignment to sellers is less random. This highlights the importance of controlling for seller fixed effects in our analysis, as we do below.

[36] We denote this vector X_i to distinguish it from Z_i in (2), as the controls in the two regressions only partially overlap. In particular, depending on the specification, X_i can include seller and buyer fixed effects, rather than just the buyer and seller leave-one-out agreement rates found in Z_i.

[37] Our data contain separate lists of buyer, seller, and mediator identities for each auction house location. We therefore treat these identities as specific to a given location. In practice, it is possible, for example, that the same seller operates at two different locations in our data, but our analysis will treat this as two distinct sellers. This is not necessarily a downside for the analysis; rather, we highlight this to clarify that any time we control for seller, buyer, or mediator identities, those indicators will naturally absorb auction house location effects. Note also that

Figure 3: Mediator Fixed Effects for Trade Probability Under Different Specifications



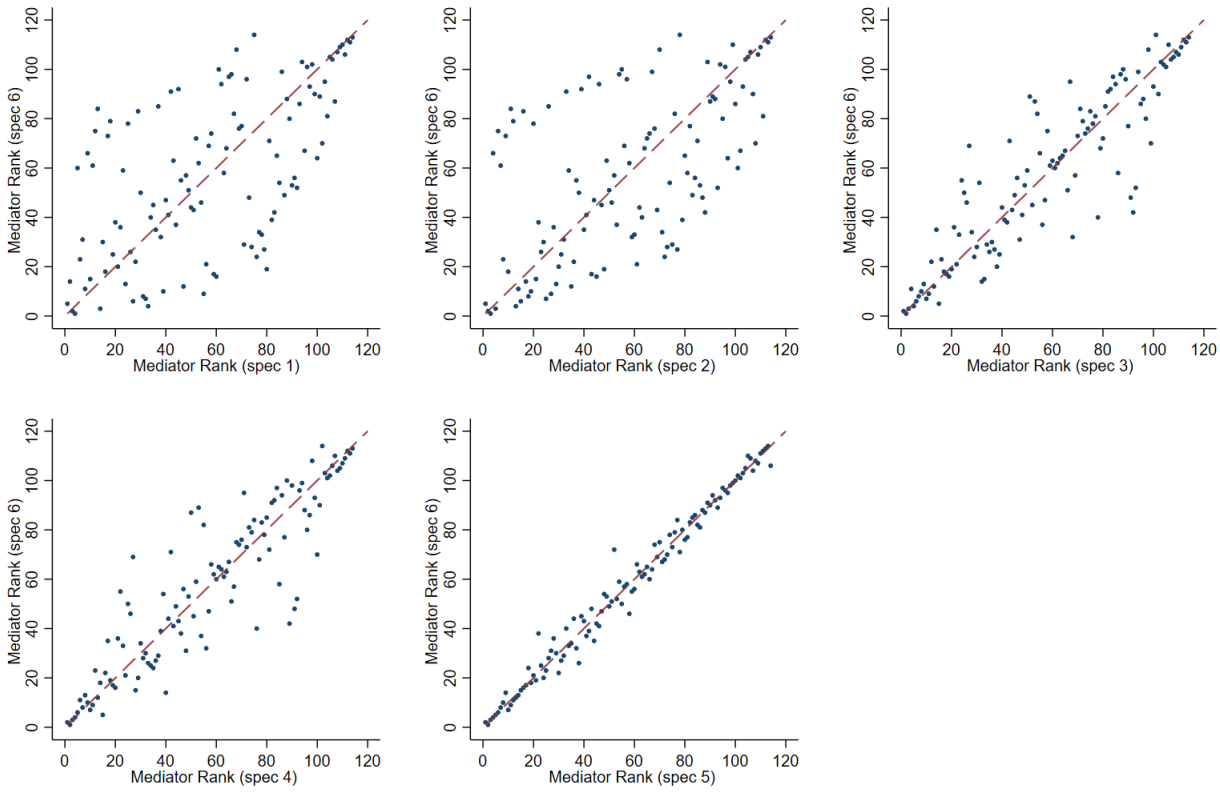
Notes: Figure shows normalized mediator fixed effect estimates for trade probability under increasingly stringent specifications for X_i from regression (3).

version (spec 6) includes all controls. Seller fixed effects are of particular interest given that Figure 2 suggests that assignment of mediators to sellers may not be entirely random, and certain sellers may be more prone to agree independent of the mediator. Seller fixed effects help absorb variation in mediators' trade probabilities that is due solely to the seller. Specs 5 and 6 allow us to examine whether mediator heterogeneity remains even after controlling for seller effects.

Figure 3 shows results comparing these different specifications, ranking mediator fixed effects from lowest to highest as in Figure 1. Controls change the shape of the mediator effects curve, but the heterogeneity remains relatively consistent and large across all specifications. The interquartile range is also large enough to be economically meaningful in each specification: the minimum value of the 75th to 25th percentile gap is 18.07 percentage points, which is large relative to the mean trade probability in the data (0.623 from Table 1). Figure 4 illustrates how mediator fixed effects are correlated across specifications. In each panel, vertical axes show mediators' ranks (from 1 to 114) as measured by her fixed effect from spec 6. Horizontal axes show mediators' ranks from other

some fixed-effect cells are singletons (e.g., some seller or buyer cells), and these observations are naturally dropped from regressions.

Figure 4: Mediator Rank Across Specifications



Notes: Figure presents mediators' rank in trade probabilities across different specifications. Each data point represents a mediator. Vertical axis shows mediator fixed effects from spec 6, horizontal axis shows fixed effects from other specifications, and dashed line shows 45-degree line.

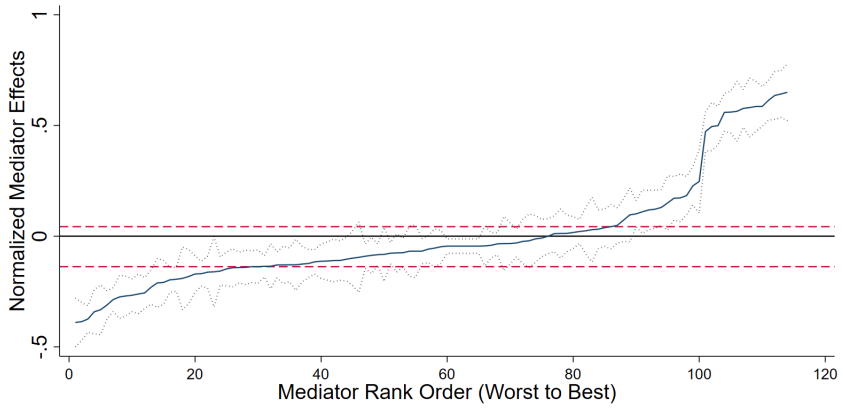
specifications (1–5). Rankings are highly correlated across specifications.[38]

4.2 Mediator Effects on Prices. We repeat estimation of (3) using final negotiated prices (normalized by book prices), denoted $Price_i$, as the outcome of interest rather than $Agree_i$. Figure 5.B shows the estimated mediator fixed effects under the saturated model (spec 6), along with pointwise 95% confidence intervals. Panel A shows the trade-probability effects (also under spec 6) for comparison. The estimated mediator effects for final prices are smaller in magnitude than those

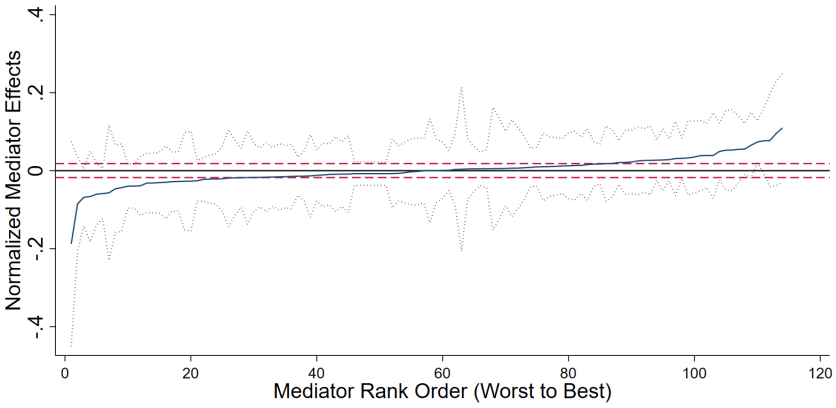
[38] Appendix B includes additional analysis of mediator heterogeneity. Appendix Table B.4 adds additional controls to the Figure 3 regressions, controlling for the number of previous interactions between a mediator-buyer pair, mediator-seller pair, or buyer-seller pair. These controls have little to no effect. Appendix Table B.6 examines incremental changes in adjusted R^2 from adding (to a regression where the right-hand side only includes X_i from spec 1) or subtracting (from spec 6) any of the following fixed effects: mediator, buyer, seller, auction house location, or make-model combination. The latter two have negligible effects; the other three help explain variation in outcomes.

Figure 5: Mediator Fixed Effects Under Most-Saturated Specification

(A) Trade probability



(B) Normalized final price



Notes: Panel A shows mediator fixed effect estimates for trade probability along with 95% confidence intervals under spec 6 of (3). Panel B shows results from the same exercise but using normalized final price as the outcome variable. Horizontal dashed lines represent the estimated fixed effect for the 25th- and 75th-percentile mediators.

for agreement. A one-standard-deviation increase in mediator performance is only associated with a 3.8 percentage point increase in the final price, but a 24.9 percentage point increase in the trade probability. Confidence intervals for each estimated effect are also wider for the price measure: for 113 out of the 114 mediators, price effects are not significantly different from zero. These results are unsurprising given mediators' objectives, which the industry clearly dictates as increasing trade.

4.3 Sampling Error. A concern with any study of heterogeneity is that some (and potentially all) heterogeneity may be driven by sampling error: even if true mediator performance were constant

across mediators, some variance in outcomes would arise in any finite sample. Here we quantify the variation due to sampling error through a parametric bootstrap approach that takes into account other factors — such as seller identity or the make and model of the car — that may also affect trade probability.[39] We begin with the null hypothesis that mediators have no effect on trade probability. We regress $Agree_i$ on X_i from spec 6 of (3), without including mediator fixed effects. Let $\widehat{Agree}_i$ denote the predicted probability of trade from this regression for observation i. We then generate a large number of simulated datasets, indexed by ℓ, of identical size to the original data. For simulated dataset ℓ, we generate a synthetic trade indicator for observation i, $\widetilde{Agree}_i^{\ell}$, which is equal to 1 with probability $\widehat{Agree}_i$ and 0 otherwise.[40]

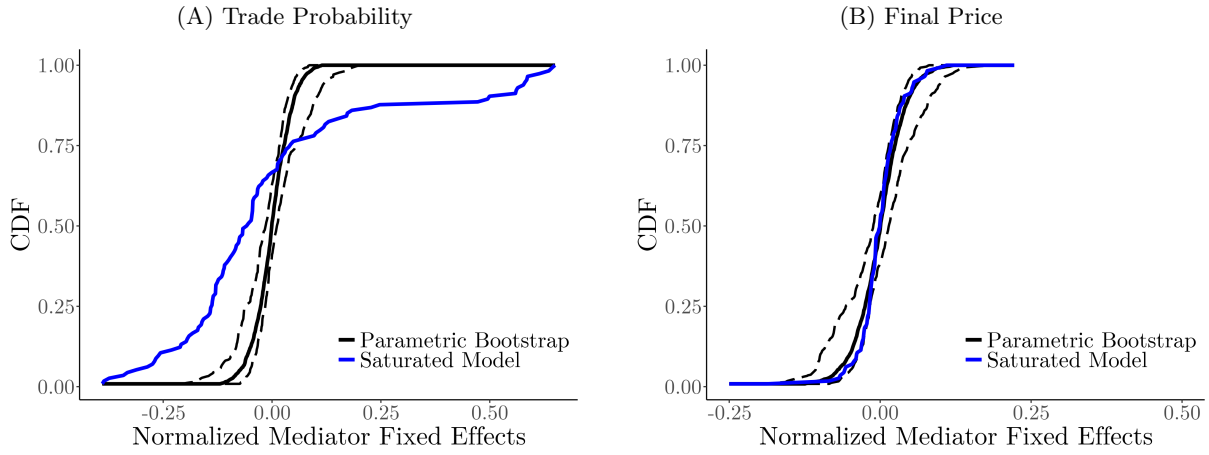
This resampling method preserves much of the data structure. It preserves, for example, the size distribution of mediators; the correlations between trade probabilities and features such as auction house locations, sellers, buyers, car models, and other controls in X_i; and the correlation between these features and mediator assignments. For example, if a certain seller tends to have a higher trade probability in the raw data, the procedure ensures that this seller will also have a higher trade probability in the simulated data. If certain mediators are systematically more likely to interact with particular sellers or car types, this is also preserved, because the *outcomes* are synthetically produced but not covariates. The main part of the data structure that is destroyed in the synthetic data is any correlation between mediator identities and trade probabilities, conditional on the controls X_i, which is exactly the variation that identifies mediator effects.

We simulate 100 such synthetic datasets and estimate mediator fixed effects in each as in spec 6 of (3). Because mediator identities have no effect on trade probabilities in the synthetic data, any variance in mediator effects in the synthetic data reflects pure sampling noise. Figure 6.A plots the CDF of the median bootstrap estimates across these 100 datasets, along with 95% confidence bands. The figure also shows the CDF of the actual mediator fixed effect estimates from spec 6 of (3). The distribution of trade probability fixed effects is quite dispersed, lying well outside the 95% confidence bands of the bootstrap distribution, suggesting that the heterogeneity we find is not purely due to noise. This same approach admits a one-sided test of the null hypothesis that the standard deviation of actual mediator fixed effects is no larger than what we would expect solely

[39] See Section 10.4 of Efron and Hastie (2021) for a description of the parametric bootstrap. We are grateful to an anonymous referee whose comments inspired this approach.

[40] Because we use a linear probability model, in a small fraction of the data, the estimated probabilities $\widehat{Agree}_i$ are greater than 1 or less than 0; in these cases, we simply set the probability equal to 1 or 0, respectively.

Figure 6: Parametric Bootstrap Test of Heterogeneity from Sampling Error



Notes: Panel A shows parametric bootstrap test results for trade probability. Panel B presents results for final price normalized by book price. To obtain the bootstrap estimates and confidence intervals for trade probabilities, we estimate a model of trade probabilities using all variables in spec 6 of (3), except that we exclude mediator fixed effects. For each bootstrap sample, we then generate simulated trade indicators, equal to 1 with probability equal to the predicted probabilities from the estimated model. We similarly estimate a model of prices using spec 6 of (3), and then simulate prices using the regression predictions combined with the regression estimate of the error variance. In each case, we use 100 bootstrap samples. Blue lines show the estimated distribution of fixed effects from the actual data and black lines show the bootstrap median and 95% pointwise confidence bands.

due to sampling error. We solidly reject this null: the standard deviation of the actual mediator fixed effect estimates (0.25) is well outside the one-sided 95% bootstrapped confidence value (0.06).

We adopt a similar method to evaluate the effects of mediators on prices, regressing $Price_i$ on X_i from spec 6 of (3). Let $\widehat{Price_i}$ be the predicted price from this regression for observation i and let $\hat{\sigma}$ be the estimated standard deviation of the residuals. Because prices are continuous rather than discrete, simulating prices requires a stronger stance on the data generating process. For simulation ℓ, we generate a synthetic price outcome for observation i, $\widetilde{Price}_i^{\ell}$, by drawing from a $N(0, \hat{\sigma})$ distribution and adding the realizations to $\widehat{Price_i}$. We create 100 such datasets and estimate mediator effects in each as in spec 6 of (3) but using $\widetilde{Price}_i^{\ell}$ as the outcome.

The results of this exercise are shown in Figure 6.B. The estimated mediator fixed effects for prices fall within the 95% confidence band of our bootstrap estimates. Thus, unlike our effects on trade probabilities, we cannot reject the null that mediators have no effect on prices, and that any apparent effect observed in Figure 5.B is due to sampling error. This is also confirmed by the standard deviation of the actual mediator fixed effects for prices, which is 0.038, whereas the 95th

Table 3: Agreement Rates in First vs. Later Rounds and Mediator Experience

	(1) $Agree_i$	(2) $Agree_i^1$	(3) $Agree_i^{>1}$	(4) $Agree_i$	(5) $Agree_i^1$	(6) $Agree_i^{>1}$	(7) $Agree_i$	(8) $Agree_i^1$	(9) $Agree_i^{>1}$
High	0.254***	0.250***	0.004						
	(0.010)	(0.010)	(0.006)						
Medium	0.102***	0.098***	0.004						
	(0.006)	(0.006)	(0.004)						
Experience (years)				0.005***	0.004***	0.001*	0.010***	0.019***	-0.009***
				(0.001)	(0.001)	(0.001)	(0.004)	(0.004)	(0.002)
Mediator FE							Y	Y	Y
R2	0.371	0.438	0.233	0.398	0.467	0.257	0.416	0.484	0.262
N	71,154	71,154	71,154	51,589	51,589	51,589	51,589	51,589	51,589

Notes: "High" and "Medium" are dummies for a given thread i being mediated by a mediator in the top and middle tercile of mediator trade probability fixed effects, respectively, where the fixed effects are spec 6 of (3). In Columns 1-3, where we analyze mediator terciles, the benchmark group is the bottom tercile of mediators. Experience is measured for a given thread as the time elapsed between the thread mediator's start date, which we observe in the data, to the date of the bargaining thread. The dependent variable in columns 1, 4, and 7 is a dummy for the thread ending in agreement. In columns 2, 5, and 8, the dependent variable is a dummy for the thread ending in agreement in the first period, and in columns 3, 6, and 9, the dependent variable is a dummy for the thread ending in agreement later than the first period. The vector X_i from spec 6 of (3) is included as a control in all regressions. The number of observations is fewer than in Table 1 because experience is not available for some mediators and because singleton fixed effect cells are dropped from the analysis. Significance levels: *: $p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$.

percentile standard deviation across parametric bootstrap samples is 0.06.[41]

4.4 Other Differences Across Mediators. This section examines several differences across mediators. First, we show that effective mediators' threads have a higher tendency to end in agreement in the first period, and that this fully explains their better performance. Let $Agree_i^1$ be equal to 1 if thread i ends in agreement in the first period (meaning the seller accepts the auction price, the first bargaining offer), and 0 otherwise. Let $Agree_i^{>1}$ be equal to 1 if thread i ends in agreement some time after the first period, and 0 otherwise. Thus, mechanically, $Agree_i = Agree_i^1 + Agree_i^{>1}$. We define three tercile groups of mediators, which we call *low*, *medium*, and *high*, ranked by their estimated fixed effects for trade probability from spec 6 of (3).[42] We then analyze the three components $Agree_i$, $Agree_i^1$, and $Agree_i^{>1}$ in these groups.

Table 3 column 1 displays estimates from regressing $Agree_i$ on dummies for tercile groups, with the lowest group as the omitted category. The vector X_i from spec 6 of (3) is also included. We find higher trade probabilities for medium and high mediators; this is mechanical given their grouping by

[41] We find similar results using final prices normalized by the auction price or reserve price rather than book price or using an alternative version of the bootstrap test relying only on seller fixed effects (which, again, absorb auction house location effects) in X_i; See Appendix Figure B.3. Appendix B also proposes a test based on splitting our main sample in half that also confirms these results.

[42] For this exercise, as well as for the structural model in Section 5, we take estimated mediator rankings as truth; that is, we treat the spec 6 estimates from (3) as though they control well enough for covariates that we have a correct understanding of who the best and worst (and middle-ground) mediators truly are.

trade probabilities. Column 2 displays results when the outcome is instead the first-round agreement indicator, $Agree_i^1$. The coefficients for the medium and high groups are positive and significant, and similar in magnitude to those in column 1. Column 3 uses instead the later-round agreement indicator, $Agree_i^{>1}$. The coefficients are insignificant in both cases, and the standard errors are tight enough to reject economically large effects. Thus, effective mediators' increased propensity to agree in the first period appears to fully explain their better performance.[43]

We next examine how mediators' experience affects their trade probabilities. First, Figure 7.A shows a scatter plot of mediator experience against mediators' estimated fixed effects from spec 6 of (3). Experience is measured, for a given thread, as the time elapsed between the mediator's employment start date and the start time of the current thread; we then average this across threads a mediator handles.[44] We observe that more experienced mediators are more effective. However, a large degree of mediator skill is unexplained by experience.

We also quantify experience effects through variants on the following regression:

$$Agree_i = X_i'\xi + \psi Experience_{k,i} + \omega_i \tag{4}$$

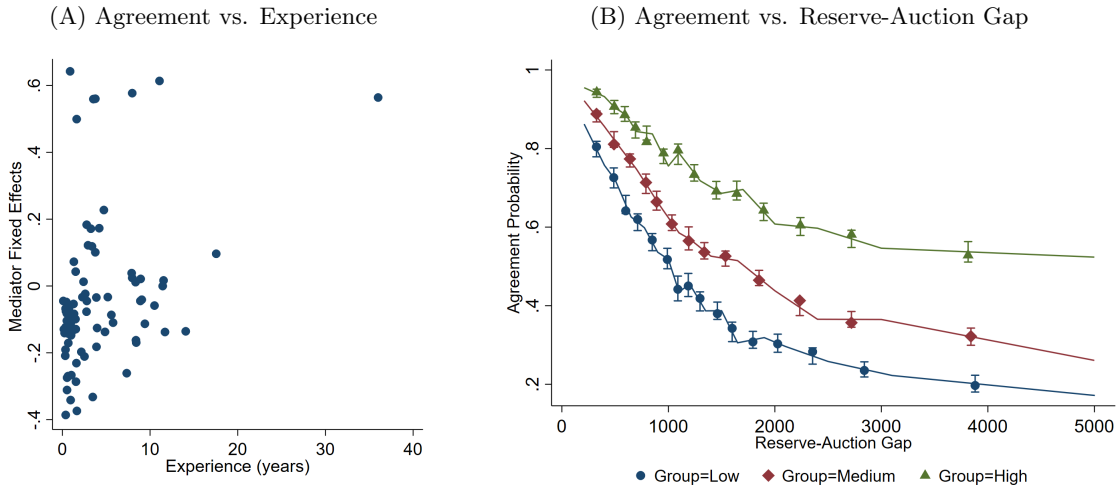
where $Experience_{k,i}$ denotes the experience level of mediator k up to the point where thread i begins, and X_i is as in spec 6 of (3). Table 3 column 4 shows the results: mediator experience is positively correlated with agreement, an effect that is small in magnitude but statistically significant.

We augment (4) to include mediator fixed effects, exploring how individual mediators' performance changes as they gain experience. Table 3 column 7 shows the results. The coefficient on experience is positive and significant, suggesting that individual mediators, even those who are poor performers on average, improve in their trade probability with additional experience, becoming more like high-skilled mediators. The effects are quantitatively modest: A one-year increase in experience is associated with a one-percentage point increase in trade probability, a movement upward of about three slots (relative to the 25th percentile) in the mediator rankings from spec 6 of (3). Mediator performance thus appears to improve with experience, but some variation in effectiveness remains that experience cannot explain. More experienced mediators have higher first-round agreement

[43] Appendix D replicates all analyses separately for fleet/lease vs. dealer sellers. We find that the null effect in column 3 of Table 3 is driven by a positive effect for dealer cars and a negative effect for fleet/lease cars. In both samples, the strong positive correlation between a mediator's overall trade probability and her first-round trade probability (column 2) remains.

[44] We observe mediators' employment start date in the data.

Figure 7: Agreement vs. Mediator Experience and Agreement vs. Reserve-Auction Gap



Notes: Panel A displays a scatter plot of mediators' estimated fixed effects from spec 6 of (3) vs. average work experience in years. Panel B shows a binned scatter plot with the probability of agreement on the vertical axis and the reserve-auction gap on the horizontal axis, with the plot created separately for negotiations handled by high-, medium-, and low-skilled mediators. The binned scatter plot sets the polynomial degree and number of smoothness constraints to 1 and uses the bias-corrected confidence interval of Cattaneo et al. (2024).

probabilities (columns 5 and 8) and lower rates of later-round agreement (column 9; the result in column 6 is not significant at the 0.05 level).[45]

Table 3 offers two takeaways. First, effective mediators appear to succeed by convincing sellers to concede in the first period more often and this channel appears to explain the entirety of high-skilled mediators' better performance.[46] Second, experience appears to be an important driver of mediators' effectiveness: more experienced mediators have higher trade probability fixed effects, and mediators (even low-skilled ones) become more effective as they gain experience. A narrative consistent with these stylized facts is that effective mediators may be able to coordinate buyers and sellers on playing an equilibrium that involves sellers conceding faster. The structural exercise in Section 5 will allow for these possibilities by viewing mediators of different skill levels as indexing different equilibrium strategies. Our results in Table 3 suggest that the skills required to implement faster agreement may be difficult to acquire: mediators improve slightly with experience, but there is

[45] In Appendix D we show that the strong positive effects in columns 7 and 8 are driven by negotiations involving fleet/lease sellers. The negative effects in column 9 appear in both the dealer and fleet/lease samples.

[46] Our result that better mediators achieve agreement in the first period also suggests that these mediators reduce delay, which is costly in many bargaining models. The structural analysis in Section 5 allows us to examine whether the realized welfare gains, including bargaining costs, are higher for these higher-trade-probability mediators.

large variation in mediators' abilities that is not explained by experience. This finding is consistent with a number of other papers documenting persistent productivity differences among individual employees (e.g., salespeople in Jindal and Newberry 2022 or auctioneers in Lacetera et al. 2016) despite strong incentives for firms to attempt to remove heterogeneity through training, suggesting that some skills may be difficult to transfer from high to low performers.

Our final investigation in this section looks for evidence regarding the *surplus*, or *gains from trade*, in trades that are consummated by different mediators. The gains from trade can be defined as the buyer's willingness to pay (the buyer's value) less the price at which the seller is willing to sell (the seller's value). While these objects are unobserved in the data, it is intuitive that a seller with a higher value would report a higher secret reserve price and that the buyer's value is related to the auction price: the last bidder standing in the auction must have a value higher than the auction price. The model in Section 5 formally derives these and other properties. Here, we simply use these arguments to motivate that the reserve price minus the auction price — we call this the *reserve-auction gap* for brevity — is negatively related to the true surplus.[47]

Figure 7.B shows a binned scatter plot with trade probability on the vertical axis and the reserve-auction gap on the horizontal. We plot this separately using high-, medium-, and low-skilled mediators. Within each group, the higher the reserve-auction gap, the lower the probability of agreement. This is because a larger reserve-auction gap reflects a greater chance that gains from trade are small or nonexistent (nonexistent meaning the seller values the car more than the buyer). Figure 7.B shows that the different mediators achieve a relatively similar probability of agreement when the reserve-auction gap is small. These low-gap cases likely reflect transactions where the surplus is large enough for any mediator, regardless of skill, to achieve agreement. Where mediator skill diverges is for higher reserve-auction gaps. Here, where the true gains from trade are likely smaller, and hence negotiators are less likely to naturally agree, high-skilled mediators substantially outperform medium-skilled mediators, who outperform low-skilled mediators. This raises the question of whether better mediators largely only capture additional *low-surplus* trades missed by other mediators. If so, better mediators may have only minimal effects on welfare, in spite of their increased trade probability. This question requires a model, to which we now turn.

[47]The reserve-auction gap is always a positive number for any observation that enters bargaining: the auction price being less than the reserve price is why the auction fails and bargaining begins.

5 The Effect of Mediators on Welfare

In this section, we construct a structural model to quantify how mediators' performance gaps affect buyer, seller, and total welfare. Our goal is to analyze whether mediators achieving a higher trade probability also achieve higher overall efficiency, or whether instead these mediators simply capture low-surplus trades (i.e., cases where the buyer values the car only slightly more than the seller, adding little to overall efficiency). For this analysis, we construct a structural model building on the incomplete-information, mechanism design framework for bilateral trade of Myerson and Satterthwaite (1983) and Williams (1987). Our empirical approach applies that of Larsen (2021), extended to allow for different mediators to index different equilibrium strategies in bargaining.

5.1 Conceptualizing Mediators as Implementing Different Equilibria. Mediators could potentially influence outcomes in a number of ways. To our knowledge, the theory literature offers no model of mediator *heterogeneity*, but we can extract possible theories that might generate heterogeneity from the literature that compares *mediated* to *unmediated* bargaining. For example, several theories discussed in Section 2 suggest that mediation can improve efficiency when a mediator elicits private information from each party and selectively discloses it (e.g., JPS, Goltsman et al. 2009, Hörner et al. 2015, Fanning 2021). Heterogeneity could therefore potentially arise if some mediators are better at information elicitation or optimal disclosure. Alternatively, heterogeneity might arise if some mediators are better informed than others about agents' values, a feature of Gottardi and Mezzetti (2024). Finally, in a setting like ours with multiple equilibria, different mediators could correspond to distinct equilibria, in the sense that equilibrium strategies in the bargaining stage of the game can depend on the mediator. For our model, we take the strong view that mediators influence outcomes only through the final channel; better mediators do not have (or extract) different information, but have the ability to coordinate agents on higher-trade equilibrium strategies. This structure turns out to be sufficiently flexible to capture large mediator effects on outcomes. This is because sequential-offer, incomplete-information bargaining games have a rich set (indeed, a continuum) of qualitatively different equilibria (Ausubel et al. 2002).[48] Even with a fixed bargaining protocol and rational agents, a backdrop where mediators implement different equilibria can yield a wide array of trade probabilities and welfare results across equilibria.

[48]This feature of multiple, qualitatively different equilibria is also discussed in Gul and Sonnenschein (1988) for sequential-offer bargaining games and in Satterthwaite and Williams (1989) with simultaneous offers.

In addition to its flexibility, this framework is empirically appealing due to several properties derived in Larsen (2021) that identify model primitives while remaining largely agnostic about equilibria. By the revelation principle (Myerson 1979), each equilibrium has a corresponding direct mechanism pinning down the probabilities with which agents trade. As we show, these different direct mechanisms are nonparametrically identified. The data also allow us to identify (or partially identify) agents' value distributions, which, combined with the direct mechanisms, are sufficient to quantify welfare under different mediators. Our identification arguments require data not only from transactions involving mediated bargaining, but also those that ended at other stages of the game. For example, we need auction and reserve price data for successful auctions. We use the estimated primitives (buyer and seller values and mediators' direct mechanisms) to estimate welfare solely within the bargaining stage of the game under different mediator skill levels.[49]

5.2 Model Setup. Our discussion follows Larsen (2021). Identification relies heavily on a basic model of the pre-bargaining (auction) stage. For each auction, we treat each bidder j as having a private value B_j drawn from F_B (with density f_B and support $(\underline{b}, \overline{b})$) and the seller as having a private value, $S \sim F_S$ (with density f_S and support $(\underline{s}, \overline{s})$), where B_j and S are independent for all j. When referring to the value of the highest bidder – who potentially enters bargaining with the seller – we use the notation B, without a j subscript. The number of bidders in a given auction is a random variable N. In addition to their own private values, agents' *total values* include a game-level heterogeneity component, $Y'\gamma + W$, where Y is a vector observable to the econometrician, W is a scalar unobserved by the econometrician, and γ is a parameter vector to be estimated. All agents in the game observe both Y and W.[50] We summarize our first set of assumptions:

(A1) $N \geq 2$ *risk-neutral bidders participate in an ascending button auction with zero participation costs. For* $j = 1, ..., N$*, where* N *is a random variable, each buyer* j *has a value* $Y'\gamma + W + B_j$*, with* $B_j \sim F_B$ *and* $W \sim F_W$*, and with* $(Y, W, N, \{B_j\}_{j=1}^N)$ *mutually independent.*

(A2) A risk-neutral seller has a value $Y'\gamma + W + S$*, with* $S \sim F_S$ *and with* S *independent of* $(Y, W, N, \{B_j\}_{j=1}^N)$*.*

(A3) Post-auction bargaining lasts for up to $T < \infty$ *periods; buyers incur a common bargaining*

[49] Larsen and Zhang (2018) studied the full auction-plus-bargaining game in this market, finding that increases in the number of bidders can improve efficiency, consistent with theoretical implications of Williams (1999).

[50] We can remain agnostic about whether N, the number of bidders, is observed or not by the agents, because, as we show below, bids do not depend on N, nor does the distribution of buyer values conditional on entering bargaining.

cost, $\eta_B > 0$ (and sellers $\eta_S > 0$), for each offer made. Buyers also incur η_B when first choosing to enter bargaining.

$A1$–$A2$ imply that B_j, S, N, Y, and W are mutually independent. We allow bidders' values to be correlated with one another and with sellers, but only through Y and W. This is consistent with evidence that an independent private values (IPV) model, or IPV with unobserved heterogeneity, fits well for used-car auctions; allowing for additional correlation — or for interdependent/common values — is beyond the current state of empirical methodologies for ascending auctions.[51] A seller's private value, S, describes how much above or below the game-level heterogeneity component $Y'\gamma + W$ her total value $(Y'\gamma + W + S)$ is, and similarly for buyers. Thus, agents' private values $\{B_j\}_{j=1}^N$ and S can take on positive and negative values, as can W. Below, we consider the game *conditional* on a realization of $Y'\gamma + W$, and agents' continuous actions (e.g., reserve prices, auction bids, bargaining offers) describe how far above or below $Y'\gamma + W$ these actions are. We omit $Y'\gamma + W$ from the discussion until Proposition 5, where we show that the game's equilibria are invariant to $Y'\gamma + W$. Appendix C contains additional model discussion.

We review the timing of the game, described in Section 3, in the context of our model. Let t represent the period of the game. At $t = 0$, the seller chooses a secret reserve price R. At $t = 1$, N bidders participate in an ascending auction (which we assume, for simplicity, follows a button auction format), resulting in an auction price P^A. If $P^A \geq R$, the seller and high-bidder trade. Otherwise, the high-bidder is given the opportunity to walk away or enter bargaining. If she enters bargaining, she incurs the cost η_B (by $A3$) for essentially making the first offer (an offer of P^A). Nature then moves, choosing κ, the mediator's identity. The realization of κ is immediately revealed to all agents, and the game continues to $t = 2$. We refer to $t = 2$ and onward as the *bargaining* stage of the game. At $t = 2$, the seller can accept, reject, or counter in response to the offer P^A. If the seller counters, the game continues to $t = 3$, the buyer's turn, and so on.[52] Bargaining lasts for up to T periods, which, by $A3$, is finite (but potentially large). The final bargained price is above P^A; this is the standard outcome in the game, and we treat it here as a requirement of the auction company (although in practice this is not explicitly imposed).

[51] Structural models of similar used-car auctions (but in Korea) in Roberts (2013) and Kim and Lee (2014) found evidence consistent with IPV. Earlier used-car auctions work (Genesove 1993) found weak evidence of a common value component at used-car auctions; sellers' information disclosure requirements have improved since then, arguably reducing potential adverse selection concerns that could arise in common values setting and rendering the environment of IPV with unobserved heterogeneity more palatable.

[52] We will use the terms "high bidder" and "buyer" interchangeably.

We allow κ to take values in $\{L, M, H, \emptyset\}$, standing for low, medium, or high skill levels; a value of $\kappa = \emptyset$ means that the negotiation stage took place without a mediator or that the negotiation failed to meet our sample criteria from Section 3.[53] The random variable κ indexes a particular set of equilibrium strategies of the bargaining stage of the game, and this is the sense in which different mediators correspond to different equilibria. At stages of the game before κ is realized, agents form expectations of their payoffs, integrating over the allocation rule associated with the strategies indexed by κ (as well as integrating over opponent types). We assume the following:

(A4) κ is independent of all random variables and actions realized before the bargaining stage.

Section 4 provides some evidence that mediator assignment is largely random conditional on observables. For our structural exercise, $A4$ is a stronger, unconditional independence assumption. Conditional independence would be weaker but would require continued conditioning on observables — or an index of observables — in subsequent estimation steps, greatly increasing the computational burden. As described below, our approach to handling game-level observables involves a first estimation step residualizing actions against a large set of observables (a standard "homogenization" approach in the auction literature; Haile et al. 2003.)

Payoffs are as follows. If trade occurs at a price P, a buyer with value B_j receives $B_j - P$ and the seller receives P. If trade fails, the buyer receives 0 and the seller receives S (her private value). Agents also face any incurred bargaining costs. These costs, described in $A3$, consist of a disutility η_B that the buyer pays for each period she chooses to continue the game (by making an offer or by initially entering bargaining). The seller incurs η_S for each offer she makes.

Let $\rho(S) = R$ be the seller's secret reserve price strategy. Let ζ_j be the drop-out price of bidder j in the auction. Let $\pi^B(P^A, B) - \eta_B$ be the high bidder's expected payoff of entering bargaining when $R > P^A$. The buyer computes the expected benefit $\pi^B(P^A, B)$ prior to knowing the mediator's type, and hence the expectation integrates over κ. Let $D_1^B = 1$ represent the buyer's decision to walk away when informed that $P^A < R$ (and $D_1^B = 0$ means not walking away). Let $D_t^B \in \{a, c, q\}$ be the buyer's decision to accept, counter, or quit at odd $t > 1$. Let P_t^B be the buyer's counteroffer (if the buyer counters) in period t. Let D_t^S and P_t^S be defined similarly for even periods t. Let

[53] As described in Section 3, a negotiation can occur without a mediator if the seller is physically present at the auction sale and accepts or rejects P^A before any mediator involvement. It is also possible for the auctioneer to reject P^A on behalf of the seller if the seller is absent. A negotiation falls outside of our Section-3 sample criteria, for example, if it is handled by multiple mediators or by a mediator who handled fewer than 50 negotiations (meaning we lack data to classify the mediator's skill level).

H_t be the history of publicly observed actions up through period $t-1$ of the game. These actions include P^A and all previous bargaining offers and period-specific decisions.

We focus on pure-strategy PBE of this game.[54] We do not allow for equilibria to condition on non-payoff-relevant variables, including Y, W, and N; Propositions 3 and 5 below discuss these variables further. The strategy of a buyer of type b_j is $\sigma_j^B(b_j,\kappa) = \{\zeta_j, \{D_t^B|H_t\}_j, \{P_t^B|H_t\}_j|b_j,\kappa\}$, where the decisions D_t^B and offers P_t^B included are those for periods in which it is the buyer's turn. The strategy of a seller of type s is a history-contingent set of actions $\sigma^S(s,\kappa) = \{\rho, \{D_t^S|H_t\}, \{P_t^S|H_t\}|s,\kappa\}$.[55] Actions that occur after mediator assignment is revealed are allowed to depend on κ. A set of strategies $\{\sigma^{B*}(b_j,\kappa)\}_{j=1}^N$ for all buyer types and $\sigma^{S*}(s,\kappa)$ for all seller types, along with buyer beliefs $\mu_{B*}(S|H_t,b,\kappa)$ and seller beliefs $\mu_{S*}(B|H_t,s,\kappa)$, constitutes a PBE if, for each player, her strategy is a best response to opponents' strategies and sequentially rational given her beliefs, and beliefs are consistent with strategies and updated using Bayes rule on path; off-path beliefs are unrestricted.[56]

The following examples (from Larsen and Zhang 2026) illustrate four of the many different possibilities of play that could occur in equilibrium in the bargaining stage (beginning with $t=2$):

***Example 1**: At $t=2$, sellers counter at a price equal to the take-it-or-leave-it offer (TIOLIO) that is optimal conditional on the current history of play. At $t=3$, buyers accept or quit. At $t=4$, (only reached off path) sellers quit.*

***Example 2**: At $t=2$, sellers accept (if $S \leq P^A$) or counter at $\bar{b}$ (an uninformative offer). At $t=3$, buyers quit if facing any (off-path) offer other than $\bar{b}$; otherwise, they counter at the optimal TIOLIO. At $t=4$, sellers accept or quit. At $t=5$ (off path), buyers quit.*[57]

***Example 3**: At $t=2$, sellers accept if $S \leq P^A$ 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 δP^A (with $\delta > 1$ fixed and known to all parties), whichever yields a higher expected payoff. At $t=3$, buyers accept or quit.*

[54] Mixed-strategy equilibria may also exist; we focus on pure strategies to facilitate the model properties we exploit for identification while still being sufficiently flexible to admit a wide range of behavior.

[55] Note that, throughout the paper, for random variables that can take on *infinitely many* values (e.g., Y, W, N, S, B_j, R, P^A), we use corresponding lowercase letters to denote realizations. We use v as a generic argument in various functions.

[56] Because bidders are symmetric in the model, and only one bidder continues in the game past the bargaining stage, an equilibrium only needs to specify the seller's beliefs about the highest bidder and vice versa. Also, while our empirical approach does not require fully specifying an equilibrium (its strategies or beliefs), we note that, in sequential bargaining games of incomplete information, a wide range of behavior can be sustained as PBE through the choice of the belief-updating rule after off-path actions. For example, many theoretical models use *optimistic conjectures* in which, after an off-path action, the opponent believes she faces the weakest opponent type (the highest buyer type or lowest seller type). See discussions in Rubinstein (1985) and Gul and Sonnenschein (1988), for example.

[57] Relekar and Ren (2025) analyzed this equilibrium in detail, illustrating many of the model properties derived in Section 5.3 and offering additional insights.

At $t = 3$, buyers quit if facing any (off-path) price other than δP^A. At $t = 4$ (off path), sellers quit.

5.3 Model Properties. We present several properties that help identify primitives.

Buyer Values. We make the following assumption on bidding behavior:

(A5) Bidders bid truthfully in the auction; that is, each bidder j chooses ζ_j equal to her value, b_j.

The following demonstrates that any behavior other than truthtelling is weakly dominated:

Proposition 1. *Suppose A1 and A3 hold. For any arbitrary bidder j with value b_j, any strategy that involves $\zeta_j \neq b_j$ is weakly dominated by a strategy that involves $\zeta_j = b_j$.*

The proof of Proposition 1 shows that the key model features for this result are (i) if a bidder learns $P^A < R$ and she is the high bidder, she can costlessly opt out of bargaining, eliminating any downward bid shading; and (ii) the bargaining price is weakly greater than P^A, eliminating incentive to shade auction bids upward.

Under $A5$, the distribution of auction prices, F_{P^A}, is equal to the distribution of the second-highest order statistic of buyer values, and hence F_B is related to F_{P^A} and $\Pr(N = n)$ (the probability mass function of N) as follows, evaluated at any number v:

$$F_{P^A}(v) = \sum_n \Pr(N = n) \left[n F_B(v)^{n-1} - (n-1) F_B(v)^n \right] \tag{5}$$

For any v, $F_{P^A}(v)$ can be estimated from data on auction prices. The right-hand side of (5) is monotonic in $F_B(v)$, and thus $F_B(v)$ is identified.[58]

Seller Values. Let $D^S = a$, without a t subscript (to distinguish this from the period-specific action described in Section 5.2), represent the event in which the seller takes an action at $t = 1$ or 2 that results in the game ending in agreement at P^A. This event occurs either when (i) $P^A \geq R$; or (ii) $P^A < R$, the high bidder does not opt out of bargaining, and the seller accepts the auction price on her first bargaining turn. Similarly, let $D^S = q$ represent the event in which the seller takes an action at $t = 2$ that results in the game ending in disagreement in that period. This event happens when $P^A < R$, the high bidder does not opt out of bargaining, and the seller rejects P^A or quits on her first bargaining turn (rather than accepting P^A or making a counteroffer).[59]

[58] Equation (5) can be equivalently written $F_{P^A}(v) = \sum_n \Pr(N = n) \left[F_B^n(v) + n \left(1 - F_B(v)\right) F_B^{n-1}(v) \right]$.

[59] For purposes of bounding F_S, we treat cases where $P^A < R$ and the auctioneer rejected P^A on behalf of the seller as cases where the seller herself rejected P^A.

In any equilibrium, the seller will not accept P^A if $P^A < S$ (the seller would be better off quitting and getting S instead of P^A) and will not quit if $S < P^A$ (the seller would be better off accepting, receiving P^A instead of S). This implies the following bounds:

Proposition 2. *Under A1–A5, for any* $v \in (\underline{s}, \overline{s})$, $F_S(v) \in [E_\kappa\{\Pr(D^S = a|P^A = v, \kappa)|P^A = v\}, E_\kappa\{\Pr(D^S \neq q|P^A = v, \kappa)|P^A = v\}] = [\Pr(D^S = a|P^A = v), \Pr(D^S \neq q|P^A = v)]$.

The notation E_κ indicates an expectation taken over mediator assignment, which yields bounds on F_S because, by A1–A5, $\Pr(S \leq v|P^A = v, \kappa) = \Pr(S \leq v)$. The property underlying these bounds is similar to that of the English-auction bounds in Haile and Tamer (2003) but applied to the seller's decision to accept or reject the first bargaining offer (P^A). As explained above, period 2 may or may not involve a mediator interaction. We do not use mediator assignment in estimating seller bounds but instead estimate bounds that average over mediator assignment (averaging over both whether a mediator is involved beginning in period 2 and, if so, which mediator is assigned). This allows us to obtain a single set of bounds on F_S and to focus on the different equilibria/mechanisms mediators implement, which we address below. See Appendix C.9 for additional discussion.

Belief Updating After Auctions. If $P^A < R$, the buyer chooses to enter bargaining if $\pi^B(P^A, B) - \eta_B \geq 0$. When the realized auction price is $P^A = p^A$, let $\chi(p^A)$ be defined as follows:

$$\chi(p^A) \equiv \begin{cases} \overline{b} & \text{if } \ \pi^B(p^A, b) - \eta_B < 0 \ \ \forall b \in (\underline{b}, \overline{b}) \\ \chi' \ \ s.t. \ \ \pi^B(p^A, \chi') - \eta_B = 0 & \text{otherwise} \end{cases} \tag{6}$$

In words, when $P^A = p^A$, $\chi(p^A)$ is the lowest buyer type who enters bargaining, unless the payoff from entering bargaining is negative for all buyer types, in which case $\chi(p^A) = \overline{b}$. When bargaining ensues, both the buyer and seller are aware that the support of buyer types who would enter bargaining is truncated: $B \geq \chi(P^A)$. Similarly, both parties know that, because $P^A < R$, S is truncated: $P^A < \rho(S)$ (recall that $\rho(S) \equiv R$.) Proposition 3 below states this result, along with a monotonicity property satisfied by ρ that relies on two additional assumptions:

(A6) In any equilibrium, the seller's expected bargaining payoff is continuous in the auction price.

(A7) The support of B *is* $(-\infty, \infty)$*, and the density* f_B *is positive on this support.*

*A*6 is a direct assumption on an equilibrium object that we make in order to ensure differentiability of the seller's payoff, allowing us to apply the Edlin and Shannon (1998) theorem to prove

strict monotonicity of $\rho(\cdot)$.[60] $A7$ is invoked in the proof of Proposition 3 to avoid dividing by zero and to ensure strict monotonicity of $\rho(\cdot)$.

Proposition 3. *(i) Under $A1$–$A5$, $\rho(\cdot)$ is weakly increasing and, conditional on $P^A = p^A$, $\chi(p^A)$ exists and the support of buyer types in bargaining is $[\chi(p^A), \overline{b})$. (ii) If, in addition, $A6$–$A7$ hold, $\rho(\cdot)$ is strictly increasing, and, conditional on $P^A = p^A$, the support of seller types in bargaining is $[\rho^{-1}(p^A), \overline{s})$.*

In addition to $A6$–$A7$, a key to obtaining strict monotonicity of $\rho(\cdot)$ is a buyers' ability to opt out of bargaining after the auction if $\pi^B(P^A, B) - \eta_B < 0$. This disciplines sellers to set finite reserve prices. By Proposition 3, the densities of seller and buyer types entering bargaining conditional on $P^A = p^A$ are $f_{S|P^A}(s|p^A) \equiv \frac{f_S(s)}{1-F_S(\rho^{-1}(p^A))}$ and $f_{B|P^A}(b|p^A) \equiv \frac{f_B(b)}{1-F_B(\chi(p^A))}$.[61] These reflect Bayes-updated beliefs at the beginning of the bargaining stage. Proposition 3 implies

$$\Pr(D_1^B = 0 | P^A = p^A, P^A < R) = \frac{1 - F_B(\chi(p^A))}{1 - F_B(p^A)}, \tag{7}$$

an expression that we use below in estimating $\chi(\cdot)$.

Direct Mechanism. We consider different mediator types κ to index different equilibrium strategies of the bargaining stage of the game. By the revelation principle (Myerson 1979), any equilibrium has an equivalent *direct mechanism* in which agents truthfully report types to a designer who ensures outcomes occur as if the original game had played out. A mechanism is described by an *allocation function* specifying the probability with which a given seller and buyer trade and a *transfer function* specifying transfers between agents. For our purposes, we do not need to identify the entire transfer function, only the expected price of consummated trades, which is observable in the data. Identifying the allocation function, however, is essential and requires more work.

The allocation function for a mediator handling negotiations between buyer and seller types b and s, when $P^A = p^A$, can be written $x_\kappa(s, b; p^A)$, as the mechanism can depend on mediator type

[60] Equilibrium examples 1–4 above have no built-in discontinuity that violates $A6$ (and Relekar and Ren 2025 showed this explicitly for Example 2), but it is possible to conceive of equilibria in which $A6$ could be violated. For example, consider a modified version of Example 4 in which, when $P^A < 0.5$, the only accepted offer on path is δP^A and, when $P^A \geq 0.5$, the only accepted offer on path is $\delta' P^A$, where $\delta \neq \delta'$.

[61] The following is a property of order statistics of n independent draws of a random variable B: the density of the highest order statistic, evaluated at a point b and *conditional* on the second-highest order statistic being equal to $p^A < b$, is given by $f_B(b)/(1 - F_B(p^A))$. This is the simplest special case of a property derived in (and used for identification of values in auctions by) Song (2004), who showed that the density of a higher-ranked order statistic conditional on a lower one does not depend on n. In our setting, because of the buyer's option to opt out of bargaining, the conditional density of buyer types faced by the seller in bargaining is instead $f_B(b)/(1 - F_B(\chi(p^A)))$.

κ and P^A. This is because outcomes in bargaining can depend on P^A, as P^A is the first bargaining offer, affecting belief updating. We assume the following property for this allocation function:

(A8) In any equilibrium, for any seller type s, if a buyer type b trades with s on path, all buyer types $b' > b$ also trade with s on path.

We do not believe this assumption is inordinately strong in our setting. The property is known to hold in many bargaining models of incomplete information, such as Myerson and Satterthwaite (1983) and Williams (1987). Ausubel and Deneckere (1993) referred to this property as the "Northwestern Criterion," as it implies that trade occurs if and only if players' types lie northwest of a particular boundary.[62] Building on Theorem 1 of Storms (2015), we prove that a sufficient condition for $A8$ to hold is that S and B have finite type spaces.[63]

Proposition 4. *Under $A1$–$A3$, S and B having finite type spaces implies $A8$.*

Proposition 4 implies that $A8$ holds for arbitrarily fine type spaces; we assume directly that it holds in our setting. $A8$ implies that the allocation function for any mediator type κ takes an indicator-function form, $x_\kappa(r, b; p^A) \equiv 1\left\{b \geq g_\kappa\left(r, p^A\right)\right\}$, where $g_\kappa(r, p^A)$ is weakly increasing in r. In writing the allocation function this way we also incorporate the result from Proposition 3 that $\rho(\cdot)$ is strictly increasing, and thus we define the allocation function as a function of the realized reserve price $R = r$ rather than s.[64] Proposition 4 implies that the function g_κ in $1\{b \geq g_\kappa(r; p^A)\}$ is related to the conditional probability of agreement as follows:

$$\Pr(\mathcal{A}|R = r, P^A = p^A, \kappa) = \frac{1 - F_B\left(g_\kappa(r, p^A)\right)}{1 - F_B\left(p^A\right)}. \tag{8}$$

where $\mathcal{A}$ is the event that trade occurs. The denominator accounts for the fact that any buyer in bargaining has $B \geq P^A$. The conditional probability on the left of (8) is the key to identifying g_κ: because F_B is identified by (5), at a given $P^A = p^A$ and $R = r$, the quantile of $(1-F_B(\cdot))/(1-F_B(p^A)$ that matches the conditional trade probability reveals the lowest buyer type that trades with a seller who reports a reserve of $R = r$. This lowest buyer type is $g_\kappa(r, p^A)$.

[62] Note that, in a more general framework of mixed strategies, equilibria could exist that correspond to random allocation functions that yield trade with some probability in $(0, 1)$.

[63] The restriction to finite type spaces facilitates a proof that $A8$ holds by allowing us to condition on elements of finite sets rather than positive probability subsets of infinite sets.

[64] This step is valuable because the econometrician observes information about R, whereas S is private information.

Game-level Heterogeneity. The above discussion describes the game conditional on a realization of the game-level heterogeneity. Let the random variable Z denote the game-level heterogeneity: $Z \equiv Y'\gamma + W$. We prove the following:

Proposition 5. *Strategies and beliefs that constitute an equilibrium at a realization z also constitute an equilibrium at a different realization z' for any z and z'.*

To elaborate on this result, consider a realization $Z = z$. We will refer to continuous actions (reserve prices, auction prices, and bargaining offers) as *markup actions* because they describe how far above or below z these objects are. For any markup action, we will define its corresponding *total* action as z plus the markup action. For example, consider an equilibrium in which the realized markup auction price is p^A and the realized markup reserve price is r; the total auction price is then $p^A + z$ and the total reserve price is $r + z$. Similarly, for a buyer of type b and seller of type s, we will refer to their *total values* as $b + z$ and $s + z$. Now consider a different realization of $Z = z'$, where the total values of these same buyer and seller types would be $b + z'$ and $s + z'$. Proposition 5 implies that, if the same equilibrium is played under $Z = z$ and $Z = z'$, total auction prices and total reserve prices will be $p^A + z'$ and $r + z'$, and thus the different realizations of Z shift total *actions* by the same amount by which they shift total *values*, leaving the equilibrium mapping between values and markup actions unchanged.

For a concrete example, suppose the realization of the game level heterogeneity Z is $z = \$10{,}000$. Consider a seller with realized type $s = -\$2{,}000$, so her total value is $z + s = \$8{,}000$ (i.e., \$2,000 below z). Suppose that, in equilibrium, she chooses a reserve price markup $r = -\$1{,}000$, meaning a total reserve price of $z + r = \$9{,}000$ (i.e., \$1,000 below z). Proposition 5 implies that, in this equilibrium, if instead $Z = z' = \$5{,}000$, this same seller type s would choose the same reserve markup $r = -\$1{,}000$; hence her total reserve price becomes $z' + r = \$4{,}000$, again \$1,000 below the game-level heterogeneity.

Proposition 5 is an equilibrium translational invariance property, similar in spirit to other *homogenization* results (Haile et al. 2003) but applied here to this specific case of a secret reserve auction followed by bargaining. This property allows us to parse out game-level heterogeneity, which we do for *observable* heterogeneity ($Y'\gamma$) using a linear regression and for *unobserved* heterogeneity (W) via a deconvolution exercise, as we describe below.

5.4 Evaluating Gains from Trade. An upper bound on the expected gains from trade in the

bargaining stage of the game, under allocation function $x_\kappa(r, b; p^A)$, can be evaluated as

$$\int_{\underline{b}}^{\overline{b}} \left[\int_{\chi(p^A)}^{\overline{b}} \int_{\rho^{-1}(p^A)}^{\overline{s}} (b - s) x_\kappa(\rho(s), b; p^A) f_{S|P^A}(s|p^A) f_{B|P^A}(b|p^A) ds\, db \right] f_{P^A}(p^A) dp^A \tag{9}$$

(9) is an upper bound on the expected trade gains because it does not incorporate bargaining costs (essentially treating them as 0); a corresponding lower bound is obtained by subtracting from (9) a bound on the loss due to bargaining costs, described in Appendix C.7. We evaluate these gains separately for low-, medium-, and high-skilled mediators. As a counterfactual benchmark, we also compute the infeasible first-best mechanism, which is simply $1\{B \geq S\}$ — trading whenever gains from trade exist — and the second-best mechanism.[65] The Myerson-Satterthwaite Theorem showed that a gap exists between the second- and first-best; the latter is infeasible because each party has incentives to retain information rents. Larsen (2021) showed that inefficiency in the real-world mechanism is even larger than this gap, as the real-world outcome lies below even the second-best. Our setting allows us to unpack some of this shortfall to examine whether certain mediators are better at reducing this inefficiency.

Expression (9) shows that the key objects required for evaluating the trade gains are the allocation functions for each mediator type, $x_\kappa(r, b; p^A)$; the conditional densities $f_{S|P^A}(s|p^A)$ and $f_{B|P^A}(b|p^A)$; the truncated support bounds for the types who enter bargaining, $\chi(p^A)$ and $\rho^{-1}(p^A)$; and the density of auction prices f_{P^A}. The truncated supports allow us to compute welfare solely within the bargaining stage of the game, which is the only stage mediators can influence.

5.5 Estimation Sample and Steps. Our estimation relies on a dataset formed by merging our indicators of mediator type from the mediated negotiation data with an auction-plus-bargaining dataset of 264,996 observations from Larsen (2021). The latter only partially overlaps with the mediated negotiation data: some observations in the auction-plus-bargaining data end at the auction stage and have no mediator identifiers, and some observations in the mediated negotiation data are missing key variables found in the auction-plus-bargaining data.

The merge yields the following number of observations handled by mediators whom we can identify as low, medium, or high-skilled, respectively: 11,764 (denote this set by $\mathcal{N}_L$), 10,235 ($\mathcal{N}_M$),

[65] Appendix Figure C.3 shows the estimated welfare under different mediators relative to the full Pareto frontier of Williams (1987), of which the Myerson-Satterthwaite second-best is the highest-utility point. We compute these objects following the linear programming problem described by Myerson and Satterthwaite (1983) and Williams (1987), without enforcing distribution regularity.

10,043 ($\mathcal{N}_H$). In 36,400 observations, a mediator identity is recorded but the negotiation does not meet our sample criteria from Section 3; e.g., the mediator handles too few observations. We denote this set $\mathcal{N}_\omega$. Let $\mathcal{N}_A$ denote the set of 196,524 observations that ended at the auction stage, ended with the buyer opting out of bargaining, or ended through bargaining but with *no* mediator recorded. Let $\mathcal{N}_B$ be the union of $\mathcal{N}_L$, $\mathcal{N}_M$, $\mathcal{N}_H$, and $\mathcal{N}_\omega$ (and thus $\mathcal{N}_A \cup \mathcal{N}_B$ is the full dataset).

The model properties yield identification arguments and corresponding estimation approaches. To conserve space, we summarize estimation steps; Appendix C contains details. Steps 1–5 do not require conditioning on mediator type, and these steps follow Larsen (2021) closely. Steps 6–7 involve estimation separately by mediator type.

1. **Observed Heterogeneity ($Y'\gamma$).** By Proposition 5, auction and reserve prices are additively separable in game-level heterogeneity, $Y'\gamma + W$. We estimate γ via a stacked linear regression of auction and reserve prices on a rich set of game-level observables Y. Residuals from this regression correspond to $\tilde{R} \equiv R + W$ and $\tilde{P}^A \equiv P^A + W$. This step uses $\mathcal{N}_A \cup \mathcal{N}_B$, in addition to observations in which auction or reserve prices are missing (but not both). This improves estimation of γ given the many controls in Y.

2. **Unobserved Heterogeneity (W).** By Proposition 5, the residuals from Step 1 are additively separable in W. By independence of R, P^A, and W, a convolution argument implies that the marginal densities f_R, f_W, and f_{P^A} are identified from the joint distribution of $\tilde{R}$ and $\tilde{P}^A$.[66] We estimate these marginal densities f_{P^A}, f_R, and f_W and their corresponding CDFs using a flexible maximum likelihood approach. This step (and Steps 3–4) uses $\mathcal{N}_A \cup \mathcal{N}_B$.

3. **Buyer Values.** By $A5$ and Proposition 1, we can use the CDF F_{P^A}, estimated in Step 2, along with $\Pr(N = n)$, to identify and estimate F_B by directly solving (5) on a grid of values.

4. **Seller Values.** Ignoring W, the lower and upper bounds on F_S in Proposition 2 are probabilities conditional on P^A. Accounting for W, the objects we can estimate in the data condition instead on $\tilde{P}^A \equiv P^A + W$. These objects are $\Pr(D^S = a|\tilde{P}^A)$ and $\Pr(D^S \neq q|\tilde{P}^A)$, which we estimate via local linear regressions of the events $D^S = a$ and $D^S \neq q$ on $\tilde{P}^A$. For the F_S lower bound, the theoretical counterpart is a convolution of $\Pr(D^S = a|P^A)$ with f_W. We use the convolution expression to estimate a lower bound on F_S, parsing out W using a flexible

[66] A similar convolution argument was first applied to auction-level heterogeneity by Krasnokutskaya (2011).

parameterization (piecewise linear splines) and minimizing the distance between the estimate of $\Pr(D^S = a|\tilde{P}^A)$ and its theoretical counterpart. The upper bound is similar.

5. **Belief Updating (Functions $\rho(\cdot)$ and $\chi(\cdot)$).** For any function $F_S(\cdot)$ lying within the estimated bounds, $\rho(\cdot)$ can be constructed as $\rho(s) = F_R^{-1}(F_S(s))$, with F_R replaced with $\hat{F}_R$ from Step 2. To estimate $\chi(\cdot)$, we modify (7) to incorporate unobserved heterogeneity W, which becomes another convolution, and minimize the distance between the left and right-hand sides of (7), parameterizing $\chi(\cdot)$ via piecewise linear splines. As indicated by the conditional probability in (7), this step conditions on observations from $\mathcal{N}_A \cup \mathcal{N}_B$ in which $R > P^A$.

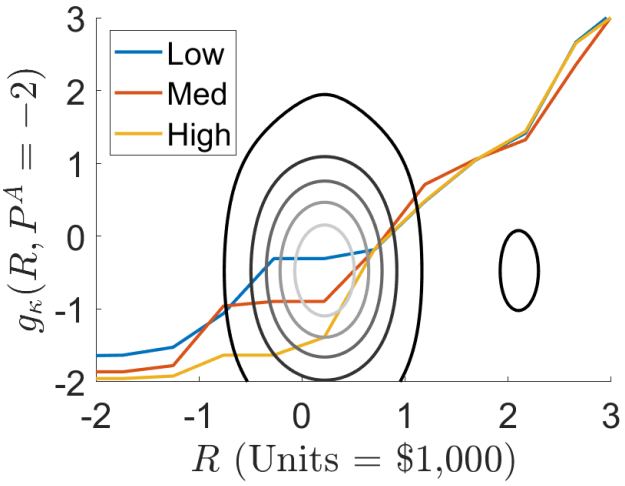
6. **Direct Mechanisms.** To estimate each $g_\kappa(\cdot)$, we modify (8) to include W – another convolution – and minimize the distance between the left and right-hand sides of (8), parameterizing $g_\kappa(\cdot)$ via bilinear splines. In estimation, we impose the constraint that $g_\kappa(\cdot)$ cannot imply trade when $S > B$. We create appropriate pseudo samples to estimate these g_κ functions separately by mediator type. To estimate g_L, for example, we create a low-skilled mediators pseudo sample by combining $\mathcal{N}_A$ with $|\mathcal{N}_B|$ observations drawn with replacement from the set $\mathcal{N}_L$. We create similar pseudo samples for medium- and high-skilled mediators. This allows each g_κ function to be estimated with a dataset of the same size and same proportion of observations ending at the auction or bargaining stages.

7. **Expected Trade Gains.** We evaluate the upper bound on expected gains from trade using (9) and obtain a lower bound by subtracting from this quantity a bound on the loss due to bargaining costs. This loss bound is estimated separately using information contained in the gap between bargaining offers. For example, a buyer making an offer of P_3^B must prefer this action (and its accompanying bargaining cost, η_B) to accepting the seller's most recent offer, P_2^S, implying that $B - P_3^B - \eta_B \geq B - P_2^S$. Rearranging yields bounds on η_B. An upper bound on the buyer's total expected loss due to bargaining costs is constructed by multiplying η_B by the expected number of bargaining rounds. A similar argument applies to the seller. We estimate these bargaining cost bounds separately by mediator type.

5.6 Estimated Allocation Threshold Functions by Mediator Type. The estimated $g_\kappa(R, P^A)$ functions for the three mediator types are shown in Figure 8, with units in \$1,000. This figure contains similar information to the model-free reserve-auction gap plot in Figure 7.B, but offers

additional insights through the lens of the structural model, which allows us to plot and interpret a function of values (e.g., B) rather than just an object related to values (P^A). To aid visualization, we plot R on the horizontal axis and B on the vertical, and hold P^A fixed at \$2,000 below the game-level heterogeneity component, $Y'\gamma + W$. After subtracting the heterogeneity components, because heterogeneity enters additively, values and prices (such as R, P^A, and B) will range from negative to positive. For example, a reserve price of $R = -1$ in Figure 8 means \$1,000 below the car's game-level heterogeneity component.[67] Each mechanism yields trade if B lies above a given $g_\kappa(R, P^A)$ line. The contour lines show the amount of mass under the joint pdf of B (conditional on $B > P^A$) and R, with lines representing where 1/6, 2/6, ..., 5/6 of the mass lies.[68]

Figure 8: Direct Mechanisms $g_\kappa(\cdot)$ Functions for Different Mediator Types (Units = \$1,000)



Notes: Figure displays estimates of $g_\kappa(R, P^A)$ for low- (blue), medium- (red), and high-skilled (yellow) mediators, where the auction price is held at a value of \$2,000 below the heterogeneity value ($Y'\gamma + W$) for the car; that is, $P^A = -2$, as units are \$1,000. Trade occurs if $B \geq g_\kappa(R, P^A)$. Contour plot shows the amount of mass under the joint pdf of B (conditional on $B > P^A$) and R, representing where 1/6, 2/6, ..., 5/6 of the mass lies.

Figure 8 shows a clear ordering of mediator types: low-skilled mediators implement a higher threshold function, consummating fewer efficient trades, than do medium-skilled mediators, who in turn consummate fewer efficient trades than high-skilled mediators (the lowest function verti-

[67] Because of this additive separability, it is possible for our estimates to imply total reserve prices $Y'\gamma + W + R$ or total auction prices $Y'\gamma + W + P^A$ that are negative, which can occur if the realization of $Y'\gamma + W$ is not sufficiently positive and R or P^A is sufficiently negative. Appendix C.2 shows that this rarely occurs for our estimates.

[68] Because the estimated F_R is multimodal, the widest contour line displays an additional amount of mass to the right of where the majority of the mass lies. Note that, because R and B denote how far above or below $Y'\gamma + W$ these actions/values are, the contour sets show mass both above and below zero.

cally).[69] This ordering implies that a higher-skilled mediator captures trades among (approximately) the same buyer-seller pairs as a lower-skilled mediator, plus additional, positive-surplus trades, foreshadowing the result we examine in detail in the following section, that higher-skilled mediators increase the overall *gains from trade*. At a reserve price equal to the game-level heterogeneity component (shown as zero on the horizontal axis), a low-skilled mediator only successfully closes a deal for buyers who have a value of at least -$307 (i.e., $307 below the heterogeneity component). Faced with the same seller (i.e., same reserve price), a medium-skilled mediator is able to generate trades for buyers with values approximately above -$896, whereas a high-skilled mediator captures trades even for lower buyer values, when the buyer's value exceeds about -$1,633.

5.7 Welfare Results. Because F_S is only partially identified, we obtain bounds on trade gains. We first evaluate bounds on the difference in trade gains from (9) between a given real-world mediator's mechanism and the first- or second-best mechanism, for the average negotiation.[70] Figure 9.A shows, in blue, bounds on the difference between the first-best and real-world outcomes separately for negotiations mediated by low-, medium-, or high-skilled mediators.[71] These represent bounds on the deadweight loss in bargaining — gains that real-world bargaining fails to capture but that would be realized in a first-best world because they correspond to a buyer valuing the car more than the seller. A larger number in this figure thus represents a larger loss relative to the first-best efficient outcome. Yellow bars show similar bounds comparing to the second-best outcome. In Figure 9.A, under each mediator type, the outcome is slightly closer to the second-best than to the first-best; this is because the second-best falls short of the first-best (the inefficiency described by Myerson

[69] The key object in the data that identifies a g_κ function is the probability of trade, for a given mediator type κ, conditional on the reserve and auction price, as shown in (8). A higher conditional probability of trade implies a lower g_κ function. In Appendix C.6, we show that the ordering of the g_κ functions in Figure 8 is driven by these conditional trade probabilities. We also show that the close alignment of the g_κ functions in Figure 8 at high values of R is driven by the constraint we impose in estimation Step 6 that the g_κ functions cannot imply trade if $S > B$.

[70] Proposition 6 of Larsen (2021) showed that, across all possible F_S within the seller CDF bounds, an upper bound on the gap between trade gains in the first-best and in the real-world bargaining is given by the difference between the upper bound of (9) evaluated at the upper bound seller CDF and the lower bound of (9) evaluated at the lower bound seller CDF. A lower bound on the trade gains difference is given analogously. This same bounding property applies to the gap between trade gains under any two different mediator types (high- vs. low-skill, say). For consistency, we follow this same approach (evaluating trade gains at the lower bound and upper bound seller CDFs) when evaluating the *gap between the real-world and second-best* outcomes, but, unlike the gap between the real-world and *first-best* outcomes, we have no theoretical result that this yields a bound on the difference across all seller CDFs. The qualitative implications from examining first- vs. second-best mechanisms (relative to any mediator mechanisms) are similar because the first- and second-best are similar in this market; see Larsen (2021).

[71] Dashed lines show 95% bootstrapped confidence sets for these bounds. Confidence bounds throughout this section are constructed via a nonparametric percentile bootstrap with 200 replications of our estimation procedure. We report the 0.025 quantile of the estimated lower bound on an object of interest and the 0.975 quantile of the estimated upper bound; these are conservative by a Bonferroni-style argument.

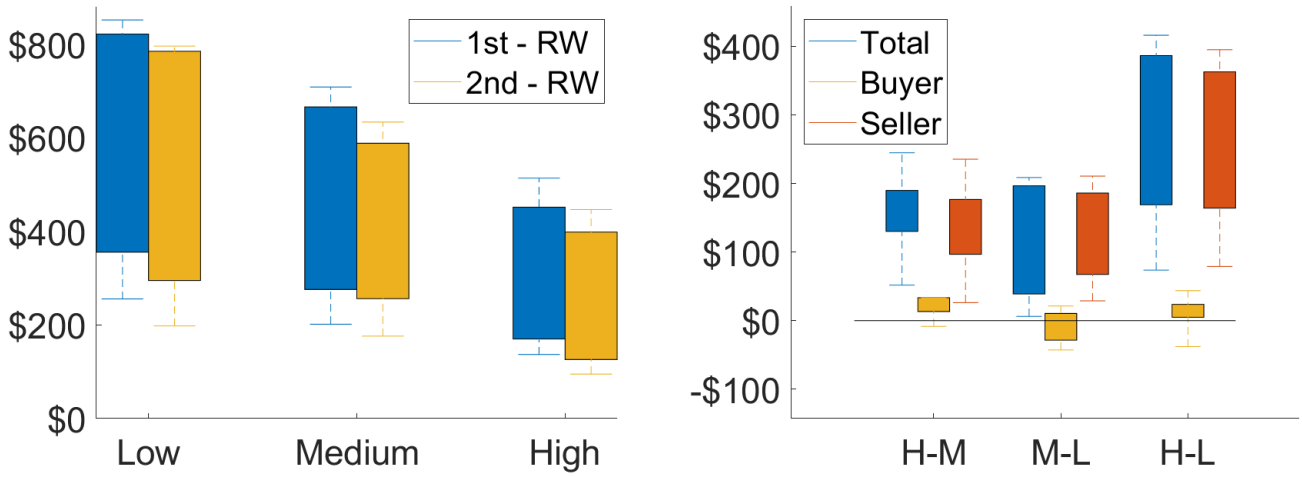
and Satterthwaite 1983).

We find that low-skilled mediators fall short of the first best by \$356–823. This is a sizeable gap relative to the total first-best surplus (not shown in Figure 9.A; we estimate bounds on this quantity of \$2,602–4,459). The losses in Figure 9.A can also be compared to the average difference between the auction price and reserve price, which can be calculated without the structural model; this quantity is \$1,775 (see Table 1). For high-skilled mediators, the first-best vs. real-world surplus gap in Figure 9.A is significantly smaller, and we can reject the largest efficiency shortfalls — those greater than \$514, which is the upper confidence bound on the gap. We fail to reject these levels for low-skilled mediators.

Figure 9: Surplus Differences Across Mediators

(A) Real-world vs. 1st/2nd Best

(B) Real-world Diff. Across Mediators



Notes: Panel A displays, in solid blue, bounds on the difference in total surplus between first-best (1st) and real-world bargaining (RW), separately for threads mediated by low, medium, and high-skilled mediators; and solid yellow displays bounds on the difference between the second-best (2nd) and real-world outcomes. Panel B displays, in solid blue, bounds on the difference in total surplus in the real-world mechanism between high vs. medium-skilled mediators (H-M), medium vs. low-skilled (M-L), and high vs. low-skilled (H-L); solid red lines display bounds on the difference in seller surplus; and solid yellow lines display bounds on the difference in buyer surplus. All bounds are computed by evaluating differences at the upper and lower bounds on F_S. Whiskers show 95% confidence intervals on these bounds constructed via bootstrapping with 200 replications.

Our primary result is in Figure 9.B, where we show realized surplus differences for the average negotiation mediated by high- vs. medium-skilled (H-M), medium- vs. low-skilled (M-L), and high- vs. low-skilled (H-L) mediators. These total surplus differences are in solid blue. Solid yellow and red bars show buyer and seller surplus separately. We observe a statistically significant difference in

total surplus for all three pairwise comparisons, with a surplus increase of about $130–190 between high- vs. medium-skilled mediators, $39–197 for medium- vs. low-skilled, and $169–386 for high- vs. low-skilled. This implies that mediators who achieve higher trade probabilities also achieve a statistically significant increase in total welfare. Relative to the shortfall from Figure 9.A, the quantities in Figure 9.B represent a meaningful efficiency improvement.

Figure 9 also offers a reasonable way to break down the limits of the real-world mechanism itself vs. the limits of the mediators who implement the mechanism. The former can be proxied by how far high-skilled mediators are below the second-best in Figure 9.A — $126 at best — while the latter is captured by how far medium- or low-skilled mediators are below high skilled mediators in Figure 9.B, discussed in the preceding paragraph.

A priori, it is unclear whether mediators' improvements in surplus would accrue mostly to buyers or sellers: a high-skilled mediator could achieve higher total surplus by affecting one or both of the agents' surplus. Figure 9.B shows that higher-skilled mediators tend to improve seller surplus and that we cannot reject a null effect on buyer surplus. This result is driven by two facts: (i) regardless of mediator assignment, more surplus accrues to sellers than to buyers and (ii) mediators have little effect on prices.[72] Therefore, an increase in trade probability from a higher-skilled mediator affects sellers' surplus more than buyers'.

Table 4: Auction House Revenue Under Different Mediator Types (Units = $)

	Low	Medium	High
Revenue	135.22	168.34	202.46
	(1.59)	(1.64)	(1.43)
	High - Medium	Medium - Low	High - Low
Revenue Difference	34.12	33.12	67.24
	(2.24)	(2.22)	(2.25)

Notes: Table displays average auction house revenue per sales attempt under different mediator types, as well as the differences between average revenues. Standard errors are from 200 bootstrap replications. Units = $.

A natural question is how large the revenue gains are for the auction house company from having a better mediator. The company collects fees from buyers and sellers only when trade occurs (see Appendix C.8). These fees are observable in the auction-plus-bargaining data, and thus we can estimate the fees collected under different mediator skill levels. Table 4 displays average auction house revenue per sales attempt under high-, medium-, and low-skilled mediators, where this

[72]Larsen and Zhang (2026) showed that, on average, sellers receive a larger surplus than buyers in this market due to competition among bidders.

revenue, for negotiations handled by a given mediator type, is simply the trade probability multiplied by the average auction house fees collected from negotiations that ended in trade. Average revenue per sales attempt is $203 for high-skilled mediators, $168 for medium-skilled mediators, and $135 for low-skilled mediators. Thus, high-skilled mediators improve auction house revenue by about 50% relative to low-skilled mediators. These numbers are also meaningful relative to the total surplus differences between high- and low-skilled mediators discussed above.

6 Conclusion

In this paper, we have shown that mediators have statistically significant and economically large effects on bargaining outcomes. The 75th-percentile intermediary is 18.07 percentage points more likely to close a deal than the 25th percentile intermediary. The estimated mediator effects are robust to a variety of different controls, and are not driven by sample noise. Effective mediators tend to achieve trade quickly, and particularly stand out above low-performing mediators in situations where agreement does not look like a foregone conclusion (i.e., when the auction price is not close to the reserve price). Mediators appear to become more effective with experience, but there is substantial variation in mediator effectiveness left unexplained by experience. Our structural exercise demonstrated that better mediators not only increase trade probabilities but also have real effects on realized gains from trade. Myerson and Satterthwaite (1983) showed that incomplete-information bilateral bargaining under overlapping supports generically results in deadweight losses — cases where a buyer values the item more than the seller but they nonetheless fail to trade. These represent potential gains from trade left on the table. Our findings suggest that real-world bargaining falls even farther from the first-best outcome than implied by Myerson and Satterthwaite (1983), but that some of this shortfall can be overcome by high-performing mediators who are able to execute a mechanism/equilibrium lying closer to the efficient outcome. We see these findings as potential starting points for future experimental, theoretical, and empirical work in incomplete-information bargaining settings. A natural next step would be to design laboratory experiments to analyze more in-depth the question of *how* good mediators achieve a higher trade probability or to identify interventions that could shift play to better equilibria. We hope these findings can also motivate theoretical investigations explicitly allowing for mediator heterogeneity.

References

Abreu, D. and Gul, F. (2000). Bargaining and reputation. *ECMA*, 68(1):85–117.

Ambrus, A., Chaney, E., and Salitskiy, I. (2018). Pirates of the Mediterranean: An empirical investigation of bargaining with asymmetric information. *QE*, 9(1):217–246.

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.

Barwick, P. J. and Pathak, P. A. (2015). The costs of free entry: An empirical study of real estate agents in greater Boston. *RAND*, 46(1):103–145.

Basak, D. (2024). Social value of public information in bargaining. *EJ*, 134(660):1356–1378.

Bazerman, M. H., Neale, M. A., Valley, K. L., Zajac, E. J., and Kim, Y. M. (1992). The effect of agents and mediators on negotiation outcomes. *Org. Beh. & Hum. Dec. Proc.*, 53(1):55–73.

Bercovitch, J. and Jackson, R. (2001). Negotiation or mediation? An exploration of factors affecting the choice of conflict management in international conflict. *Neg. J.*, 17(1):59–77.

Biglaiser, G., Li, F., Murry, C., and Zhou, Y. (2020). Intermediaries and product quality in used car markets. *RAND*, 51(3):905–933.

Bruno, B., Garcia-Appendini, E., and Nocera, G. (2018). Experience and brokerage in asset markets: Evidence from art auctions. *Fin. Man.*, 47(4):833–864.

Casella, A., Friedman, E., and Archila, M. P. (2020). Mediating conflict in the lab. NBER 28137.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2024). On binscatter. *AER*, 114(5):1488–1514.

Copic, J. and Ponsati, C. (2008). Robust bilateral trade and mediated bargaining. *JEEA*, 6(2-3):570–580.

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

Dindaroglu, B. and Ertac, S. (2024). An empirical study of sequential offer bargaining during the Festival of Sacrifice. *J. Econ. Psych.*, 101:102707.

Dixon, W. J. (1996). Third-party techniques for preventing conflict escalation and promoting peaceful settlement. *Int. Org.*, 50(4):653–681.

Edlin, A. and Shannon, C. (1998). Strict monotonicity in comparative statics. *JET*, 81(1):201–219.

Efron, B. and Hastie, T. (2021). *Computer Age Statistical Inference, Student Edition: Algorithms, Evidence, and Data Science*, volume 6. Cambridge University Press.

Eilat, R. and Pauzner, A. (2021). Bilateral trade with a benevolent intermediary. *The. Econ.*, 16(4):1655–1714.

Eisenkopf, G. and Bachtiger, A. (2013). Mediation and conflict prevention. *J. Con. Res.*, 57(4):570–597.

Emery, R. E., Matthews, S. G., and Wyer, M. M. (1991). Child custody mediation and litigation: Further evidence on the differing views of mothers and fathers. *J. Con. and Clin. Psych.*, 59(3):410.

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

Fanning, J. (2021). Mediation in reputational bargaining. *AER*, 111(8):2444–72.

Fanning, J. (2023). Optimal dynamic mediation. *JPE*, 131(11):2969–3002.

Fey, M. and Ramsay, K. W. (2010). When is shuttle diplomacy worth the commute? Information sharing through mediation. *World Politics*, 62(4):529–560.

Freyberger, J. and Larsen, B. (2022). Identification in ascending auctions, with an application to digital rights management. *QE*, 13(2):505–543.

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.

Genesove, D. (1993). Adverse selection in the wholesale used car market. *JPE*, 101(4):644–665.

Gilbukh, S. and Goldsmith-Pinkham, P. (2024). Heterogeneous real estate agents and the housing cycle. *RFS*, 37(11):3431–3489.

Glode, V. and Opp, C. (2016). Asymmetric information and intermediation chains. *AER*, 106(9):2699–2721.

Goltsman, M., Hörner, J., Pavlov, G., and Squintani, F. (2009). Mediation, arbitration and negotiation. *JET*, 144(4):1397–1420.

Gottardi, P. and Mezzetti, C. (2024). Shuttle diplomacy. *JET*, 216:105794.

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. WP, Yale.

Haile, P. and Tamer, E. (2003). Inference with an incomplete model of English auctions. *JPE*, 111(1):1–51.

Hörner, J., Morelli, M., and Squintani, F. (2015). Mediation and peace. *RESTUD*, 82(4):1483–1501.

Jarque, X., Ponsatı, C., and Sákovics, J. (2003). Mediation: Incomplete information bargaining with filtered communication. *J. Math. Econ.*, 39(7):803–830.

Jindal, P. and Newberry, P. (2022). The profitability of revenue-based quotas under price negotiation. *Man. Sci.*, 68(2):917–940.

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

Kim, J. Y. (2017). Interim third-party selection in bargaining. *GEB*, 102:645–665.

Kim, K. and Lee, J. (2014). Nonparametric estimation and testing of the symmetric IPV framework with unknown number of bidders. WP, Michigan State.

Kim, S. (2024). Informed agent's advice in bargaining under two-sided incomplete information. *Journal of Economics*, 143(2):183–205.

Kong, Y., Silveira, B. S., and Tang, X. (2025). Risk and information in dispute resolution: An empirical study of arbitration. *JPE*, 133(9).

Kotlarski, I. (1967). On characterizing the gamma and normal distribution. *PJM*, 20(1):69–76.

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

Kydd, A. (2003). Which side are you on? Bias, credibility, and mediation. *AJPS*, 47(4):597–611.

Lacetera, N., Larsen, B., 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.

Larsen, B. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *RESTUD*, 88(2):851–882.

Larsen, B. and Zhang, A. L. (2018). A mechanism design approach to identification and estimation. NBER 24837.

Larsen, B. and Zhang, A. L. (2026). Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry. *RESTUD*, forthcoming.

Li, H. and Liu, N. (2015). Nonparametric identification and estimation of double auctions with bargaining. WP, Penn State.

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

Li, Y. and Zhang, X. (2024). Collective decision through an informed mediator. *JET*, 218:105817.

McEwen, C. A. and Maiman, R. J. (1981). Small claims mediation in Maine: An empirical assessment. *Me. L. Rev.*, 33:237.

Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *JET*, 29(2):265–281.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *ECMA*, 47(1):61–73.

Perry, M. (1986). An example of price formation in bilateral situations: A bargaining model with incomplete information. *ECMA*, 54(2):313–321.

Relekar, H. and Ren, Y. (2025). The extra shot: How post-auction bagaining rewrites auction outcomes. WP, University of Chicago.

Roberts, J. W. (2013). Unobserved heterogeneity and reserve prices in auctions. *RAND*, 44(4):712–732.

Robles-Garcia, C. (2020). Competition and incentives in mortgage markets: The role of brokers. WP, Stanford.

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

Rubinstein, A. (1985). A bargaining model with incomplete information about time preferences. *ECMA*, 53(5):1151–1172.

Saran, R. (2011). Bilateral trading with naive traders. *GEB*, 72(2):544–557.

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

Silveira, B. S. (2017). Bargaining with asymmetric information: An empirical study of plea negotiations. *ECMA*, 85(2):419–452.

Song, U. (2004). Nonparametric estimation of an ebay auction model with an unknown number of bidders. WP, University of British Columbia.

Stipanowich, T. J. and Lamare, J. R. (2014). Living with ADR: Evolving perceptions and use of mediation, arbitration, and conflict management in Fortune 1000 corporations. *HNLR*, 19:1.

Storms, E. (2015). Monotonicity of trade in PSBNE of bargaining with alternating offers. Mimeo, Stanford.

Treece, J. B. (2013). Manheim exec: Why do so few sales at auctions close? *Automotive News*. Url: http://www.autonews.com/article/20131202/RETAIL04/312029992/manheim-exec-why-do-so-few-sales-at-auctions-close. Posted Dec. 2, 2013. Retrieved Dec. 2, 2013.

Williams, S. (1987). Efficient performance in two-agent bargaining. *JET*, 41(1):154–172.

Williams, S. (1999). A characterization of efficient, Bayesian incentive compatible mechanisms. *Economic Theory*, pages 155–180.

A Proofs

Before proving the main results we introduce notation and prove preliminary lemmas.

A.1 Preliminaries. For a seller of type s and buyer of type b, $\pi^S(p^A, s)$ and $\pi^B(p^A, b)$ are the *equilibrium expected bargaining payoffs* for the seller and buyer, respectively, conditional on the events that (i) the auction price is equal to p^A, (ii) p^A lies below the reserve price, and (iii) the buyer chooses to enter bargaining. We will refer to an agent's *bargaining strategy* as her strategy from the beginning of the bargaining stage of the game and onward conditional on these events. $\pi^S(p^A, s)$ and $\pi^B(p^A, b)$ constitute expected bargaining payoffs before the mediator is assigned; thus, these objects incorporate expectations over κ. Let $\pi^S(p^A, s, \kappa)$ and $\pi^B(p^A, b, \kappa)$ (with three arguments) be the analogs of these objects conditional on given realization of κ.

Agents' expected bargaining payoffs from playing a given strategy can be summarized by three quantities: the trade probability, the expected price when trade occurs, and the expected period in which the game ends under that strategy. When the mediator is κ and a seller plays the equilibrium strategy $\sigma^{S*}(s, \kappa)$ corresponding to a type s, and all buyers play their equilibrium strategies σ^{B*}, let $\Pr\left(\mathcal{A}|\sigma^{S*}(s, \kappa)\right)$ denote the trade probability, $price\left(\sigma^{S*}(s, \kappa)\right)$ denote the price, and $\mathcal{T}\left(\sigma^{S*}(s, \kappa)\right)$ denote the expected period in which the game ends. Note that these expectations average over buyer types, which are unknown to the seller. Similarly, when a buyer of type b plays the equilibrium strategy $\sigma^{B*}(b, \kappa)$ corresponding to a type b and all sellers play their equilibrium strategies σ^{S*}, let $\Pr\left(\mathcal{A}|\sigma^{B*}(b, \kappa)\right)$ denote the trade probability, $price\left(\sigma^{B*}(b, \kappa)\right)$ denote the price, and $\mathcal{T}\left(\sigma^{B*}(b, \kappa)\right)$ denote the expected period in which the game ends.

Conditional on bargaining occurring (that is, a period $t > 1$ is reached), the seller's total expected cost due to bargaining when playing the equilibrium strategy of a type s is given by $\eta_S \lfloor (\mathcal{T}(\sigma^{S*}(s, \kappa)) - 1)/2 \rfloor$, where $\lfloor \cdot \rfloor$ represents the floor function; this is because, by any period t, the seller has made a total of $\lfloor (t-1)/2 \rfloor$ offers. The buyer, on the other hand, faces a cost of η_B upon initially entering bargaining as well as any expected bargaining costs from later periods, $\eta_B \lfloor (\mathcal{T}(\sigma^{B*}(b, \kappa)) - 2)/2 \rfloor$. For brevity, the following combines the expected price and costs:

$$\mathcal{P}_S(\sigma^{S*}(s, \kappa)) \equiv price(\sigma^{S*}(s, \kappa)) \Pr\left(\mathcal{A}|\sigma^{S*}(s, \kappa)\right) - \eta_S \lfloor (\mathcal{T}(\sigma^{S*}(s, \kappa)) - 1)/2 \rfloor \tag{10}$$

$$\mathcal{P}_B(\sigma^{B*}(b, \kappa)) \equiv price(\sigma^{B*}(b, \kappa)) \Pr(\mathcal{A}|\sigma^{B*}(b, \kappa)) + \eta_B \lfloor (\mathcal{T}(\sigma^{B*}(b, \kappa)) - 2)/2 \rfloor \tag{11}$$

Thus, $\mathcal{P}_S(\sigma^{S*}(s,\kappa))$ is the expected monetary amount accruing to the seller when she plays $\sigma^{S*}(s,\kappa)$ and $\mathcal{P}_B(\sigma^{B*}(b,\kappa))$ is the expected monetary amount incurred by the buyer when he plays $\sigma^{B*}(b,\kappa)$.

Using (10) and (11), when the mediator's type is κ, the expected bargaining payoff for a seller of type s playing her equilibrium strategy $\sigma^{S*}(s,\kappa)$ can be written

$$\pi^S(p^A,s,\kappa) = s\left(1-\Pr(\mathcal{A}|\sigma^{S*}(s,\kappa))\right) + \mathcal{P}_S(\sigma^{S*}(s,\kappa)) \tag{12}$$

Similarly, the expected bargaining payoff for a buyer of type b playing his equilibrium strategy $\sigma^{B*}(b,\kappa)$ can be written

$$\pi^B(p^A,b,\kappa) = b\Pr(\mathcal{A}|\sigma^{B*}(b,\kappa)) - \mathcal{P}_B(\sigma^{B*}(b,\kappa)) \tag{13}$$

The objects $\pi^S(p^A,s)$ and $\pi^B(p^A,b)$ are constructed by averaging (12) and (13) over κ. Importantly, whenever we use $\pi^S(p^A,s)$ and $\pi^B(p^A,b)$, it is in the context of decisions made before κ is realized, including the buyer's decision to enter bargaining and the seller's reserve price choice. Note that we write the buyer's expected bargaining payoff as not including the initial η_B paid upon entering bargaining. Thus, prior to incurring this cost, a buyer's expected payoff from *entering bargaining* is given by $\pi^B(p^A,b)-\eta_B$.

Lemma 4 and Proposition 4 below use the notation $\Pr(\mathcal{A}|\sigma^{S*}(s,\kappa),h_t)$, which denote the probability that trade eventually occurs conditional on the history of play being h_t at period t and conditional the seller playing the strategy corresponding to a type s from that point onward (and buyers continue to play their equilibrium strategies). Similarly, the object $\mathcal{P}_S(\sigma^{S*}(s,\kappa),h_t)$ represents the expected monetary amount received by the seller with the same conditioning. The analogous notation applies for the buyer.

Lemma 1. *Under A1–A3, for any p^A, $\pi^S(p^A,s)$ weakly increases in s and $\pi^B(p^A,b)$ weakly increases in b.*

Proof. We prove the case for the seller; the buyer is analogous. Consider a fixed realization of κ, the mediator type. We show that, for s and s' such that $s>s'$, s can attain a weakly higher payoff than s' simply by mimicking the strategy of s'; thus, s does weakly better than s' in equilibrium.

If, instead of playing her own equilibrium strategy, a type s plays $\sigma^{S*}(s',\kappa)$, mimicking the

equilibrium bargaining strategy of s', where $s' < s$, she receives

$$s\left(1-\Pr\left(\mathcal{A}|\sigma^{S*}\left(s',\kappa\right)\right)\right)+\mathcal{P}\left(\sigma^{S*}\left(s',\kappa\right)\right)\geq s'\left(1-\Pr(\mathcal{A}|\left(\sigma^{S*}\left(s',\kappa\right)\right)\right)+\mathcal{P}\left(\sigma^{S*}\left(s',\kappa\right)\right) \quad (14)$$

The right-hand side of (14) is the equilibrium expected bargaining payoff of s', $\pi^S(p^A, s', \kappa)$, given by (12) with s' replacing s.[73] Now, s must achieve a weakly higher expected payoff from her equilibrium strategy $\sigma^{S*}(s,\kappa)$ than from that of any other type, implying that

$$s\left(1-\Pr\left(\mathcal{A}|\sigma^{S*}\left(s,\kappa\right)\right)\right)+\mathcal{P}\left(\sigma^{S*}\left(s,\kappa\right)\right)\geq s\left(1-\Pr\left(\mathcal{A}|\sigma^{S*}\left(s',\kappa\right)\right)\right)+\mathcal{P}\left(\sigma^{S*}\left(s',\kappa\right)\right) \quad (15)$$

The left-hand side of (15) is the equilibrium expected bargaining payoff of s, $\pi^S(p^A, s, \kappa)$, given by (12). Together, (14) and (15) imply that $\pi^S(p^A, s, \kappa)$ is weakly increasing in s. Averaging over κ yields the desired result. □

Proof of Proposition 1: Consider an arbitrary bidder of type $B = b$. A bidder's strategy is the price at which to stop bidding, ζ. Suppose $\zeta > b$, so the buyer drops out above her value. Consider an auction price $P^A = p^A$. If $p^A \in (b, \zeta)$, the buyer's payoff is nonpositive with probability 1, because the final bargained price is required to be greater than p^A (as stated in Section 5.2). If $p^A > b$ and exceeds the reserve price, the buyer is required to purchase the car at p^A, which delivers a strictly negative payoff. If p^A is below the reserve price, the buyer's payoff conditional on entering bargaining will also be nonpositive, again because the final bargained price is required to be greater than p^A. As a result, dropping out at $\zeta > b$ is weakly dominated by dropping out at $\zeta = b$.

Now suppose $\zeta < b$, so the buyer drops out below her value. The buyer can then guarantee a weakly greater payoff by instead dropping out at her value b, and never entering bargaining if $p^A \in (\zeta, b)$. When $p^A \in (\zeta, b)$, this strategy delivers a payoff of zero when p^A is below the reserve price, because the buyer immediately exits bargaining; it delivers the strictly positive payoff $b - p^A$ if p^A exceeds the reserve price, and thus the game ends at the auction stage. Thus, dropping out at $\zeta < b$ is also weakly dominated by dropping out at $\zeta = b$. □

Lemma 2. *Under A1–A5, in any equilibrium, for any realization of $P^A = p^A$, $\chi(p^A)$ exists and satisfies $\chi(p^A) > p^A$. Moreover, if $\chi(p^A) < \overline{b}$, then $\pi^B(p^A, b) - \eta_B > 0 \ \ \forall b > \chi(p^A)$.*

[73] When a type s mimics s', the buyer's beliefs evolve as they would when the buyer faces a type s'.

Proof. We first show that, for any p^A, it is not possible for $\pi^B(p^A, b) - \eta_B$ to be positive for all $b \in (\underline{b}, \bar{b})$. Note that $\pi^B(p^A, b)$ is bounded, because the best price a buyer can achieve is p^A and the probability of trade is bounded by 1. Thus, by Lemma 1, any buyer with value $b < p^A + \eta_B$ is guaranteed negative surplus from entering bargaining. This implies both that $\chi(p^A)$ as defined in (6) exists (because it is either the case that $\pi^B(p^A, b) - \eta_B < 0$ for all $b \in (\underline{b}, \bar{b})$ or that $\pi^B(p^A, b) - \eta_B = 0$ for some b) and that the lowest-valued buyer who enters bargaining satisfies $\chi(p^A) \geq p^A + \eta_B > p^A$. The second part of Lemma 2 follows by Lemma 1. □

Lemma 3. *(i) Under A1–A5, the seller's optimal secret reserve price is weakly increasing in s. (ii) Under A6–A7, the seller's optimal secret reserve price is strictly increasing.*

Proof. A seller of type s chooses ρ to solve $\rho^*(s) = \max_\rho V(s, \rho)$, where $V(s, \rho)$ is given by

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

We need not specify any dependence of the expected bargaining payoffs on κ because we are evaluating a decision taken before κ is realized, and hence the forms of expected bargaining payoff we use is $\pi^S(p^A, s)$ from Appendix A.1, which integrate over possible realizations of κ.

We show that $V(s, \rho)$ has weakly increasing differences in s and ρ. Consider $\rho > \rho'$ and the difference $V(s, \rho) - V(s, \rho')$, which is

$$-\int_{\rho'}^{\rho} p^A f_{P^A}\left(p^A\right) dp^A + \int_{\rho'}^{\rho} \left[\underbrace{s\left(F_B\left(\chi\left(p^A\right)\right) - F_B\left(p^A\right)\right)}_{(i)} + \underbrace{\pi^S\left(p^A, s\right)\left(1 - F_B\left(\chi\left(p^A\right)\right)\right)}_{(ii)}\right] \frac{f_{P^A}\left(p^A\right)}{1 - F_B\left(p^A\right)} dp^A \tag{16}$$

To show that (16) is weakly increasing, we show that both (16.i) and (16.ii) are. For (16.i), Lemma 2 shows that $\chi(p^A) > p^A$, and thus, because F_B is a CDF (and hence a weakly increasing function) $F_B\left(\chi\left(p^A\right)\right) - F_B\left(p^A\right)$ is non-negative. For (16.ii), Lemma 1 shows that $\pi^S\left(p^A, s\right)$ is weakly increasing in s, and hence (16.ii) is also weakly increasing in s. These results imply, by Topkis' Theorem, that $\rho^*(s)$ is weakly increasing in s, proving Lemma 3.i.

With the additional assumptions *A6–A7*, we can apply the result of Edlin and Shannon (1998)

to show further that $\rho^*(s)$ is strictly increasing in s. Differentiating $V(s,\rho)$ yields

$$\frac{\partial V}{\partial \rho} = f_{P^A}(\rho)\left[-\rho + \underbrace{s\frac{F_B(\chi(\rho)) - F_B(\rho)}{1-F_B(\rho)}}_{(i)} + \underbrace{\pi^S(\rho,s)\frac{1-F_B(\chi(\rho))}{1-F_B(\rho)}}_{(ii)}\right] \tag{17}$$

First, note that f_{P^A} is everywhere positive because f_B is (by $A7$) and hence, by $A5$, f_{P^A} is as well. Next, by Lemma 2, $\chi\left(p^A\right) > p^A$. This fact, combined with f_B being everywhere positive ($A7$), implies that $\frac{1-F_B(\chi(\rho))}{1-F_B(\rho)}$ and $\frac{F_B(\chi(\rho))-F_B(\rho)}{1-F_B(\rho)}$ are everywhere positive. Thus, the term (17.i) is positive and strictly increasing in s. By Lemma 1, $\pi^S(\rho,s)$ is weakly increasing in s, and thus (17.i) is weakly increasing in s. Thus, $\frac{\partial V}{\partial \rho}$ is strictly increasing in s, and $V(s,\rho)$ satisfies the increasing marginal returns property of Edlin and Shannon (1998). This implies that the optimal reserve price $\rho^*(s)$ is strictly increasing in s as long as in lies in $(\underline{b},\overline{b})$, which, by $A7$, is the whole real line, completing the proof of Lemma 3.ii.

Note that $A6$ — $\pi^S(\rho,s)$ being continuous in ρ — is needed to satisfy the differentiability assumption of Edlin and Shannon (1998). Intuitively, if $\pi^S(\rho,s)$ discontinuously jumps downwards in ρ at some optimum $\rho^*(s)$, $V(s,\rho)$ is not differentiable at this $\rho^*(s)$, and $\frac{\partial V}{\partial \rho^+}|_{\rho=\rho^*(s)}$, the right-derivative at the optimum, can be negative rather than zero. (17) shows that an increase from s to s' increases this right-derivative; however, if $V(s,\rho)$ is discontinuous, $\frac{\partial V}{\partial \rho^+}$ can still be negative after this increase, implying that an increase in ρ above $\rho^*(s)$ may not increase V. □

Proof of Proposition 3: Lemma 3.i and Lemma 2 imply (i). Lemma 3.ii implies (ii). □

Lemma 4. *If $A1$–$A3$ are satisfied, then, for any κ, the following hold: (i) For any types b and b', with $b' > b$, and any history h_t reached on path by both b and b',*

$$\Pr\left(\mathcal{A} \mid \sigma^{B*}(b',\kappa), h_t\right) \geq \Pr\left(\mathcal{A} \mid \sigma^{B*}(b,\kappa), h_t\right) \tag{18}$$

(ii) For any types s and s', with $s' < s$, and any history h_t reached on path by both s and s',

$$\Pr\left(\mathcal{A} \mid \sigma^{S*}(s',\kappa), h_t\right) \geq \Pr\left(\mathcal{A} \mid \sigma^{S*}(s,\kappa), h_t\right) \tag{19}$$

Proof. We prove this for the buyer; the case for the seller is symmetric. Buyer type b can attain the

trade probability $\Pr\left(\mathcal{A} \mid \sigma^{B*}(b', \kappa), h_t\right)$ and pay the expected amount $\mathcal{P}_B(\sigma^{B*}(b', \kappa), h_t)$ — attained by b' in equilibrium — simply by mimicking the behavior of b' from h_t onward. But, in equilibrium, b must prefer playing her strategy from h_t onward to playing that of b', implying that

$$b \Pr\left(\mathcal{A} \mid \sigma^{B*}(b, \kappa), h_t\right) - \mathcal{P}_B(\sigma^{B*}(b, \kappa), h_t) \geq b \Pr\left(\mathcal{A} \mid \sigma^{B*}(b', \kappa), h_t\right) - \mathcal{P}_B(\sigma^{B*}(b', \kappa), h_t)$$

Likewise, b' must prefer her strategy to that of b, implying that

$$b' \Pr\left(\mathcal{A} \mid \sigma^{B*}(b, \kappa), h_t\right) - \mathcal{P}_B(\sigma^{B*}(b, \kappa), h_t) \leq b' \Pr\left(\mathcal{A} \mid \sigma^{B*}(b', \kappa), h_t\right) - \mathcal{P}_B(\sigma^{B*}(b', \kappa), h_t)$$

Subtracting these inequalities and dividing by the strictly positive $b - b'$ yields (18). □

Proof of Proposition 4: This proof follows that of Theorem 1 of Storms (2015) with little variation. Fix κ and p^A throughout this proof. Note that, in any pure-strategy equilibrium, the history of play h_t up to any period t is a deterministic function of the buyer and seller types, and any pair of buyer and seller types eventually end up trading or not trading.

The proof proceeds by contradiction. Such a contradiction would be a triple (s, b, b') with $b' > b$ such that the deterministic path that types s and b follow in equilibrium eventually results in trade and that of s and b' does not. For the sake of clarity, we will give such triples a name, referring to them as *Type A* triples. For any such triple, the equilibrium paths followed by (s, b) and (s, b') diverge at some history h_t in some period t. The game length is bounded, so, for any Type-A triple, this t is finite. The set of types is also finite, so there are a finite number of Type-A triples. Thus, there is some triple that takes the longest to diverge (or multiple triples tied for longest-to-diverge). Among the set of triples tied for longest-to-diverge, we choose an arbitrary triple. Throughout the remainder of the proof, let (s, b, b') denote this triple. Call the time period after which histories diverge for this longest Type-A triple $\breve{t}$, and call the history $\breve{h}_{\breve{t}}$.

Equation (18) of Lemma 4 implies that b' must trade with weakly greater probability than b, conditional on reaching $\breve{h}_{\breve{t}}$. Because b' does not ultimately trade with s, but b does, there must be at least one other seller type $s' \neq s$ that reaches $\breve{h}_{\breve{t}}$ on the equilibrium path with both b and b' such that s' ultimately trades with b' and s' ultimately does not trade with b. In other words, in order for (18) to be true and for it to simultaneously be true that s trades with b but not with b', there must be some additional type s' that eventually trades with b' in equilibrium.

Now note that $\breve{h}_{\breve{t}}$ cannot be a terminal history, because the outcomes of b and b' diverge thereafter, so there must be at least one additional period $\breve{t}+1$ on the equilibrium paths for both b and b', at which point the histories diverge. Importantly, the divergence in histories from $\breve{t}$ to $\breve{t}+1$ must be driven by differences in actions taken by buyers (b and b') rather than actions by the seller. This is because, up to $\breve{h}_{\breve{t}}$, the actions taken by b and b' are identical, so the seller thus cannot distinguish b from b' based on $\breve{h}_{\breve{t}}$, and cannot tailor her action based on that distinction. Because s and s' both reach $\breve{h}_{\breve{t}}$ on path with both b and b', this *buyer-driven history divergence* implies that, on path, both s and s' reach the history that b diverges to (denote this history $h^b_{\breve{t}+1}$) and both also reach the history that b' diverges to (denote this $h^{b'}_{\breve{t}+1}$).

We now show that both possible cases, $s' > s$ and $s' < s$, lead to contradictions.

1. $s' > s$. Consider $h^{b'}_{\breve{t}+1}$, the history reached by (s, b') and (s', b') on the equilibrium path at $\breve{t}+1$. s cannot trade with any types $\tilde{b} < b'$ that reach $h^{b'}_{\breve{t}+1}$ on the equilibrium path because otherwise the triple $(s, \tilde{b}, b')$ would be a Type-A triple that diverges after $\breve{t}+1$, contradicting our construction of (s, b, b') as the longest Type-A triple. Now, because $s' > s$, by (19),

$$\Pr\left(\mathcal{A} \mid \sigma^{S*}(s,\kappa), h^{b'}_{\breve{t}+1}\right) \geq \Pr\left(\mathcal{A} \mid \sigma^{S*}(s',\kappa), h^{b'}_{\breve{t}+1}\right) \tag{20}$$

   That is, the probability of trade for s must be higher than the probability of trade for s', conditional on reaching $h^{b'}_{\breve{t}+1}$. We stated, in constructing s', that s' ultimately trades with b'. If (20) holds, and s neither trades with b' nor any $\tilde{b} < b'$, then there must be some type $b''' > b'$ that reaches $h^{b'}_{\breve{t}+1}$ on path such that s ultimately trades with b''' and s' does not. We then have $b''' > b'$; both b''' and b' reach $h^{b'}_{\breve{t}+1}$ with s'; s' ultimately trades with b'; and s' ultimately does not trade with b'''. Thus, (s', b', b''') constitutes a Type-A triple of length at least $\breve{t}+1$, contradicting our choice of (s, b, b') as the longest type-A triple.

2. $s' < s$. Consider $h^{b}_{\breve{t}+1}$, the history reached by (s, b) and (s', b) on path at $\breve{t}+1$. Because s ultimately trades with b, s must ultimately trade with any $\tilde{b} > b$ that reaches $h^{b}_{\breve{t}+1}$ with s. Otherwise, the triple $(s, b, \tilde{b})$ would constitute a Type-A triple that diverges after $\breve{t}+1$, contradicting our construction of (s, b, b') as the longest Type-A triple. Now, because $s' < s$, the probability of trade under s' must be greater than the probability of trade under s,

conditional on reaching $h^b_{\breve{t}+1}$:

$$\Pr\left(\mathcal{A} \mid \sigma^{S*}(s',\kappa), h^b_{\breve{t}+1}\right) \geq \Pr\left(\mathcal{A} \mid \sigma^{S*}(s,\kappa), h^b_{\breve{t}+1}\right) \tag{21}$$

We stated, in constructing s', that s ultimately trades with b and s' ultimately does not trade with b. If (21) holds, and s trades with both b and all $\tilde{b} > b$, then there must be some type $b''' < b$, such that b''' reaches $h^b_{\breve{t}+1}$, and b''' ultimately trades with s', but not with s. We then have $b''' < b$; both b''' and b reach $h^b_{\breve{t}+1}$ with s; s ultimately trades with b'''; and s ultimately does not trade with b. Thus, (s, b''', b) constitutes a Type-A triple of length at least $\breve{t}+1$, contradicting our choice of (s, b, b') as the longest type-A triple. □

Lemma 5. *For a fixed κ, conditional on any on- or off-path history of play, including the reserve price and all public actions, the resulting buyer payoffs and the resulting seller payoff (net of her value) are invariant to the realization of game-level heterogeneity.*

Proof. Consider a fixed seller of type s and a fixed highest-value buyer of type b. We will consider two settings: one in which $Y'\gamma + W = z$ (which we will refer to as the z game) and one in which $Y'\gamma + W = z' \neq z$ (the z' game), with z and z' fixed throughout. We consider first the z game. We will use here the terms *markup action*, *total action*, and *total value* as defined in the body of the paper after Proposition 5.

Recall that h_t is a realization of the public history; we will use $fh_t \equiv \{r, h_t\}$ to denote a realization of what we will call the *full history*, which includes the reserve price in addition to the public history. Note that we are explicit about the inclusion of the reserve price because it is not publicly observed, unlike other actions in the game.

Let a *full terminal history* refer to a history of markup actions (on or off the equilibrium path), including the reserve price and all public actions, resulting in the game's termination. Any full terminal history either ends in trade or no trade. Trade occurs through the auction price exceeding the reserve price or through an agent accepting a bargaining offer. A no-trade outcome occurs either through an agent quitting in bargaining or through the buyer opting out of bargaining before it begins. Any full terminal history fh_t has a corresponding *final period* t in which the game ends and a *markup final price* p at which the agents trade if the full terminal history involves trade.

Consider the z game and an arbitrary full terminal history fh_t with a markup final price p and final period t. We can write the buyer's total payoff from the markup actions taken in this full terminal history as

$$(b+z)-(p+z)-\eta_B\lfloor t/2\rfloor \quad \text{or} \quad -\eta_B\lfloor t/2\rfloor \tag{22}$$

in the trade or no trade cases, respectively. The expressions in (22) show that z falls out of the buyer's payoffs, yielding

$$b-p-\eta_B\lfloor t/2\rfloor \quad \text{or} \quad -\eta_B\lfloor t/2\rfloor \tag{23}$$

Consequently, if agents follow the same full terminal history fh_t (i.e., the same markup actions) in the z' game, the buyer's payoffs will still be given by (23). For example, consider the z game and a full terminal history involving a markup auction price of p^A (meaning a total auction price that is p^A+z), the buyer entering bargaining, the seller countering a markup offer of p_2^S (meaning a total price of p_2^S+z), and the buyer accepting. Then, in the z' game, if agents follow the same markup actions, the total auction price would be p^A+z' and the total second-period offer would be p_2^S+z', but the buyer's payoff would be unchanged because z and z' fall out of payoffs.[74]

We now walk through the same argument for the seller. Again consider the z game and an arbitrary full terminal history fh_t with a markup final price p and final period t. The total payoffs for the seller are

$$(p+z)-\eta_S\lfloor (t-1)/2\rfloor \quad \text{or} \quad (s+z)-\eta_S\lfloor (t-1)/2\rfloor$$

We can subtract $s+z$ from the seller's payoff at full terminal history fh_t and *all other* full terminal histories of the game without affecting the seller's preferences over full terminal histories, yielding

[74]For additional insight into this payoff invariance result, note that it would not hold if game-level heterogeneity were to enter values *multiplicatively* rather than additively. For example, suppose a buyer's and seller's total values when the game-level heterogeneity realization is z are given by bz and sz, respectively, and suppose, at any markup final price of p, the total price is also multiplicative, given by pz. In this alternative model, a buyer's total payoff from a full terminal history ending in trade in final period t at markup final price p would be $bz-pz-\eta_B\lfloor t/2\rfloor$. Because of bargaining costs, this multiplicative form for z does not fall out of the expression. This opens up the possibility that the buyer prefers one sequence of actions in the z game and a different one in the z' game.

payoffs

$$(p+z)-(s+z)-\eta_S\lfloor(t-1)/2\rfloor \quad \text{or} \quad -\eta_S\lfloor(t-1)/2\rfloor \tag{24}$$

As in the buyer's payoff, z falls out, yielding

$$p-s-\eta_S\lfloor(t-1)/2\rfloor \quad \text{or} \quad -\eta_S\lfloor(t-1)/2\rfloor \tag{25}$$

□

Proof of Proposition 5: This proof uses terms defined in Lemma 5. Consider a fixed κ, a fixed seller of type s, and a fixed highest-value buyer of type b. Recall that $\sigma^B(b,\kappa)$ and $\sigma^S(s,\kappa)$ describe the strategy of a buyer of type b and seller of type s at every history. We introduce new notation to denote a subset of these strategies (an agent's period t action) at a particular history. For odd t, let $\sigma_t^B(b,\kappa,h_t)$ denote the period-t action of a buyer of type b following strategy $\sigma^B(b,\kappa)$ when the (non-terminal) public history up to time period t is h_t. Similarly, for even t, we use $\sigma_t^S(s,\kappa,fh_t)$ for the seller's period-t action. As in Section 5.2, $\mu_B(S|h_t,b,\kappa)$ and $\mu_S(B|fh_t,s,\kappa)$ represent the buyer's and seller's beliefs about their opponent's value. Note that the seller's action and belief condition on the full history fh_t and the buyer's condition on the public history h_t because only the seller knows her choice of r. That is, the information sets of the buyer do not include r but those of the seller do. As in the proof of Lemma 5, these histories denote sets of markup actions.

Let $(\sigma^{B*},\sigma^{S*})$ and (μ_{B*},μ_{S*}) be a particular set of strategies and beliefs that are an equilibrium of the z game. Proposition 5 implies that these same strategies and beliefs constitute a PBE in the z' game. To see that this is true, note that the properties of PBE are satisfied by $(\sigma^{B*},\sigma^{S*})$ and (μ_{B*},μ_{S*}) in the z' game:

1. Payoff invariance. As shown in Lemma 5, payoffs at any full history depend only on the full history of markup actions fh_t, the buyer and seller types, b and s, and the bargaining cost terms. Therefore, payoffs after the same markup actions fh_t will be the same in the z game and the z' game.

2. Information sets. Because players observe only their own value (b or s) and the markup actions h_t (or fh_t in the case of the seller), the fact that total actions are $z'-z$ higher in the

z' game than in the z game does not change the information (h_t, b) of the buyer or (fh_t, s) of the seller.

3. Belief preservation. At each history h_t, the buyer's beliefs about the seller type are the same in the z and z' games. Similarly for the seller at each fh_t.

4. Sequential rationality & Bayes Rule. For any full history fh_t, no profitable deviations exist under $\sigma_t^{B*}(b, \kappa, h_t)$ and $\sigma_t^{S*}(b, \kappa, fh_t)$ in the z game, and because payoffs, information sets, and beliefs are the same in the z and z' game, it is also true that no profitable deviations exist under the z' game. Bayes' Rule holds for (μ_{B*}, μ_{S*}) at every on-path full history in the z' game because it holds for (μ_{B*}, μ_{S*}) at every on-path full history in the z game.

Note that r, the seller's secret reserve price, plays a unique role relative to other continuous actions in the game. It is chosen before any publicly observed actions take place (so it only depends on the seller's private value) and it never directly determines the final price. Rather, it functions as a lever determining whether trade occurs through the auction (when $p^A \geq r$) or not). This lever is also unaffected by the realization of $Y'\gamma + W$. In the z game, the condition determining whether trade occurs through the auction, written in terms of the total auction price and reserve price, is $p^A + z \geq r + z$, and in the z' game, this condition is $p^A + z' \geq r + z'$. As in the payoff arguments in Lemma 5, z and z' fall out of the condition. □

B Additional Results Related to Sections 3–4.

We first describe additional data-cleaning steps. In 9.66% of threads, a buyer who is not the high bidder contacts the auction house with an offer to be considered if negotiations with the high bidder fall through or if the seller (or auctioneer) already rejected the auction price. We drop these threads. We also drop observations in which the following variables lie outside their respective 0.005 and 0.995 percentiles: the auction price, reserve price, book price, and final price normalized by book price. We drop observations for which key variables are missing or a bargaining sequence is clearly misrecorded or incomplete. Finally, we drop a small number of bargaining threads (less than 0.02% of threads) that extend longer than one week.

As highlighted in Section 3, when the auction price falls short of the reserve price, the buyer may sometimes walk away before negotiations begin. In observations where this is indicated in the data, a mediator identity is sometimes recorded, suggesting this decision may be influenced by a

Table B.1: Descriptive Statistics at Bargaining-Thread Level by Auction House Location

	Auction House 1	Auction House 2	Auction House 3	Auction House 4	Auction House 5	Auction House 6
Agreement reached	0.600	0.560	0.514	0.700	0.645	0.630
Final price ($)	5,193	7,603	5,657	5,054	5,803	5,760
Book price ($)	6,089	9,082	7,184	6,561	6,852	7,135
Auction price ($)	4,890	7,662	5,543	4,994	5,597	5,880
Reserve price ($)	6,574	9,228	7,690	6,682	7,122	7,703
No reserve	0.265	0.126	0.275	0.264	0.154	0.312
# Offers in a thread	1.346	1.577	1.593	1.215	1.435	1.426
Length of a thread (hours)	7.719	5.928	5.357	5.777	4.131	5.073
Fleet/lease car	0.396	0.446	0.440	0.536	0.530	0.475
Car age (years)	7.057	4.938	6.584	6.529	6.623	5.985
Mileage	97,257	75,367	94,090	97,548	97,549	91,050
Engine displacement (liters)	3.359	3.789	3.523	3.618	3.803	3.646
No. Mediators	16	19	11	31	10	27
No. Sellers	975	746	614	1,802	778	1,311
No. Buyers	1,065	725	441	1,894	701	1,432
No. Threads	13,846	6,998	9,057	19,130	6,766	19,293

Notes: This table shows the means from Table 1 separately by auction house, as well as the number of mediators, sellers, buyers, and threads at each auction house.

mediator, but text notes (available only for some observations in the data) often indicate that, in such cases, the mediator is recording a walk-away decision that happened prior to any mediator involvement. We focus our analysis in Section 4 only on stages of the game occurring *after* the buyer has passed the opt-out stage of the game. Section 5 models (and uses data from) these walk-away decisions as a strategic decision of the buyer taking place before mediator involvement.

B.1 Descriptive Statistics by Auction House Location. In Table B.1, we report means from Table 1 separately by the six auction house locations. These locations differ somewhat in their volume, inventory, and negotiation outcomes. Locations 1, 4, and 6 have the most negotiations and most distinct buyers and sellers. About half of all mediators are at locations 4 or 6. Average agreement rates differ across locations, ranging from 0.514 to 0.70. Location 2 stands out with newer (lower car age and higher mileage) and higher-priced cars. The fraction of cars sold by fleet/lease sellers is lowest at location 1 (0.396) and highest at location 4 (0.536). Table B.2 reports means at the mediator level, as in Table 2, separately by auction house location. The average years of employment for a mediator range from 1.547 in location 6 to 8.806 in location 5. Location 3 has the smallest fraction of female mediators (0.222), and location 5 the largest (0.6).

B.2 Threads Handled by Multiple Mediators. Our main sample limits to threads for which a single mediator handles a given thread. We now examine statistics for threads where one mediator handles the first offer of the bargaining thread but at some point in the thread a different mediator

Table B.2: Descriptive Statistics at Mediator-Thread Level by Auction House Location

	Auction House 1	Auction House 2	Auction House 3	Auction House 4	Auction House 5	Auction House 6
Agreement reached	0.572	0.540	0.518	0.663	0.632	0.625
Final price/book price ($)	0.816	0.884	0.816	0.808	0.868	0.833
Final price/reserve price ($)	0.788	0.839	0.766	0.764	0.793	0.795
Final price/auction price ($)	1.013	1.027	1.032	1.024	1.022	1.017
Fleet/lease car	0.421	0.385	0.324	0.491	0.436	0.475
Female	0.500	0.353	0.222	0.517	0.600	0.462
# Threads mediated	865.375	368.316	823.364	617.097	676.600	714.556
Years of employment	4.218	3.451	2.191	6.684	8.806	1.547
No. Mediators	16	19	11	31	10	27

Notes: This table shows the means from Table 2 separately by auction house.

takes over. Table B.3 replicates Table 1 among threads handled by multiple mediators, showing that the latter correspond to newer and more expensive cars; tend to take longer, both in terms of the number of offers in the thread and the length of time; and are also less likely to reach agreement, by about 41 percentage points. This agreement disparity persists even after controlling for all controls (other than mediator identities) in spec 6 of (3).[75] These threads may represent more challenging (less-likely-to-agree) cases that are passed from one mediator to another.

Table B.3: Threads with Multiple Mediators

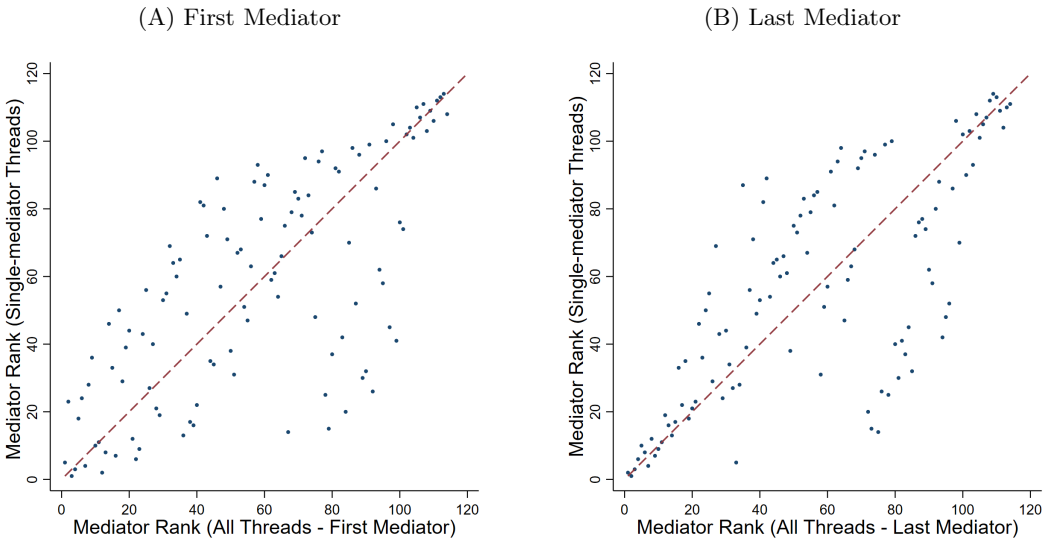
	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.212	0.408	0	1
Final price	6,958	4,962	1,500	13,750
Bluebook price	7,680	5,098	1,925	14,750
Auction price	5,896	4,664	950	12,500
Reserve Price	7,853	5,117	2,300	15,000
# Offers in a thread	2.401	0.802	2	4
Length of a thread (hours)	12.626	21.529	1.47	28.2
Fleet/lease car	0.575	0.494	0	1
Car age (years)	5.394	3.533	1	11
Mileage	86,313	50,100	26,912	152,006
Engine displacement (liters)	3.648	1.443	2	5.7
No. Threads	37,486			

Notes: Table shows descriptive statistics as in Table 1 but for threads handled by multiple mediators.

As a test of whether the exclusion of multiple-mediator threads is consequential, we replicate spec 6 of (3) using an augmented sample that pools single-mediator and multiple-mediator threads.

[75]This is based on a regression (not shown) of the agreement indicator on these controls using all single- and multiple-mediator threads.

Figure B.1: Multiple-Mediator Threads vs. Single-Mediator Threads



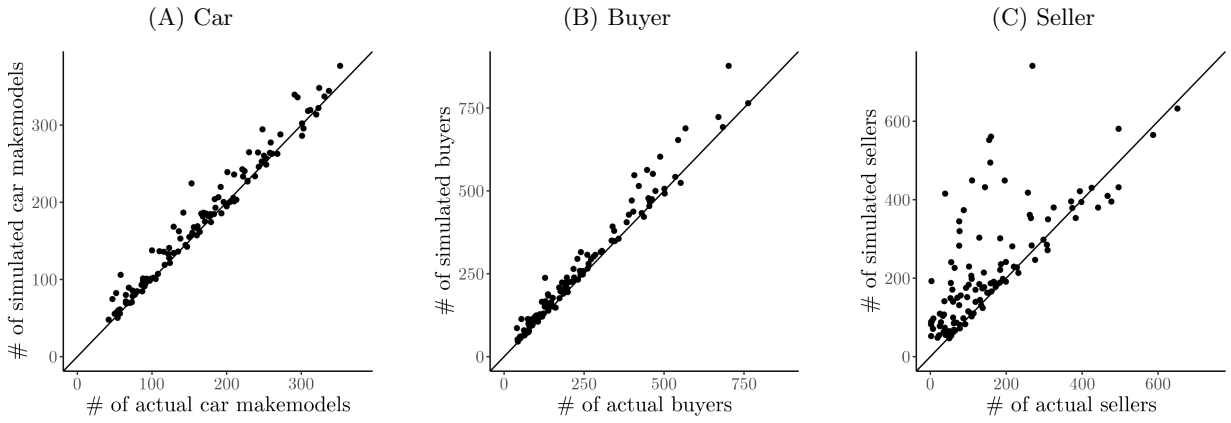
Notes: Each plot shows mediators' ranks based on the estimated mediator fixed effects from spec 6 of (3) using single-mediator threads (on the vertical axis) and compares it to rankings (on the horizontal axis) constructed from an augmented sample that includes threads handled by a single mediator and threads handled by multiple mediators. For estimates in the left (right) panel, we code the mediator who starts (ends) a thread as though she handles the whole thread. We add a 45 degree dashed line as a reference.

We do this in two different ways. First, we code the mediator who *starts* a thread as though she handles the whole thread. We then repeat but code the mediator who *ends* thread as though she handles the whole thread. Figure B.1 shows a scatter plot of the main mediator rankings (spec 6 of (3)) against these alternative versions, with the starting-mediator results in the left panel and the ending-mediator results in the right panel. In both cases, we observe a strong positive correlation with our main mediator rankings, suggesting that the main rankings are not substantially driven by our choice to omit multiple-mediator threads. We focus only on single-mediator threads in the body of the paper because these are the threads where it is more straightforward to attribute outcome variation to mediator variation.

B.3 An Alternative Assignment Test. Here we randomly shuffle a car-type (make-by-model) identifier across observations within a given auction house location, year, and month. We repeat this 500 times. We perform an analogous procedure shuffling buyers or sellers instead of car types. Figure B.2 plots the number of unique car makes, buyers, and sellers that a mediator interacts with in the simulated data, against the mean number of unique car makes, buyers, and sellers that the

mediator interacts with in the real data. Each data point represents one mediator. If mediators are indeed randomly assigned to buyers, for example, mediators should interact with roughly the same number of unique buyers in the real data as in the shuffled data, so all points in panel (B) should lie close to the 45-degree line; conversely, if buyers are assigned mediators non-randomly, a mediator should see *more* unique buyers in the shuffled dataset than in the real dataset.

Figure B.2: Random Assignment Test



Notes: Each data point is a mediator. The x-axis shows the number of car make-models (panel A), buyers (panel B), or sellers (panel C) that a mediator interacts with in the data. The y-axis shows the corresponding numbers of interactions in the randomly reshuffled data.

In panels (A) and (B), most points lie close to the 45-degree line, implying that assignment of mediators to car types and to buyers is approximately random by this measure. Panel (C) shows that mediators interact with fewer sellers in the real data than they do in the shuffled data, suggesting that assignment of mediators to sellers is less likely to be completely random, and highlighting the importance of the seller fixed effects we include in our preferred specification, spec 6 of (3).

B.4 Controlling for Previous Interactions Between Mediators/Agents. Here we examine the extent to which heterogeneity is due to repeated relationships between mediators and agents. For example, it is possible that mediators we identify as highly skilled are only effective when working with particular buyers. To explore this possibility, we re-estimate spec 6 of (3) (from Figure 3) with additional controls for the number of interactions (other than the current thread) between the current mediator and buyer, mediator and seller, or seller and buyer. Table B.4 shows that the coefficients on these interaction terms are quantitatively small and, for most of the terms,

Table B.4: Effects of Repeated Interactions

	(1)	(2)	(3)	(4)
Mediator-buyer interactions	-0.000165*			-0.000226**
	(0.000093)			(0.000105)
Mediator-seller interactions		0.000025		0.000026
		(0.000024)		(0.000025)
Buyer-seller interactions			0.000234	0.000549
			(0.000302)	(0.000350)
Mean(interactions)	2.772	4.183	1.403	.
Median(interactions)	1	1	1	.
N	71,093	71,093	71,093	71,093

Notes: Table shows estimates of (3) under spec 6 with additional controls for number of previous interactions between pairs of agents. An observation is a thread. The outcome variable is the agreement dummy. Mean(interactions) refers to the mean number of interactions between a given mediator-buyer pair, mediator-seller pair, or buyer-seller pair (and similarly for the median). The number of observations is fewer than in Table 1 because singleton fixed effect cells are dropped from the analysis. Significance levels: *: $p < 0.10$, **: $p < 0.05$, and ***: $p < 0.01$.

Table B.5: Mediator Performance Variation by Gender

	(1) Agreement	(2) Agreement	(3) Agreement
Gender(Female=1)	0.112***	-0.003	-0.003
	(0.004)	(0.005)	(0.014)
Specification	No Controls	All Controls	All Controls
R2	0.013	0.386	0.386
N	73,262	69,290	69,290

Notes: Table shows differences in agreement by mediator gender. In column 1 we regress thread-level agreement on mediator gender, with no controls. Columns 2–3 include all controls from spec 6 of (3) except for mediator fixed effects. Columns 1–2 report heteroskedasticity-robust standard errors, and column 3 mediator-level clustering. The number of observations is fewer than in Table 1 both because the gender indicator is not available for some observations and because singleton fixed effect cells are dropped from the analysis.

not statistically significant at conventional levels. We also find that mediator fixed effects estimated under the specifications in Table B.4 are nearly perfectly correlated with our main estimates.

B.5 Differences in Agreement by Gender. Table B.5 shows differences in agreement by mediator gender. In column 1 we regress thread-level agreement on mediator gender, with no other controls. Columns 2–3 include all controls from spec 6 of (3) except for mediator fixed effects. Columns 1–2 report heteroskedasticity-robust standard errors, and column 3 uses clustering at the mediator level. Column 1 suggests that female mediators outperform male mediators, but after including other controls this effect is small and statistically indistinguishable from zero.

B.6 Quantifying Variation Explained by Mediators and Other Factors. In Table B.6 we display the adjusted R^2 of regression (3) under various scenarios, both with the agreement indicator

as the outcome variable and with final price (normalized by book price) as the outcome variable. The controls in our regressions, including the fixed effect indicators, are correlated, and thus the implications for the adjusted R^2 from additional controls vary greatly depending on the *order* in which we add them. In this analysis, we take two different approaches to quantify the contribution of a given set of controls to the adjusted R^2. Panel A reports the adjusted R^2 when we only include our baseline controls (reserve price, auction price, book price, car age, mileage, engine displacement, and date fixed effects), as in spec 1 of (3) and then when we add only one set of fixed effects to those baseline controls. Panel B shows the adjusted R^2 from spec 6 of (3) and with one set of fixed effects *removed* from that full model. Qualitative implications are similar in both panels. For the agreement regressions, we find that buyer, seller, and mediator fixed effects all change the adjusted R^2 noticeably, with seller fixed effects having the largest impact, and other fixed effects have little impact on the adjusted R^2. For example, panel A shows that adding seller fixed effects to the baseline controls increases the adjusted R^2 from 0.0603 to 0.2612, whereas adding mediator fixed effects to the baseline controls increases the adjusted R^2 from 0.0603 to 0.195. For the price regressions, make-model fixed effects and buyer fixed effects have the largest effects, but the shifts are smaller than for the agreement regressions.

Table B.6: Adjusted R^2 Under Various Specifications

Panel A	Agreement Adjusted R^2	Final Price Adjusted R^2
Baseline	0.0603	0.4077
Baseline + Make-Model FE Only	0.0689	0.4486
Baseline + Auction House Location FE Only	0.0681	0.4081
Baseline + Buyer FE Only	0.1091	0.4481
Baseline + Seller FE Only	0.2612	0.4339
Baseline + Mediator FE Only	0.195	0.4099

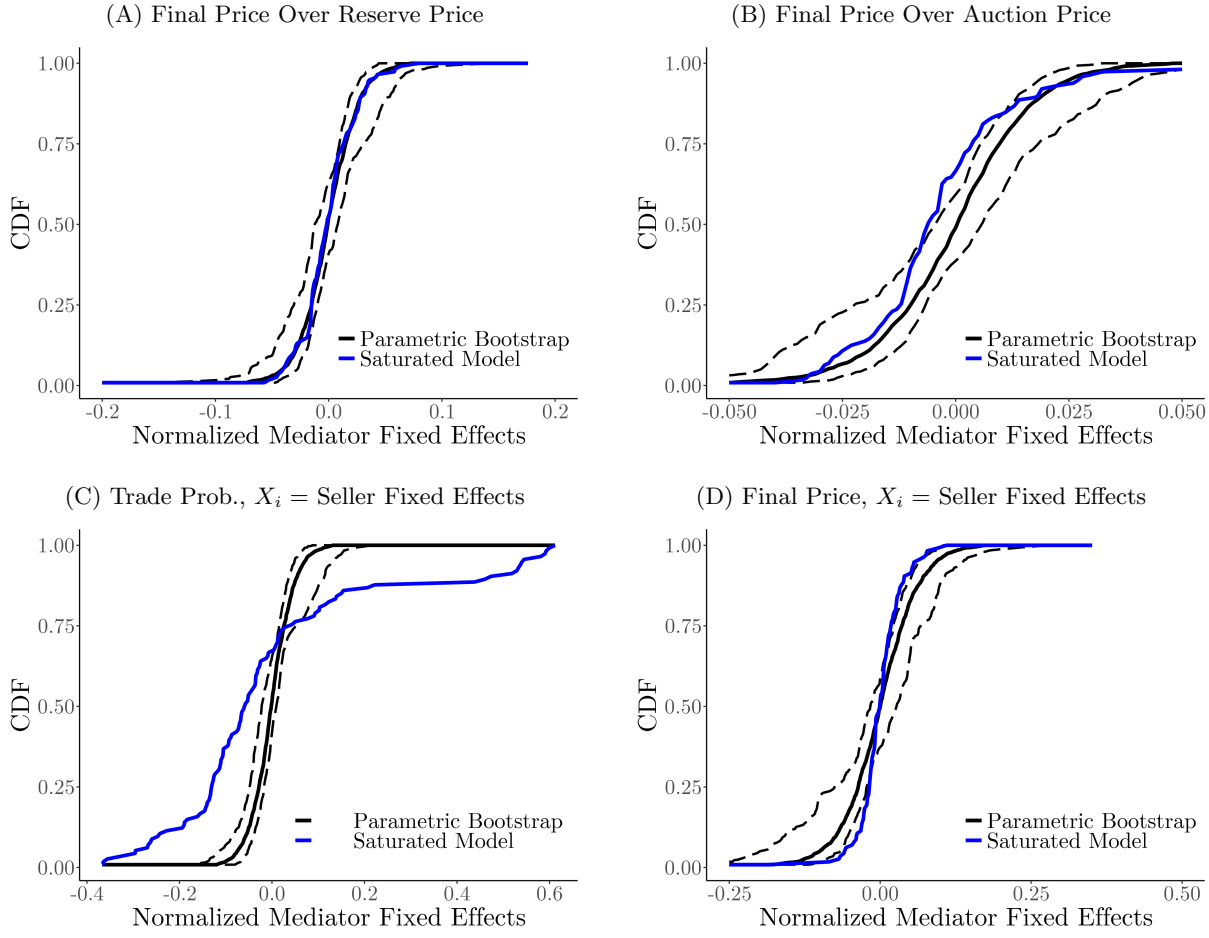
Panel B	Agreement Adjusted R^2	Final Price Adjusted R^2
All variables from Spec 6	0.3074	0.4908
All except Make-Model FE	0.3046	0.4632
All except Auction House Location FE	0.3074	0.4908
All except Buyer FE	0.286	0.4674
All except Seller FE	0.2247	0.4798
All except Mediator FE	0.2861	0.4906

Notes: Table shows adjusted R^2 from estimating (3) under various scenarios. Panel A shows adjusted R^2 when we only include our baseline controls (reserve price, auction price, book price, car age, mileage, engine displacement, and date fixed effects) and then when we add various sets of fixed effects to those baseline controls. Panel B shows the adjusted R^2 in the most saturated specification (spec 6) and then again with one set of fixed effects removed.

B.7 Alternative Bootstrap Tests of Variation Due to Sampling Error. Figure B.3 panels A and B show results of the bootstrap procedure for prices as described in Section 4.3 (using X_i from spec 6 of (3)) but normalizing by reserve prices (panel A) or auction prices (panel B) rather than book prices. Figure B.3 panels C and D show results of the bootstrap test from Section 4.3 but using only seller fixed effects in X_i.[76] The results are similar to those in the main text: the distribution of mediator fixed effects for trade probabilities lies well outside the bootstrap confidence intervals, whereas the distribution of fixed effects for prices lies largely within the confidence intervals.

[76]Recall that seller (and buyer) identifiers in our data are unique for each auction house, and thus seller (or buyer) fixed effects also absorb auction house fixed effects.

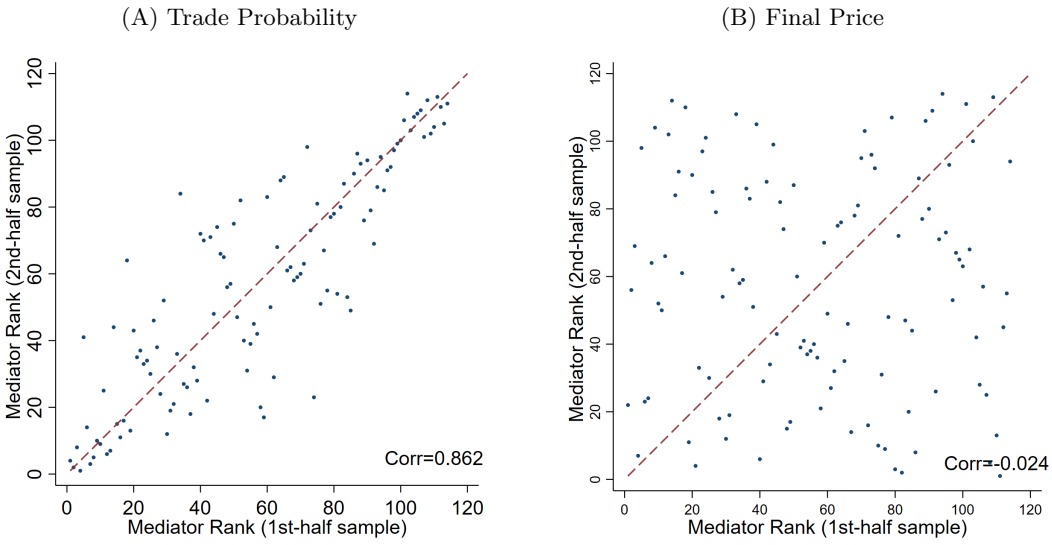
Figure B.3: Alternative Bootstrap Tests of Heterogeneity from Sampling Error



Notes: Panels A and B show results as in Figure 6.B but where the outcome is final price divided by reserve price in panel A and final price divided by auction price in panel B. Panels C and D show results as in Figure 6 but where we only use seller fixed effects (which also absorb auction house location fixed effects) to predict trade probabilities and prices. In each case, we use 100 bootstrap samples. Blue lines show the estimated distribution of fixed effects from the actual data and black show the bootstrap median and 95% pointwise confidence bands.

B.8 Split-sample Test of Fixed Effect Significance. As an alternative test of whether estimated mediator heterogeneity is due to random chance, we randomly split the main dataset into two halves and estimate mediator fixed effects using the first and second halves separately. Figure B.4 plots the estimated mediators' rank from the first subsample on the x-axis and the rank from the second subsample on the y-axis. Under the null hypothesis that mediators have no effect on trade outcomes, the entirety of the estimated dispersion in mediator fixed effects would be driven by sampling error. As a result, in the split-sample experiment, mediators' ranks in the two samples

Figure B.4: Split-sample Test of Fixed Effect Significance



Notes: Each data point is a mediator. The x-axis shows mediators' rank in the first-half sample. The y-axis shows mediators' rank in the second-half sample. Dashed line is 45-degree line. Correlation coefficients are also shown. Trade probability fixed effects are in panel A and price fixed effect in panel B.

would be uncorrelated. Panel A shows a strong positive correlation (with a correlation coefficient of 0.862) for trade probability fixed effects. Panel B shows no correlation (-0.024) for price fixed effects. These results are consistent with our findings that mediators primarily differ in their ability to help agents reach agreement, not in the prices at which trade occurs.

C Additional Model and Estimation Details for Welfare Analysis

Before fleshing out estimation details, we first make several remarks regarding the model. While we must choose a specific ordering for the purposes of our model, several aspects of the end of period 1 and the beginning of period 2 can sometimes differ in practice from the way we model them. First, in some cases, the high bidder chooses to walk away soon after the auction, while in other cases he waits several hours to call the auction company to opt out (as long as the seller has not yet accepted or rejected P^A). Our data do not allow us to distinguish between cases, but we can see whether agents agree or not, and at what price, which suffices for estimating model primitives.

Second, there are some sales in which the auction house states that P^A is considered *binding*, meaning the highest bidder is not allowed the initial walk-away option. We do not have clear indicators in the data as to which observations had this binding feature, although we are told this

is quite common for fleet/lease sales; this is one motivation for our exploration of results separately by dealers and fleet/lease sellers in Appendix D. Incorporating binding P^A would complicate the model by making bids depend on anticipated bargaining outcomes, potentially incentivizing bidders to shade their bids to avoid bargaining costs. A binding auction price scenario also complicates the proof of strictly increasing secret reserve prices, as buyers no longer have the option to opt out. For simplicity, we treat all auctions as though, when $R > P^A$, the highest bidder is always given the option to walk away from bargaining *before* the seller can accept or reject P^A.

Third, as described in Section 3, there are some cases where the seller does not report a reserve price, usually meaning the seller wishes the auction house to contact her regardless of the outcome, similar to a situation with a binding P^A. Several steps in our identification and estimation strategy require observing reserve prices, so we do drop these observations except in controlling for observable heterogeneity (as described in Section 5.5) and we do not model the seller's extensive-margin decision to report a secret reserve price vs. not. As described in Section 3, agreement is equally likely in observations with or without reserve prices reported.

We now turn to estimation details. Estimation requires several assumptions beyond those described in the model setup. Below, let F_R, F_{P^A}, and F_W represent CDFs of R, P^A, and W. Define the *full game* as the game beginning with period 0 (as described in Section 5.2) and the *bargaining stage* of the game as portion of the game beginning after a buyer chooses not to opt out of bargaining and the mediator's type κ is chosen by nature and revealed to all parties.

(A9) *F_R, F_{P^A}, and F_W have densities f_R, f_{P^A}, and f_W satisfying the following: (i) the characteristic functions of f_R and f_W have only isolated real zeros; (ii) the real zeros of the characteristic function of f_{P^A} and the real zeros of its derivative are disjoint; and (iii) E[W]=0.*

(A10) *Random variables $(S,\{B_j\}_{j=1}^N,W,N)$ are independently and identically distributed across instances of the game.*

(A11) *All observations in the data arise from the same equilibrium of the full game. All observations in the data with the same assigned mediator type κ are generated by the same equilibrium strategies of the bargaining stage of the game.*

*A*9 lists sufficient conditions from Evdokimov and White (2012) for proving identification of f_R, f_{P^A}, and f_W from the joint distribution of $R + W$ and $P^A + W$.[77] Evdokimov and White

[77]For a random variable W the characteristic function evaluated a point ν is given by $q_W(\nu) = E[e^{i\nu W}]$, where

(2012) note that these assumptions are weaker than conditions originally assumed in the empirical auction literature to handle game-level heterogeneity (Krasnokutskaya 2011). $A10$ is common in the empirical literature on games, abstracting away from dynamics *across* instances of the game. Versions of $A11$, a data generation assumption, are common in structural work. In particular, in a setting with multiple equilibria, to estimate equilibrium objects, the researcher must typically assume the data are generated by a single equilibrium. In our setting, the assumption is that equilibrium objects that do not depend on mediator assignment (such as ρ and χ) are generated by a single equilibrium across all observations in the data and that equilibrium objects that do depend on mediator assignment (g_κ) are generated by the same strategies within the set of observations in which the assigned mediator was κ.

C.1 Observed Heterogeneity. For observation i, the total reserve and auction prices, denoted R_i^{total} and $P_i^{A,total}$, are modeled as

$$\begin{bmatrix} R_i^{total} \\ P_i^{A,total} \end{bmatrix} = \begin{bmatrix} Y_i'\gamma \\ Y_i'\gamma \end{bmatrix} + \begin{bmatrix} \tilde{R}_i \\ \tilde{P}_i^A \end{bmatrix}, \tag{26}$$

where $\tilde{R}_i = R_i + W_i$, $\tilde{P}_i^A = P_i^A + W_i$. As in the standard *homogenization* approach (Haile et al. 2003), we estimate γ through a linear regression of reserve and auction prices on observables Y_i, and the residuals provide an estimate of $\tilde{R}_i$ and $\tilde{P}_i^A$. Variation in these two quantities can then be attributed to unobserved game-level heterogeneity and to players' private values.

For this linear regression, we include a rich set of controls, far more than in our specifications in Section 4 because it comes from the broader dataset of all cars running through the auction mechanism (not just observations that reached mediated negotiation), where more controls are available, and because a rigorous set of controls helps to make the independence assumptions of the model as plausible as possible. As in Larsen (2021), Y_i includes the following: fifth-order polynomials in the book price, odometer reading, run number within an auction-house-by-day combination, and run number within an auction-house-by-day-by-lane combination (where run number refers to the order in which cars are auctioned); the number of previous attempts to sell the car; the number of pictures displayed for the car on the auction company website; a dummy for whether or not the odometer reading is considered accurate, and the interaction of this dummy with the odometer

$i = \sqrt{-1}$. The characteristic function completely characterizes the density of W. The term *isolated real zeros* implies that, at any point ν at which $q_W(\nu) = 0$, a neighborhood exists containing no other points at which $q_W(\cdot)$ is zero.

reading; the interaction of the odometer reading with car-make dummies; dummies for each make-model-year-trim-age combination (where age refers to the age of the vehicle in years); dummies for condition report grade (ranging from 1-5, observed only for fleet/lease vehicles); dummies for the year-month combination and for auction house location interacted with hour of sale; dummies for 32 different vehicle damage categories recorded by the auction house; dummies for each seller who appears in at least 500 observations; dummies for discrete odometer bins (four equally sized bins for mileage in $[0, 20000)$, eight equally sized bins for mileage in $[20000, 100000)$, four equally sized bins for mileage in $[100000, 200000)$, one bin for mileage in $[200000, 250000)$, and one bin for mileage greater than 250000); several measures of the thickness of the market during a given sale.[78]

Note that, in this homogenization regression, observable heterogeneity, $Y_i'\gamma$, enters both auction and reserve prices identically, a result that follows from Proposition 5. To examine the validity of this property, we perform the homogenization step in two separate regressions, one using only auction prices as the outcome and one using only reserve prices, using the same vector of observables Y_i in each case. This exercise yields predicted values that we denote $Y_i'\hat{\gamma}_{P^A}$ and $Y_i'\hat{\gamma}_R$. The correlation between the $Y_i'\hat{\gamma}$ we use in the body of the paper and either of these two different predicted values is about 0.99. The correlation between residual auction prices using these two different regressions (that is, the correlation between $P_i^A - Y_i'\hat{\gamma}_{P^A}$ and $P_i^A - Y_i'\hat{\gamma}_R$) is 0.7812. The analogous correlation for reserve price residuals is 0.7732. Thus, the outputs from this process (which are then used as inputs for Step 2 of our estimation) are highly correlated regardless of whether we use reserve or auction prices as the regression outcome with which we estimate γ. Our main observable heterogeneity results pool both outcomes together.

C.2 Unobserved Heterogeneity. We apply a result due to Kotlarski (1967), which implies that observations of $\tilde{R} = R + W$ and $\tilde{P}^A = P^A + W$ are sufficient to recover f_W, f_R, and f_{P^A}. We estimate these using a flexible likelihood approach. The likelihood of the joint density of $(\tilde{R}, \tilde{P}^A)$ is

$$\mathcal{L}(f_{P^A}, f_R, f_W) = \prod_i \left[\int f_{P^A}(\tilde{p}_i^A - w) f_R(\tilde{r}_i - w) f_W(w) dw \right] \tag{27}$$

We approximate each of the densities f_{P^A}, f_R, and f_W using fifth-order Hermite polynomials (as in

[78] These market thickness measures are computed as follows: for a given car on a given sale date at a given auction house, we compute the number of remaining vehicles still in queue to be sold at the same auction house on the same day lying in the same category as the car under consideration. The six categories we consider are make, make-by-model, make-by-age, make-by-model-by-age, age, or seller identity.

Gallant and Nychka 1987). A similar identification argument was applied in Li and Vuong (1998) and then applied to unobserved auction-level heterogeneity first by Krasnokutskaya (2011). The authors use kernel estimation approaches rather than a likelihood approach. We found the flexible likelihood approach stable and straightforward, as in Freyberger and Larsen (2022).

We now remark on a point related to game-level heterogeneity. As highlighted in Sections 5.2 and 5.6, the objects S, B_j, R, and P^A can be positive or negative, as they describe how far above or below the game-level heterogeneity component ($Y'\gamma + W$) the *total* values and reserve/auction prices are. If $Y'\gamma + W$ is not sufficiently positive and R or P^A is sufficiently negative, the total reserve or auction price may be negative, which is unrealistic. To evaluate how often this occurs, we simulate random draws of $Y'\gamma$, W, R, and P^A from our estimates. We find negative total reserve prices in 0.8% of cases and negative total auction prices in 2.54% of cases. A similar simulation implies negative total *values* for any buyer (including losing bidders) in 3.7% of cases, negative total values for the highest-value bidder in 0.7% of cases, and negative total values for sellers in 17.8–33.0% of cases, depending on whether the lower or upper bound on F_S is used to simulate. Thus, the occurrence of negative total seller values is nontrivial under our estimates, but these total seller values very rarely imply negative total reserve prices. Moreover, these negative quantities do not lead to problematic or unrealistic predictions from our structural exercise because the key outcome of this exercise is surplus (e.g., Figure 9), not prices or values alone, and Lemma 5 shows that buyer and seller surplus in this game are invariant to the realization of $Y'\gamma + W$. Specifically, a seller's surplus involves the price minus the seller's value and a buyer's surplus involves the buyer's value minus the price. The game-level heterogeneity term falls out of surplus calculations and the surplus conditional on trade occurring is positive because trade only occurs at prices that lie between the seller's and buyer's value.

C.3 Buyer Values. The left-hand side of (5) is F_{P^A}, which is estimated in Step 2. The object $\Pr(N = n)$ is the distribution of the number of bidders in the auction. Larsen (2021) discussed several approximations for $\Pr(N = n)$ and showed that welfare estimates are relatively insensitive to these choices. In particular, Larsen (2021) demonstrated that, in a certain class of specifications for $\Pr(N = n)$ (Poisson), the distribution of valuations for the bidder who enters the bargaining game is analytically invariant to the distribution of N. Moreover, even under non-Poisson distributions, Larsen (2021) showed that the approximation for $\Pr(N = n)$ makes little difference for final welfare

estimates: the choice of $\Pr(N = n)$ affects estimates of F_B, but has little to no effect on the distribution of buyer values *conditional* on the auction price, $\frac{F_B(b)}{1-F_B(\chi(p^A))}$, which is what matters for evaluating (9). Following Larsen (2021), we approximate $\Pr(N = n)$ using the distribution of the lower bound on the number of bidders in each auction. This lower bound is observable on an auction-by-auction basis for a subset of auctions for which we have *bid logs*, recording each bid placed during the auction process. We observe these bid logs for a total of 113,497 auctions (primarily fleet/lease sales). With estimates of F_{P^A} and $\Pr(N = n)$, F_B is nonparametrically estimable by solving (5) on a grid of points v. The estimate of F_B is shown in Figure C.2.A.

C.4 Seller Values. To estimate lower and upper bounds on F_S from Proposition 2, but taking into account unobserved heterogeneity, we parameterize each bound as a flexible piecewise linear spline, denoted $F_S^L(\cdot, \theta^{S,L})$ and $F_S^U(\cdot, \theta^{S,U})$. Denote the fixed vector of spline knots $\{v_k^S\}_{k=1}^{K_S}$, which we choose to be $K_S = 200$ uniformly spaced knots between the 0.001 and 0.999 quantiles of $\tilde{P}^A$. We use minimum distance to estimate $\theta^{S,L}$ and $\theta^{S,U}$:

$$\min_{\theta^{S,L},\theta^{S,U}} \sum_{k=1}^{K_S} \left\{ \left[\widehat{\Pr}(D^S = a|\tilde{P}^A = v_k^S) \left(\int \widehat{M}_S(v_k^S, z)dz \right) - \int F_S(v_k^S - w; \theta^{S,L}) \widehat{M}_S(v_k^S, w)dw \right]^2 \right.$$
$$\left. + \left[\widehat{\Pr}(D^S \neq q|\tilde{P}^A = v_k^S) \left(\int \widehat{M}_S(v_k^S, z)dz \right) - \int F_S(v_k^S - w; \theta^{S,U}) \widehat{M}_S(v_k^S, w)dw \right]^2 \right\} \quad (28)$$

Here, $M_S(v, w) \equiv f_{P^A}(v - w) f_W(w)$ is the joint density of P^A and W, an estimate of which can be constructed using the estimated densities from Step 2. The objects $\Pr(D^S = a|\tilde{P}^A = v)$ and $\Pr(D^S \neq q|\tilde{P}^A = v)$ are, respectively, the probabilities that a seller accepts or does not quit, conditional on the realization of $\tilde{P}^A$. We estimate these objects via local linear regressions.

Following Larsen (2021), we impose the following constraints in (28): (i) F_S^L lies graphically above F_R and graphically below F_S^U; (ii) F_S^L and F_S^U lie in $[0, 1]$; (iii) F_S^L and F_S^U are weakly increasing; and (iv) $F_S^L(v)$ and $F_S^U(v)$ are equal to 0 for any $v < v_1^S$ and equal to 1 for any $v > v_{K_S}^S$. These last three constraints ensure that F_S^L and F_S^U will correspond to proper distribution functions. Condition (iv) merits additional discussion. By construction, the seller bounds are conditional on a buyer action (the auction price). If the support of auction prices in the data is wide enough that sellers accept or reject extreme auction prices with probability 1, the bounds will be surjective, mapping to each point in [0,1]. Otherwise, while still valid, the bounds may be wide near the tails, failing to reach 0 or 1. We find that this is not an issue for F_S^L, which is surjective, or for the right

tail of F_S^U, which attains a value of 1, meaning condition (iv) does not bind for these cases. It does bind, however, for the left tail of F_U^S. This means that, without condition (iv), we would be unable to reject the true F_S containing mass well below v_1^S. We choose to assign all of this mass exactly at v_1^S. While somewhat arbitrary, this choice has the following motivation: Because v_1^S approximates the lowest auction price (and hence, the lowest realization of B), condition (iv) essentially imposes that $\underline{s} \geq \underline{b}$, i.e., the lowest price any seller would accept is bounded below by the lowest buyer value. More importantly, however, Figure 8 suggests that the left tail of F_S is relatively inconsequential for our comparison of efficiency for different mediator skill levels, as differences across mediators are small when dealing with very low-value sellers. Figure C.2.A shows estimates of F_S^L and F_S^U, as well as F_B. F_B (and hence, F_{P^A}) has very little mass below -$5,000, illustrating the point that, at these low values, an assumption is needed to truncate F_S^U.

C.5 Belief Updating. Taking unobserved heterogeneity into account, (7) becomes

$$\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R}) = \int \frac{1 - F_B(\chi(\tilde{p}^A - w))}{1 - F_B(\tilde{p}^A - w)} \left(\frac{M_\chi(\tilde{p}^A, w)}{\int M_\chi(\tilde{p}^A, z) dz} \right) dw \tag{29}$$

where $M_\chi(\tilde{p}^A, w) \equiv f_{P^A}(\tilde{p}^A - w)(1 - F_R(\tilde{p}^A - w)) f_W(w)$ is the likelihood of the event $(P^A = \tilde{p}^A - w, \tilde{P}^A < \tilde{R}, W = w)$. We approximate $h_\chi(\cdot) \equiv 1 - F_B(\chi(\cdot))$ as a flexible piecewise linear spline. We estimate $\Pr(D_1^B = 0 | \tilde{P}^A = \tilde{p}^A, \tilde{P}^A < \tilde{R})$ via local linear regression and estimate $h_\chi(\cdot)$ by minimizing the distance between the left and right-hand sides of (29). We then obtain $\hat{\chi}(p^A) = \hat{F}_B^{-1}(1 - \hat{h}_\chi(p^A))$.

C.6 Direct Mechanisms. Taking unobserved heterogeneity into account, (8) becomes

$$\Pr(\mathcal{A} | \tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A, \kappa) = \int \frac{1 - F_B\left(g_\kappa(\tilde{r} - w, \tilde{p}^A - w)\right)}{1 - F_B\left(\tilde{p}^A - w\right)} \left(\frac{M_{g_\kappa}(\tilde{r}, \tilde{p}^A, w)}{\int M_{g_\kappa}(\tilde{r}, \tilde{p}^A, z) dz} \right) dw \tag{30}$$

where $M_{g_\kappa}(\tilde{r}, \tilde{p}^A, w) \equiv f_R(\tilde{r} - w) f_{P^A}(\tilde{p}^A - w) f_W(w)$ is the joint density of (R, P^A, W).[79] We approximate $h_{g_\kappa}(r, p^A) \equiv \frac{1 - F_B(g_\kappa(r, p^A))}{1 - F_B(p^A)}$ using a flexible bilinear spline parameterized by 25 knots in each dimension, uniformly spaced between the 0.001 and 0.999 quantiles of $\tilde{R}$ and $\tilde{P}^A$. We estimate the spline parameters by minimizing the distance between the left- and right-hand sides of (30). We impose several constraints in estimation: (i) $h_g(r, p^A) \in [0, 1]$; $h_g(r, p^A)$ being decreasing in r and

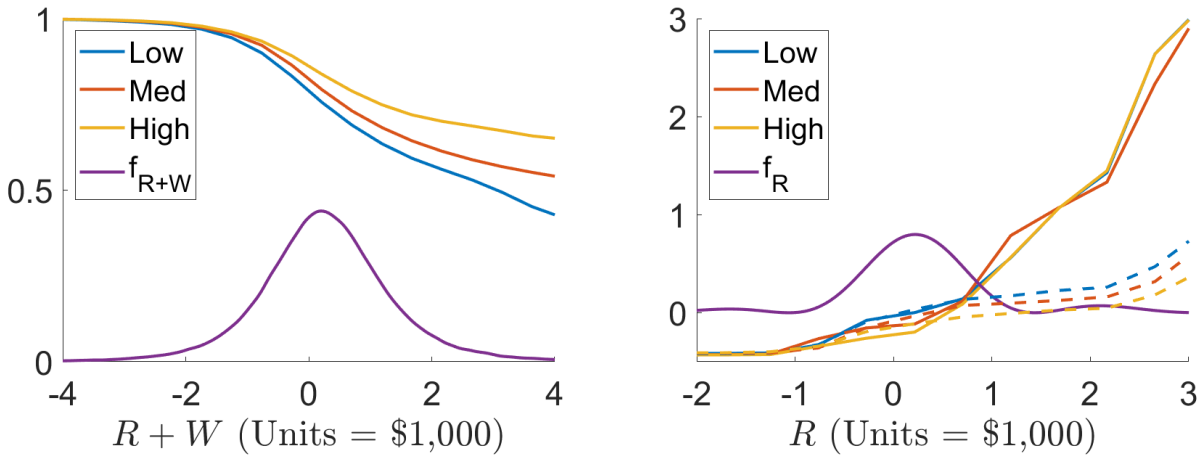
[79] Note that (30) uses an implication of Proposition 5 that, for any κ, trade occurs if and only if $\tilde{b} \geq g_\kappa(\tilde{r}, \tilde{p}^A) \Rightarrow b \geq g_\kappa(\tilde{r} - w, \tilde{p}^A - w) = g_\kappa(r, p^A)$.

(ii) $g_\kappa(r,p^A) \geq \underline{g}_\kappa(r,p^A) \equiv \max\{p^A, \rho^{-1}(r)\} \Rightarrow h_{g_\kappa} \leq \frac{1-F_B(\underline{g}_\kappa(r,p^A))}{1-F_B(p^A)}$. We enforce (ii) by evaluating $\rho^{-1}(r) = F_S^{-1}(F_R(r))$ at $F_S = \hat{F}_S^L$. Condition (ii) ensures that the estimated g_κ does not allow trade when $S > B$.

This requires first estimating the conditional probability $\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A, \kappa)$, which we do using a tensor product of cubic b-spline functions with fifteen uniformly spaced knots in each dimension. With estimates of $\hat{h}_{g_\kappa}(r,p^A)$, we then obtain $\hat{g}_\kappa(r,p^A) = \hat{F}_B^{-1}(1-(1-\hat{F}_B(p^A))\hat{h}_{g_\kappa}(r,p^A))$.

Figure C.1: Average Conditional Trade Probabilities and g_κ Functions (Units = \$1,000)

(A) Conditional Trade Probability Averaged Over $\tilde{p}^A$ (B) $g_\kappa(r,p^A)$ Functions Averaged Over p^A



Notes: Panel A displays estimated $\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A, \kappa)$, integrated over $\tilde{p}^A$, and panel B shows estimated $g_\kappa(r,p^A)$ functions, integrated over p^A, for low- (solid blue), medium- (solid red), and high-skilled (solid yellow) mediators. In panel B, the dashed blue, red, and yellow lines represent the estimated $g_\kappa(r,p^A)$ functions, integrated over p^A, when we relax the constraint that the g_κ functions cannot imply trade when $S > B$. The purple line shows the density of $R+W$ in panel A and of R in panel B. Units on the horizontal axes are \$1,000, and units on the vertical axes for the blue, red, and yellow lines are also \$1,000.

The conditional probability $\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A, \kappa)$ is the key data object that identifies g_κ. Figure C.1.A shows the estimates of this conditional probability, integrated over $\tilde{p}_A$, for each mediator type. Thus, the lines in Figure C.1.A show the estimated probability of agreement conditional on $\tilde{R}$ for different mediator types. The purple line shows the density of $R + W$ to illustrate where the mass lies. Other than at very low values of $R + W$, where the trade probability is similar for different mediator types, we find that a clear ordering of mediator types in terms of their probability of trade, as in Figure 7.B.

In Figure C.1.B we display the estimated $\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A, \kappa)$ functions, integrated over

p^A, in solid blue, red, and yellow lines.[80] The dashed lines in Figure C.1.B show the estimated $g_\kappa(r, p^A)$ functions, integrated over p^A, when we relax the constraint that the g_κ functions cannot imply trade when $S > B$. Together, panels A and B demonstrate that (i) the ordering of the g_κ functions in Figure 8 is driven by these conditional trade probabilities; (ii) the close alignment of the g_κ functions in Figure 8 at low values of R is driven by these conditional trade probabilities; and (iii) the close alignment of the g_κ functions in Figure 8 at high values of R is instead driven by the constraint.

To evaluate the fit of our exercise, and the extent to which the constraint affects this fit, we compute the integrated absolute error (IAE), which is the absolute value of the difference between the left- and right-hand sides of (30), integrated over $\tilde{r}$ and $\tilde{p}^A$. Because both sides of (30) are probability statements, the IAE is a measure of the average probability by which our model fit differs from the conditional probabilities in the data. We find that the IAE is 0.042, 0.049, and 0.058 for low-, medium-, and high-skilled mediators, respectively. When we relax the constraint that g_κ functions cannot imply trade when $S > B$, these IAE metrics improve to 0.014, 0.014, and 0.010, respectively. The improvement in fit from dropping the constraint is small because the region where the constraint makes the largest difference (i.e., where the solid and dashed lines differ most in Figure C.1.B) coincides with a region of little mass in terms of reserve prices.

C.7 Bargaining Cost Bounds. To incorporate bargaining costs into the surplus estimates, we subtract an *upper* bound on expected losses due to bargaining costs from our surplus upper bound in (9). This yields a *lower* bound on surplus measures incorporating bargaining costs. In contrast, the surplus upper bound in (9) does not need to be adjusted, because treating the expected loss due to bargaining costs as though it were zero still yields an upper bound on surplus.

We bound these losses following Larsen (2021). First, bounds on η_B and η_S can be obtained from the difference between consecutive bargaining offers. Consider offers made in periods 1 through 3 of the bargaining game. Because an agent who counters must prefer the situation in which (i) she must pay the cost of countering (η_S or η_B) but her counteroffer is accepted with probability one to the situation in which (ii) she accepts the most recent opponent offer, it must be true that $P_2^S - \eta_S \geq P_1^B$ for the seller and $B - P_3^B - \eta_B \geq B - P_2^S$ for the buyer. Note that if (i) is not preferred to (ii), the agent should not counter. Rearranging these inequalities yields $p_2^S - p_1^B \geq \eta_S$ and $p_2^S - p_3^B \geq \eta_B$. We

[80] In contrast, Figure 8 in the body of the paper shows the estimated $g_\kappa(r, p^A)$ functions evaluated at a particular realization of P^A (specifically, at \$2,000 below the game-level heterogeneity component).

compute the minimum (across all observations) of $p_2^S - p_1^B$ in 200 nonparametric bootstrap samples and use as our (conservative) upper bound on η_S the 0.95 quantile of these minima across bootstrap replications. We follow a similar procedure with $p_2^S - p_3^B$ for an upper bound on η_B.

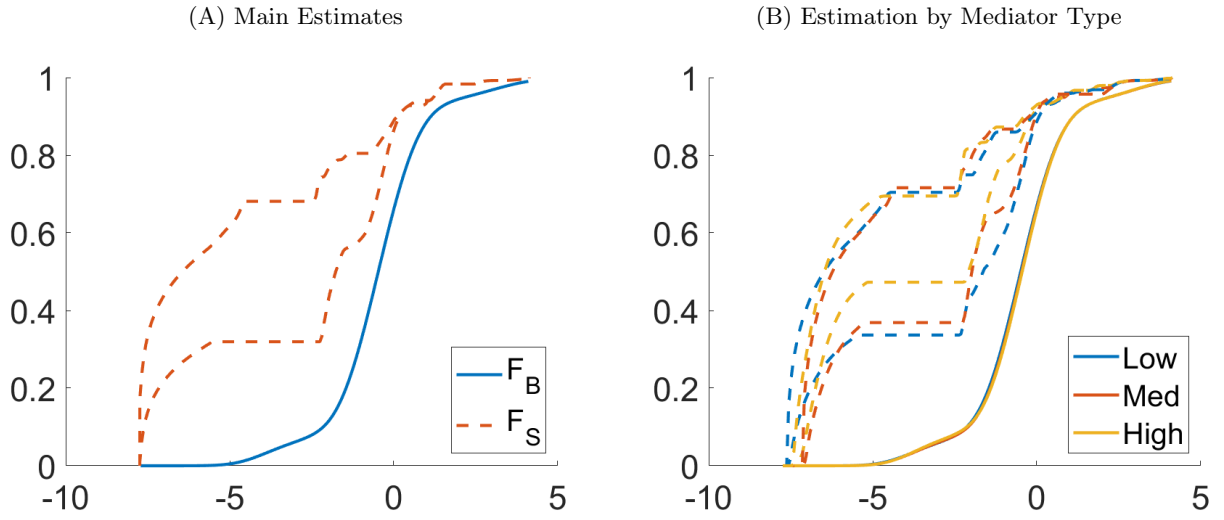
With these bounds on η_B and η_S, we compute upper bounds on the total expected loss for a buyer or seller due to bargaining costs as follows. Recall that the bargaining costs are incurred each time an agent makes an offer. Let $\mathcal{T}$ be a random variable representing the period the game ends. The buyer's and seller's expected loss incurred due to bargaining costs are given by $\eta_B E\left[\lfloor \mathcal{T}/2 \rfloor\right]$ and $\eta_S E\left[\lfloor (\mathcal{T}-1)/2 \rfloor\right]$, respectively, because by period t of the game the buyer has made a total of $\lfloor \frac{t}{2} \rfloor$ offers (where $\lfloor \cdot \rfloor$ is the floor function), and similarly for the seller.[81] Because different mediators may take a different number of periods in bargaining, we compute this upper bound on losses separately for each κ. Our estimated upper bounds on the total loss to buyers due to bargaining costs are \$1, \$40, and \$20 for low-, medium-, and high-skilled mediators. The corresponding estimated bounds for sellers are \$4, \$6, and \$2. The expected number of periods that the games lasts (the delay) are 1.96 periods, 1.91 periods, and 1.84 periods for low-, medium-, and high-skilled mediators.

C.8 Auction House Fees. As explained in Appendix E of Larsen (2021), these fees are largely fixed fees (with only a small price-based commission), and we treat them as fully comprised of a fixed fee. Our notion of buyer and seller values, B and S, embed these fees. Specifically, let h^B and h^S represent the fee paid to the auction house by the buyer and seller if trade occurs, and let B^* be the buyer's willingness to pay ignoring this fee, and S^* the seller's willingness to sell ignoring this fee. The random variables of which we identify (or bound) distributions of $B = B^* - h^B$ and $S = S^* + h^S$. Our estimates of total surplus under different mediators (as well as our calculations of first-best and second-best surplus) do not include any surplus accrued to the auction house from fees. We treat fees as fixed, to be paid in any mechanism, including theoretical benchmarks.

C.9 Estimating Distributions Separately by Mediator Type. Figure C.2.A displays the estimated F_B in the solid blue line and the estimated bounds on F_S in the dashed red line. To explore whether mediator types indeed appear to face the same distribution of buyer and seller values (as assumed in *A*4), Figure C.2.B displays these same objects estimated separately using the pseudo samples described in Section 5.5, one for each mediator type. The different estimates

[81] This expression for the total expected bargaining costs of the buyer includes the initial η_B incurred upon entering bargaining. In contrast, in Appendix A.1, expression (11) writes the expected bargaining cost of the buyer *beyond* this initial η_B. The proof of Proposition 5 contains a similar expression.

Figure C.2: Estimated Value Distributions (Units = \$1,000)



Notes: Panel A displays, in solid blue, the estimated F_B and, in dashed red, the bounds on F_S from the body of the paper. Panel B shows these same objects estimated separately using the different pseudo samples (described in Section 5.5) for different mediator types. In panel B, the solid lines represent estimates of F_B and the dashed lines bounds on F_S estimated separately using the different mediator pseudo samples (represented by different colors). The estimates of F_B in panel B are indistinguishable across mediator types. Units on the horizontal axes are \$1,000.

of F_B in panel B are indistinguishable from one another, lending support to *A*4. For F_S, we only have bounds, and thus we cannot concretely confirm or deny whether the true F_S is the same across mediator types. However, if we were to find large regions of the CDF bounds that do not overlap between different mediator types, this would be evidence that the F_S is *not* the same across mediator types. We find the opposite: the estimated bounds on F_S differ slightly by mediator type but overlap substantially, allowing for a single F_S (faced by all mediators) contained in the bounds.

The small differences we do see in the F_S bounds across mediator types can be explained by how the bounds are constructed. The bounds use all observations of the game, including those that end by a successful auction (where $P^A \geq R$). As explained in the discussion of Proposition 2, the lower bound corresponds to the probability, conditional on P^A, that the seller takes an action at or before the second period of the game that leads to trade; this action is either to choose a low enough R such that $R \leq P^A$ or, when $R > P^A$, to choose to accept the P^A in person or through mediated bargaining over the phone. Only in mediated bargaining (reached by 26% of observations; see Section 3) can mediators influence trade and, consequently, estimated F_S bounds.

Mediator-specific F_S bounds are less useful for our main evaluation of efficiency, where we

evaluate welfare at different choices of F_S within the bounds. Evaluating welfare at *different* bounds for different mediator types, when these bounds overlap, would mistakenly attribute *differences in the bounds* at which we evaluate welfare to differences in actual welfare across mediator types. Instead, we use a fixed set of F_S bounds in evaluating welfare; these bounds average over mediator types, which Section 5.3 shows contain the true CDF of seller values under the model's assumptions.

C.10 Pareto Frontier. Figure C.3 displays the second-best (ex ante efficient) Pareto frontier, which can be referred to as the Myerson-Satterthwaite and Williams frontier, along with the first-best frontier.[82] Dashed lines indicate 95% confidence intervals. We also show, in the colored dots, the estimated surplus/gains from trade in the real world mechanisms corresponding to each mediator type, with lines protruding from those points indicating 95% confidence intervals in each dimension (expected buyer gains on the vertical axis and expected seller gains on the horizontal axis). Panels on the left show the full Pareto frontier, and panels on the right zoom in to show the region surrounding the mediator-specific outcomes. Panels on top are evaluated at the F_S lower bound and panels on the bottom at the F_S upper bound. The results are consistent with higher-ability mediators achieving outcomes that are more efficient (i.e., further to the northeast).

The up-close results using the F_S lower bound (panel B) show some outcomes lying beyond the Pareto frontier, but with confidence intervals overlapping those of the frontier, and thus we cannot reject that the ordering is due to sampling error. However, it is possible to obtain such an ordering within our framework because theory only suggests that the real-world mechanism should lie within the frontier at the *true* distributions F_B and F_S, but we only have bounds on the latter. We evaluate the frontier at the upper and lower bounds while holding fixed the conditional probabilities of trade in that data that are used to estimate the g_κ functions. We estimate F_B and bounds on F_S (and these together determine the second-best frontier) separately from our estimation of the real-world mechanism, without enforcing that the latter lie within the frontier. The point estimates in Figure C.3.B suggest that enforcing such a constraint could potentially tighten the lower bound on F_S. Imposing such a constraint in estimating F_S bounds would not be straightforward, however, as the second-best frontier is itself a solution to an optimization problem (described in Myerson and Satterthwaite 1983 and Williams 1987) that depends on F_B as well as F_S.

[82] Myerson and Satterthwaite (1983) characterized the mechanism on the frontier that is closest to the first-best (ex post efficient) outcome line. Williams (1987) characterized the full Pareto frontier. Our numerical computation of these mechanisms follows the construction in these papers.

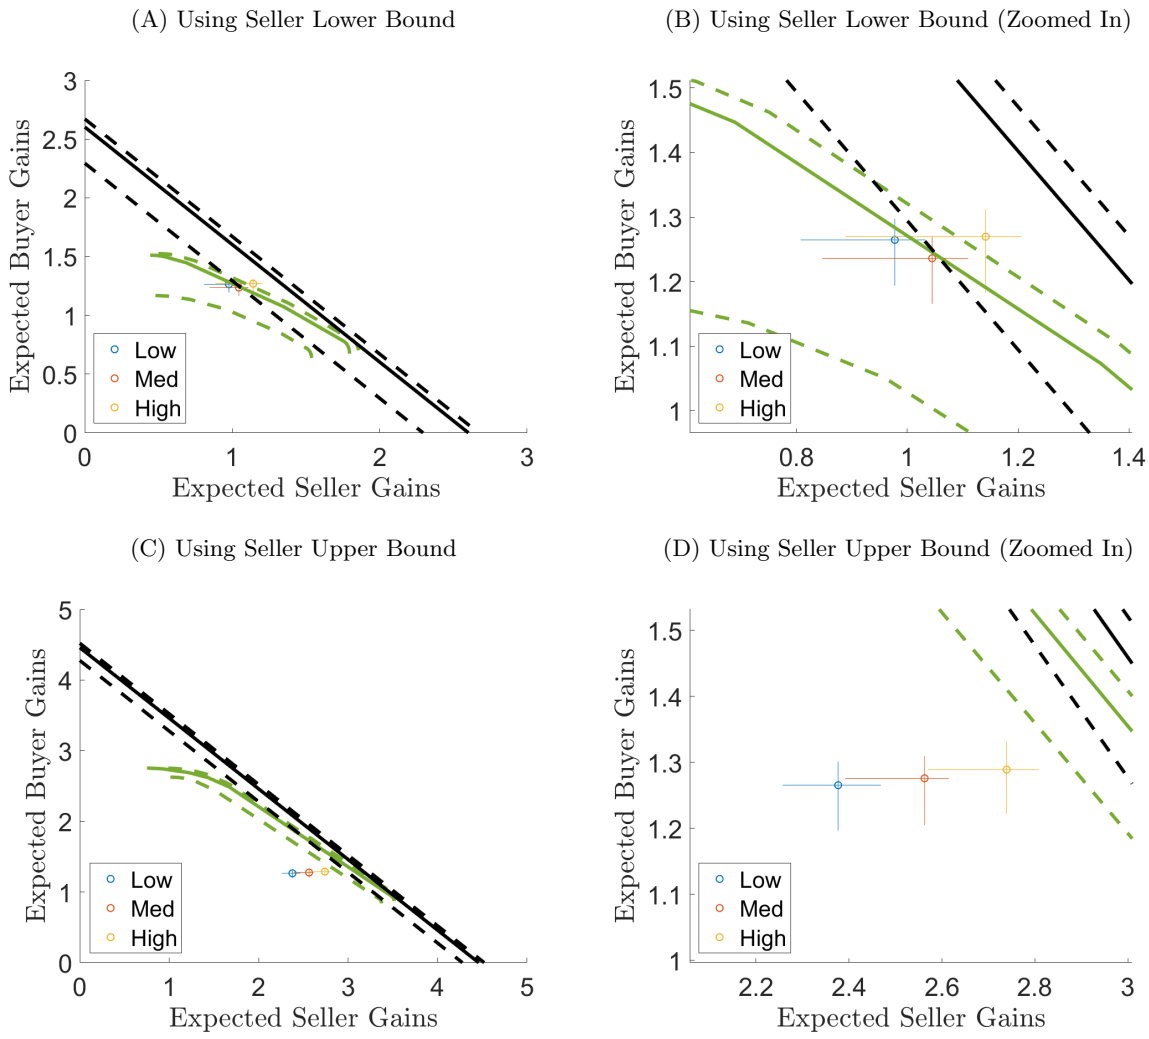
The extreme left and right points on the frontier correspond respectively to the buyer-optimal mechanism (which can be implemented by a take-it-or-leave-it offer by the buyer) and seller-optimal mechanism (a take-it-or-leave-it offer by the seller); see Williams (1987). Under the lower bound on F_S, the real-world outcomes lie relatively close to the middle of the frontier, whereas they lie relatively close to the seller-optimal mechanism under the F_S upper bound. This result is at least partially explained by the fact that the theoretical buyer-optimal mechanism could involve the buyer making an offer that lies well below the auction price, whereas the real-world mechanism enforces that post-auction bargaining offers lie above the auction price, as explained in Section 5.2.

D Analyzing Fleet/lease and Dealer Sales Separately

This section analyzes dealer vs. fleet/lease sellers separately. This analysis is motivated by several comments we gathered in conversations with industry participants. First, as described in Section 4.1, we are told that fleet/lease sellers are more likely than dealer sellers to be assigned their default mediator in negotiations. Second, in at least one auction house location, mediators handling the two different types of sellers do so in different rooms, with mediators handling fleet/lease sellers being instructed to be less aggressive. Third, as described in Appendix C, auction sales for fleet/lease cars often have a binding auction price, meaning that the highest bidder is not allowed to walk away from negotiations before the seller has the opportunity to decide whether to accept or reject the auction price. Finally, fleet/lease sellers generally have more cars to sell on a given day and hence may have less information about any given car. Any of these reasons could lead to different implications for mediator heterogeneity for the two different types of sellers.

The sample in the body uses mediators who handled at least 50 total negotiations, whether these correspond to dealers or fleet/lease sellers. Thus, some mediators in the main sample may handle few dealer cars or few fleet/lease cars. Here we exploit a sample (the *dealers sample*) involving only dealer sellers and limiting to mediators who handled at least 25 such dealer sales. There may be mediators in our main dataset who handled 50 total sales but only 10 dealer sales, and thus these mediators would not be included in the dealers sample. There may also be mediators who handled 30 dealer sales but less than 50 sales in total, and thus the dealers sample is not a strict subset of the main sample. The same construction and arguments apply to the *fleet/lease sample* we use here. We impose this 25-negotiations criterion to create samples with a sufficient number of observations per mediator. Tables D.1–D.4 and Figures D.1–D.9 are replications of all tables and figures from

Figure C.3: Pareto Frontier (Units = $1,000)



Notes: Each panel displays estimated expected seller and buyer surplus on the ex-post efficient frontier (in black), on ex-ante efficient frontier (in green), and in real-world bargaining under different mediator types (low in blue, medium in red, and high in yellow). Top panels use F_S lower bound and bottom panels use upper bound. Panels on right are zoomed-in versions of the left. Units = $\$1,000$.

the body of the paper using these dealers and fleet/lease samples separately.

The results are largely qualitatively similar in the two samples. In both samples, and in the main sample, we document significant heterogeneity in mediator fixed effects that persists after including a rich set of controls and that is not solely due to sampling error. We also find that better mediators (those who achieve a higher trade probability) achieve more efficient outcomes. Some results differ between the two samples. We now discuss similarities and differences in more detail.

Table D.1 shows that dealer sales involve less agreement, lower prices, older (and higher-mileage) cars, and slightly more bargaining offers per thread. Table D.2 shows that mediators of the two different types of cars handle a similar number of threads on average. The average mediator in the fleet/lease sample has slightly more years of employment and is more likely to be female. The ratio of the final price to the auction price or reserve price is similar in the two samples. Figures D.1–D.6 show strikingly similar qualitative and quantitative results in the two samples.

In Table D.3, we find that, in both samples, higher-skilled mediators are more likely than lower-skilled mediators to end the negotiation with agreement in the first period of bargaining (column 2), and that the effects of mediator skill on first-round agreement are much larger than those on later agreement (comparing columns 2 and 3). However, the small effects we do see in column 3 are different in sign between the two samples, suggesting that high-skilled mediators are more likely than low-skilled mediators to achieve later round agreement when interacting with dealer sellers but not with fleet/lease sellers. The effects on later periods in the main analysis (Table 3) average over these two effects, finding a null effect of mediator skill on later period agreements. A key takeaway is that the effect sizes are small in column 3 relative to column 2 in both panels, consistent with the main results. Table D.3, columns 4–6 (as well as Figures D.7.A and D.7.C) confirm that mediator experience correlates positively with agreement. Columns 7–8 of Table D.3, which include mediator fixed effects, show a significant and positive effect for fleet/lease cars of within-mediator experience on the overall probability of agreement or first-round agreement. We find an insignificant effect in these same columns for dealer cars. In both samples, column 9 shows that mediator experience is associated with a lower likelihood of ending in agreement at later bargaining periods.

Figure D.7 panels B and D show a larger gap between the agreement probabilities (as a function of the reserve-auction gap) for different mediator types in the fleet/lease sample than in the dealers sample. In both cases, the agreement probability is decreasing as the reserve-auction gap increases, and in both cases the gap between the lines corresponding to high-skilled vs. medium-skilled mediators appears to widen as the reserve-auction gap increases. The gap between the lines corresponding to high- and low-skilled mediators appears to be closer to constant.

Figure D.8 shows the estimated g_κ functions. Among dealer cars (panel A), we find the expected ordering between high- and low-skilled mediators in the region where more mass lies, with high-skilled mediators facilitating more efficient trades; this can be seen by the yellow line lying below the blue line in this region, meaning more buyer types trade under the high-skilled mediator mechanism.

Medium-skilled mediators are not clearly ordered in dealer sales, appearing to overlap substantially with g_L or g_H. Among fleet/lease sales (panel B), the order is instead quite clear across the low, medium, and high g_κ functions, with higher-skilled mediators allowing more buyer types to trade.[83] Finally, the key takeaways from comparing efficiency at different mediator skill levels are similar in the two samples (Figures D.9.A and D.9.C). One distinction between the two samples is that, in the dealers sample (Figures D.9.B), we find high-skilled mediators increase seller surplus, with no statistically significant effect on buyers, and we find the opposite in the fleet/lease sample (Figures D.9.D). The effect on total surplus of high-skilled vs. low-skilled mediation is positive in both samples. The implications for fees (Table D.4) are also similar in the two samples.

[83]Note, as mentioned in the body of the paper, this figure only shows cross-sections (evaluated at P^A=-\$2,000) of the two-dimensional surfaces corresponding to the g_κ functions. Across the full surface, the g_M function implies more trade than the g_L function even in the dealers sample, just not at the cross-section in Figure D.8.A.

Table D.1: Dealer vs. Fleet/lease Sales: Descriptive Statistics at Bargaining-Thread Level

Panel A: Dealer Sales

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.475	0.499	0	1
Final price ($)	5,578	5,209	800	13,300
Book price ($)	6,795	5,318	1,500	14,600
Auction price ($)	5,551	5,114	800	13,100
Reserve Price ($)	7,212	5,574	1,800	15,400
No reserve	0.193	0.394	0	1
# Offers in a thread	1.602	0.803	1	3
Length of a thread (hours)	6.358	16.374	0.298	21.3
Car age (years)	7.036	3.705	2	12
Mileage	97,602	51,253	31,820	162,679
Engine displacement (liters)	3.619	1.490	2	5.7
No. Threads	39,563			

Panel B: Fleet/lease Sales

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.789	0.407	0	1
Final price ($)	5,618	4,741	900	12,200
Book price ($)	7,133	5,189	1,650	14,350
Auction price ($)	5,591	4,728	850	12,200
Reserve Price ($)	7,521	5,166	2,000	14,700
No reserve	0.323	0.467	0	1
# Offers in a thread	1.159	0.476	1	2
Length of a thread (hours)	5.066	13.005	.376	6.68
Car age (years)	5.602	3.329	2	10
Mileage	88,707	50,437	28,718	151,612
Engine displacement (liters)	3.575	1.563	2	5.7
No. Threads	35,369			

Notes: Statistics at the thread level in the dealers (panel A) vs. fleet/lease (panel B) samples, as in Table 1.

Table D.2: Dealer vs. Fleet/lease Sales: Descriptive Statistics at Mediator Level

Panel A: Dealer Sales

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.484	0.146	0.325	0.639
Final price/book price	0.854	0.0872	0.76	0.951
Final price/reserve price	0.804	0.0584	0.739	0.869
Final price/auction price	1.031	0.0441	1	1.05
Female	0.357	0.482	0	1
# Threads mediated	439	524	30	1,090
Years of employment	3.933	4.469	0.375	10.5
No. Mediators	90			

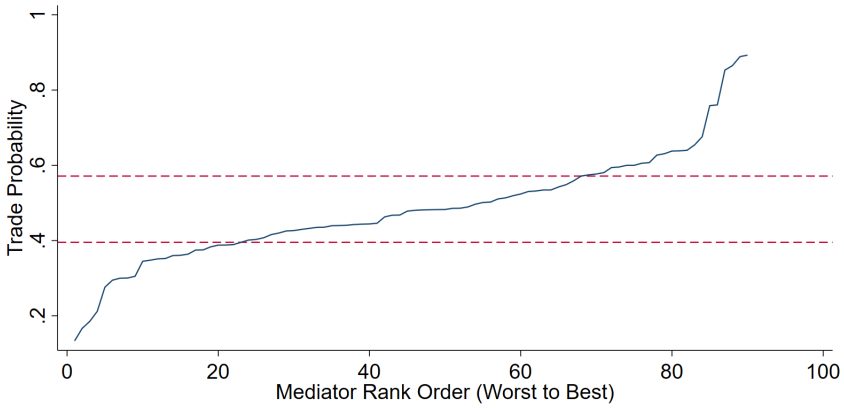
Panel B: Fleet/lease Sales

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	0.705	0.199	0.444	0.929
Final price/book price	0.813	0.0885	0.712	0.916
Final price/reserve price	0.799	0.077	0.711	0.898
Final price/auction price	1.009	0.016	1	1.02
Female	0.525	0.502	0	1
# Threads mediated	431	654	33	1,091
Years of employment	4.753	5.625	0.41	10.6
No. Mediators	82			

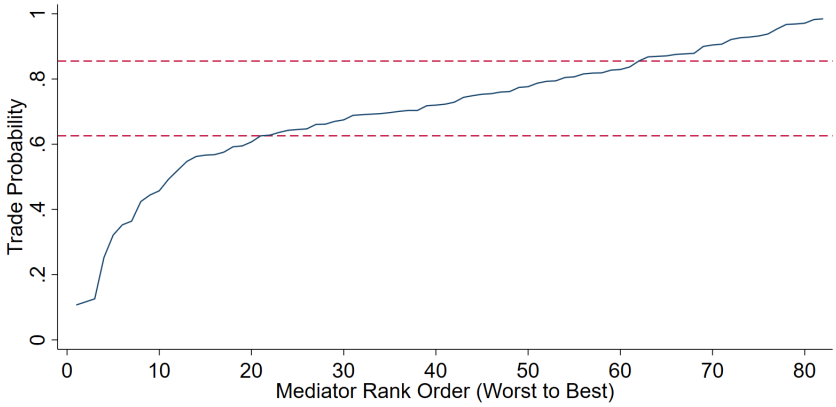
Notes: Statistics at the thread level in the dealers (panel A) vs. fleet/lease (panel B) samples, as in Table 2.

Figure D.1: Dealer vs. Fleet/lease Sales: Mediator Differences in Trade Probability

(A) Dealers

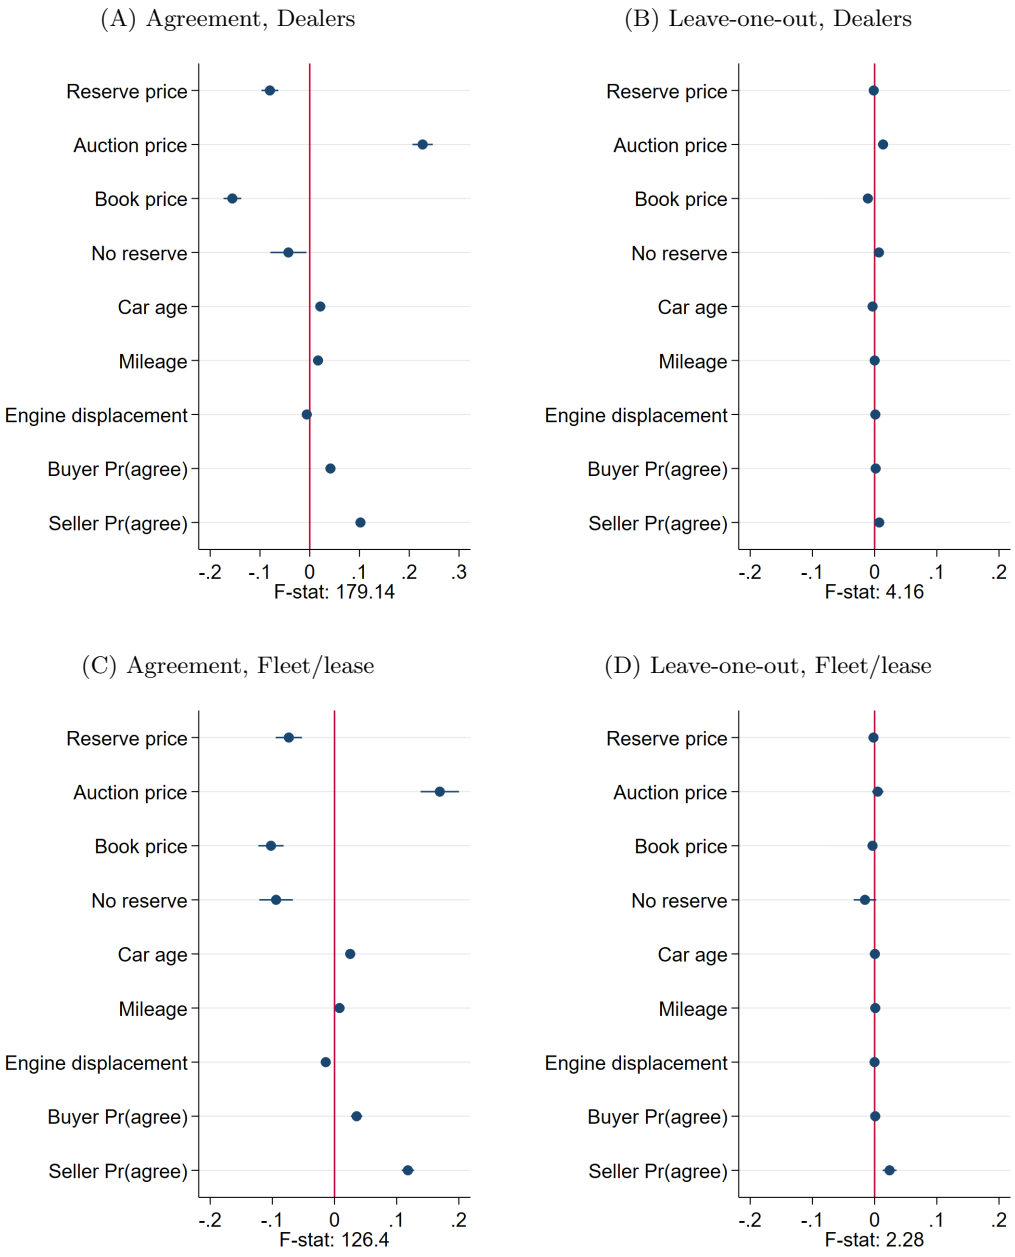


(B) Fleet/lease



Notes: Replication of Figure 1 on the dealers and fleet/lease samples separately.

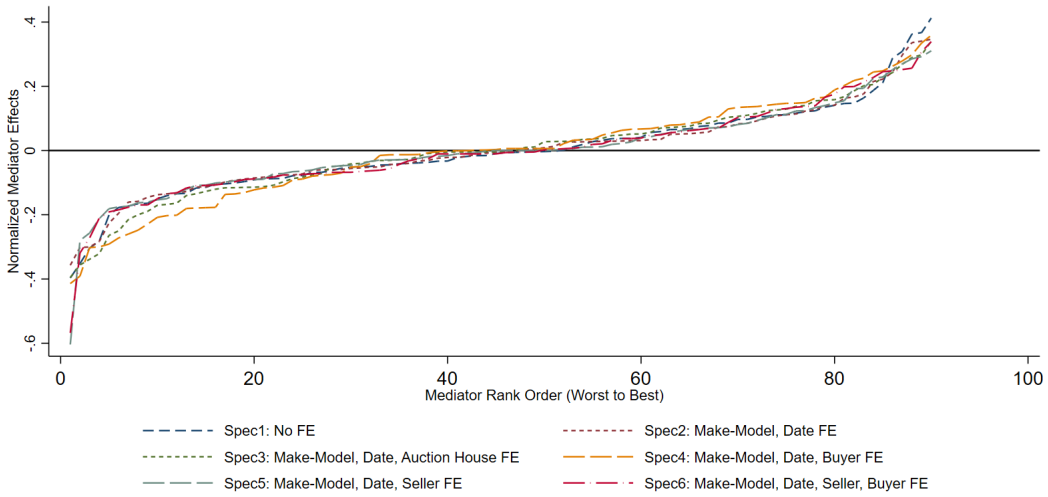
Figure D.2: Dealer vs. Fleet/lease Sales: Assignment Test



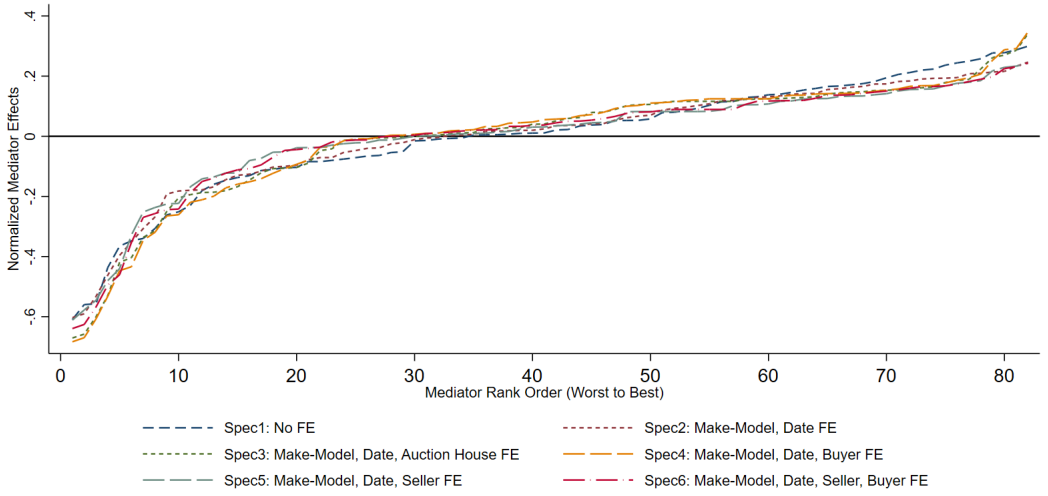
Notes: Coefficients of (2), as in Figure 2 but estimated separately on the dealers sample (panels A and B) and fleet/lease sample (panels C and D).

Figure D.3: Dealer vs. Fleet/lease Sales: Mediator Fixed Effects for Trade Probability Under Different Specifications

(A) Dealers



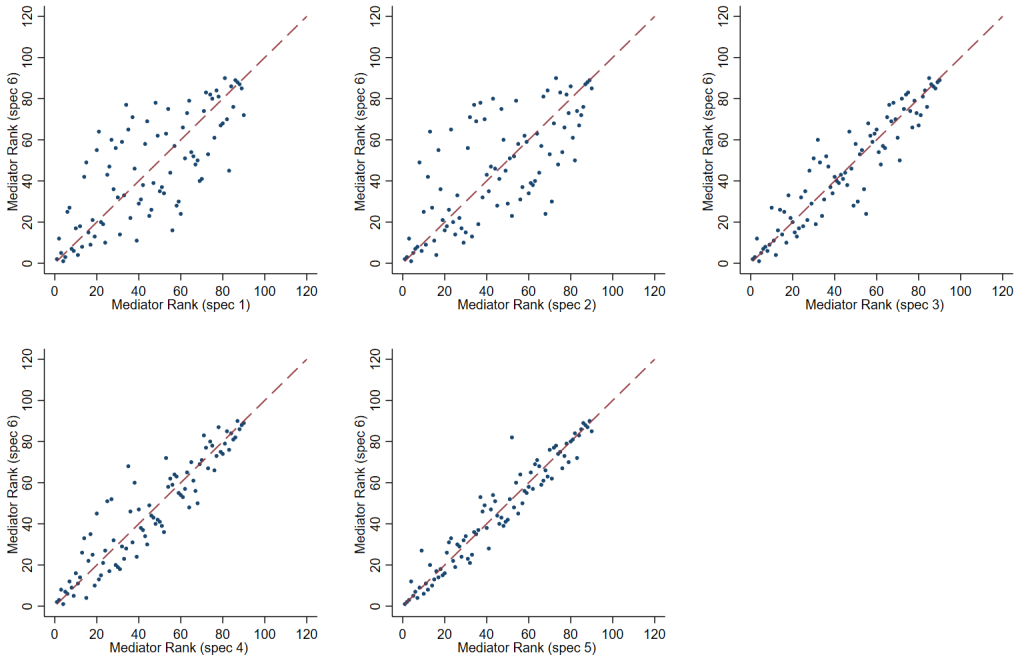
(B) Fleet/lease



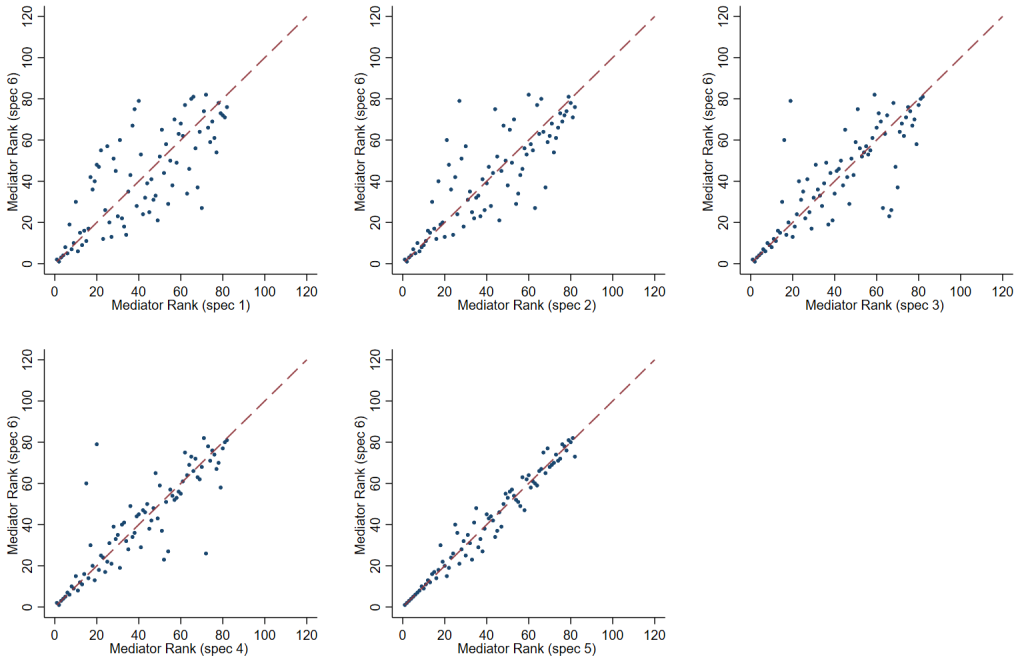
Notes: Replication of Figure 3 separately on the dealers and fleet/lease samples.

Figure D.4: Dealer vs. Fleet/lease Sales: Mediator Rank Across Specifications

(A) Dealers



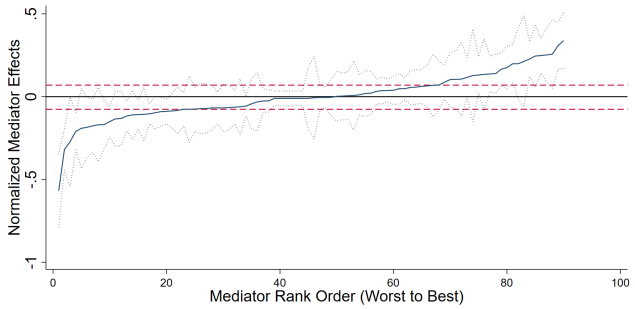
(B) Fleet/lease



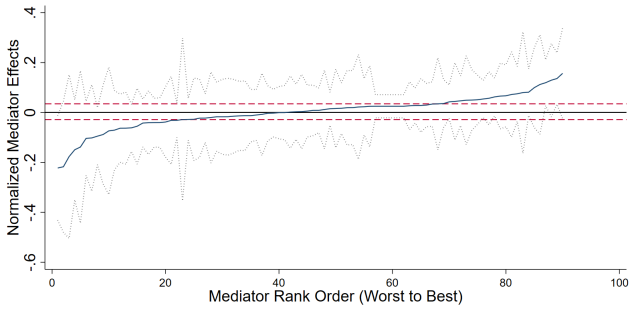
Notes: Replication of Figure 4 separately for the dealers and fleet/lease samples.

Figure D.5: Mediator Fixed Effects Under Most-Saturated Specification

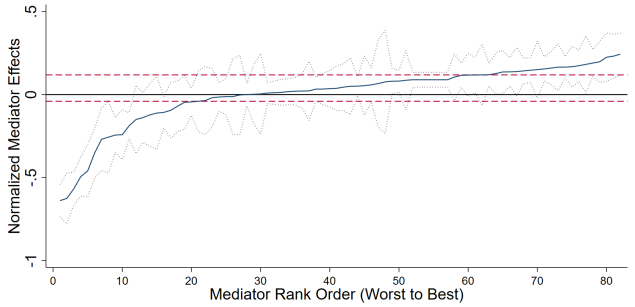
(A) Trade Probability, Dealers



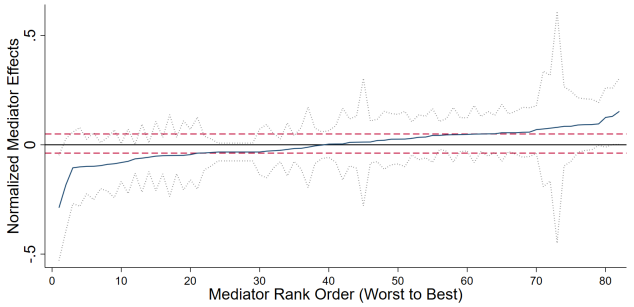
(B) Normalized Final Price, Dealers



(C) Trade Probability, Fleet/lease

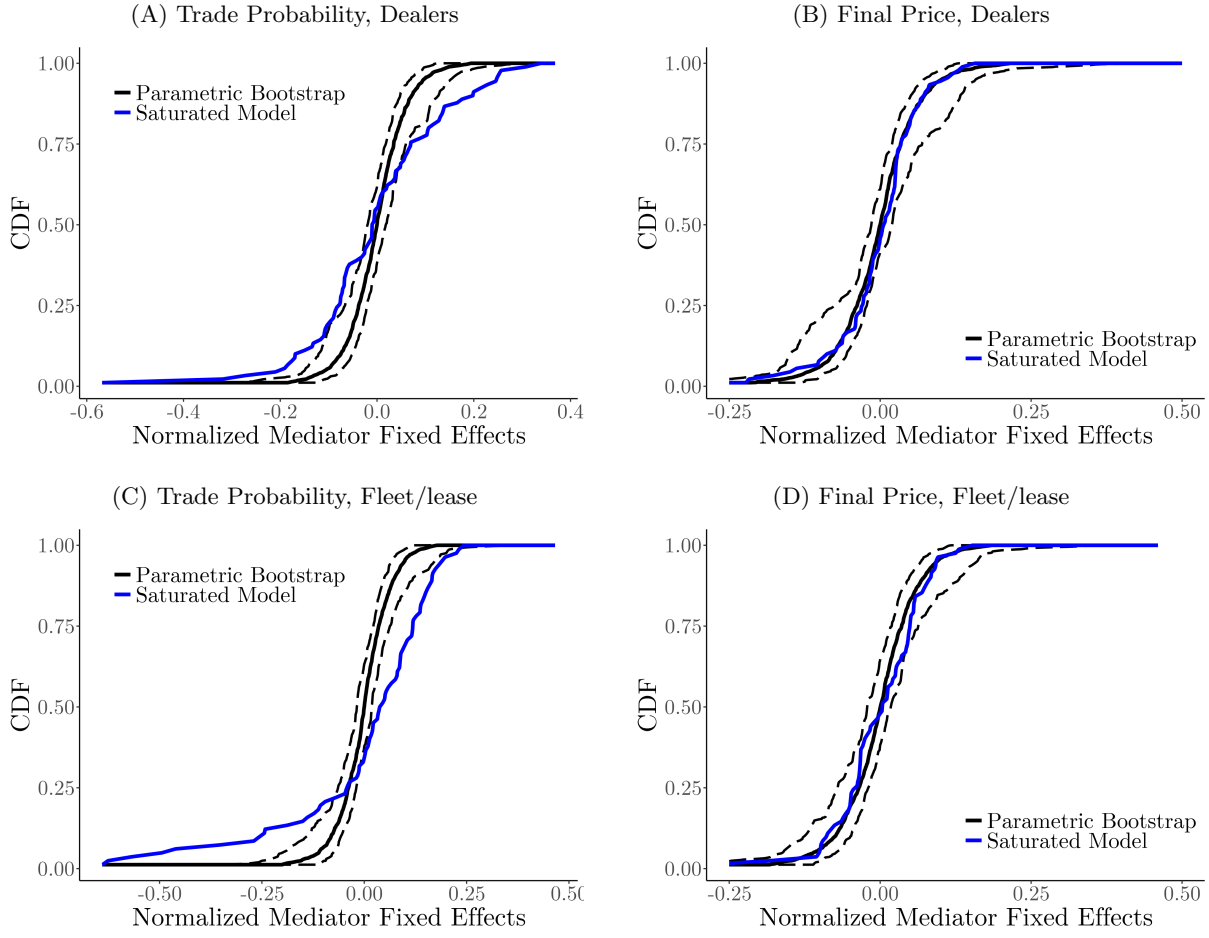


(D) Normalized Final Price, Fleet/lease



Notes: Replication of Figure 5 separately on the dealers and fleet/lease samples.

Figure D.6: Dealer vs. Fleet/lease Sales: Bootstrap Test of Heterogeneity from Sampling Error



Notes: Parametric bootstrap test results, as in Figure 6, but using the dealers sample (panels A and B) and fleet/lease sample (panels C and D).

Table D.3: Agreement Rates in First Later Rounds and Mediator Experience

(a) Dealers

	(1) $Agree_i$	(2) $Agree_i^1$	(3) $Agree_i^{>1}$	(4) $Agree_i$	(5) $Agree_i^1$	(6) $Agree_i^{>1}$	(7) $Agree_i$	(8) $Agree_i^1$	(9) $Agree_i^{>1}$
High	0.210***	0.181***	0.028***						
	(0.012)	(0.011)	(0.009)						
Medium	0.103***	0.085***	0.018***						
	(0.008)	(0.007)	(0.006)						
Experience (years)				0.010***	0.008***	0.002*	-0.006	0.004	-0.011**
				(0.001)	(0.001)	(0.001)	(0.006)	(0.006)	(0.005)
Mediator FE							Y	Y	Y
R2	0.333	0.348	0.255	0.363	0.385	0.293	0.373	0.394	0.298
N	36,481	36,481	36,481	25,392	25,392	25,392	25,392	25,392	25,392

(b) Fleet/lease

	(1) $Agree_i$	(2) $Agree_i^1$	(3) $Agree_i^{>1}$	(4) $Agree_i$	(5) $Agree_i^1$	(6) $Agree_i^{>1}$	(7) $Agree_i$	(8) $Agree_i^1$	(9) $Agree_i^{>1}$
High	0.307***	0.335***	-0.028***						
	(0.010)	(0.011)	(0.005)						
Medium	0.197***	0.233***	-0.035***						
	(0.010)	(0.011)	(0.005)						
Experience (years)				0.002**	0.001	0.001*	0.023***	0.030***	-0.007***
				(0.001)	(0.001)	(0.001)	(0.004)	(0.005)	(0.002)
Mediator FE							Y	Y	Y
R2	0.350	0.383	0.202	0.361	0.392	0.213	0.397	0.429	0.219
N	33,066	33,066	33,066	24,730	24,730	24,730	24,730	24,730	24,730

Notes: Estimates as in Table 3, but separately using dealers and fleet/lease samples.

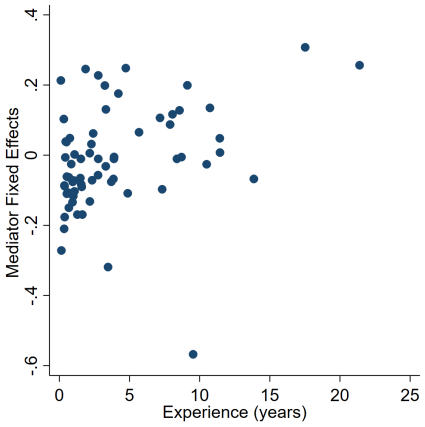
Table D.4: Auction House Revenue Under Different Mediator Types (Units= $)

A. Dealers	Low	Medium	High
Revenue	117.10	136.26	148.65
	(2.17)	(1.66)	(2.18)
	High - Medium	Medium - Low	High - Low
Revenue Difference	12.39	19.16	31.55
	(2.81)	(2.56)	(3.15)
B. Fleet/lease	Low	Medium	High
Revenue	151.47	210.34	246.91
	(4.71)	(2.04)	(1.21)
	High - Medium	Medium - Low	High - Low
Revenue Difference	36.57	58.87	95.44
	(2.38)	(5.32)	(4.86)

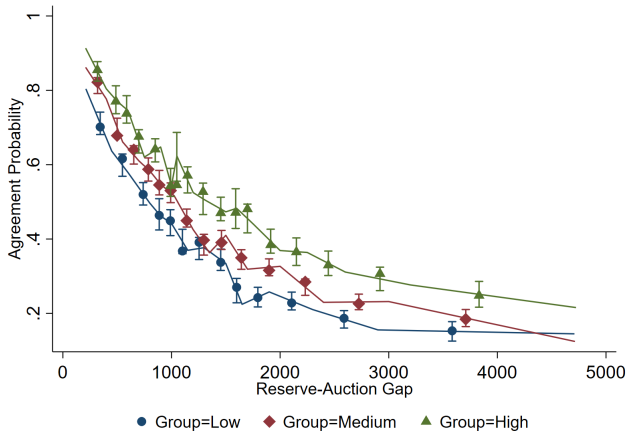
Notes: Results as in Table 4 but estimated separately using the dealers vs. fleet/lease samples. Units= $

Figure D.7: Agreement vs. Mediator Experience and Pr(Agreement) vs. Reserve-Auction Gap

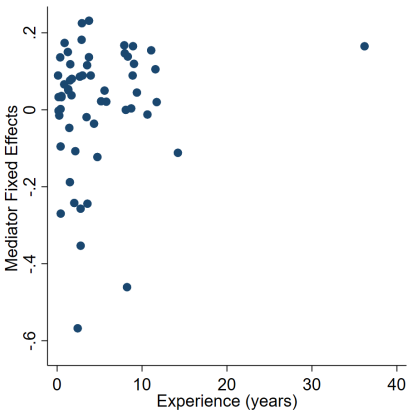
(A) Dealers, experience



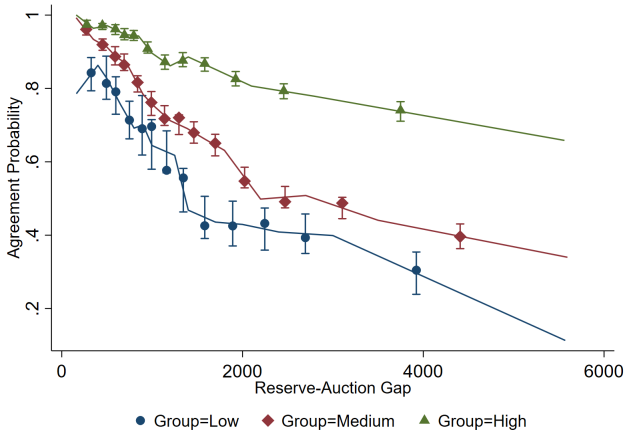
(B) Dealers, reserve-auction gap



(C) Fleet/lease, reserve-auction gap

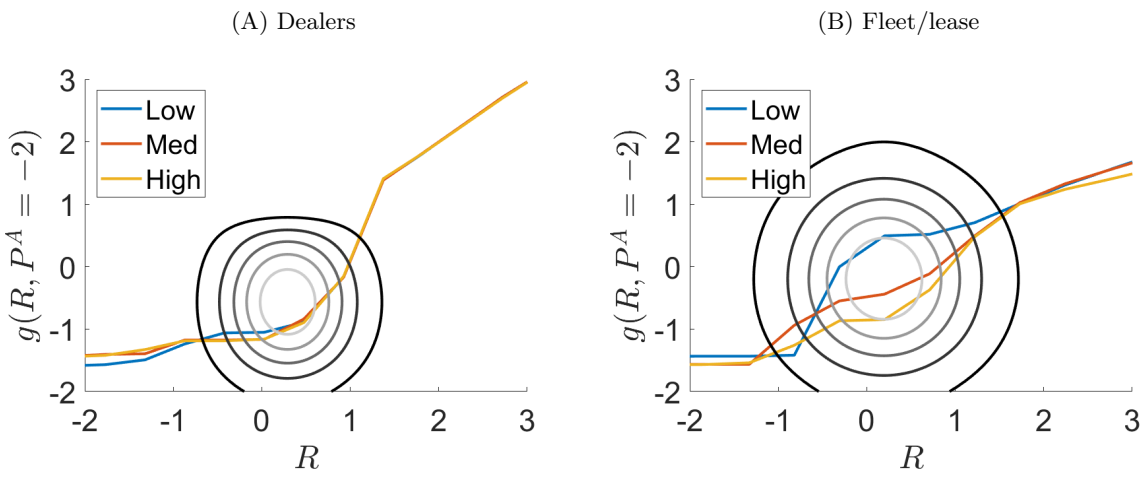


(D) Fleet/lease, experience



Notes: Estimates as in Figure 7, but separately using the dealers vs. fleet/lease samples.

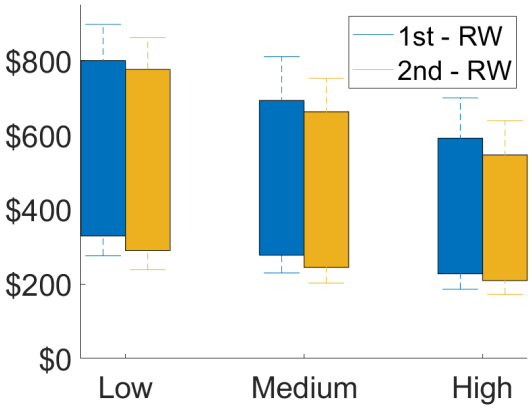
Figure D.8: Dealer vs. Fleet/lease Sales: Direct Mechanisms $g_\kappa(\cdot)$ Functions for Different Mediator Types (Units = \$1,000)



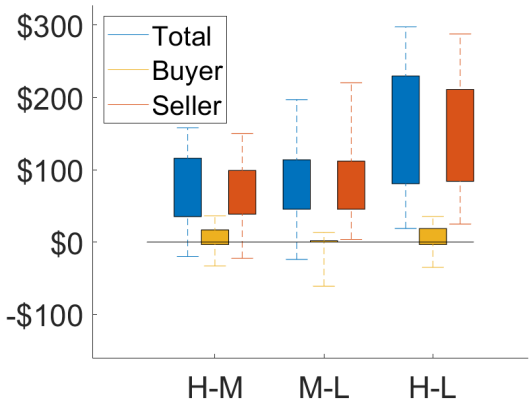
Notes: Estimates of $g_\kappa(R, P^A)$, as in Figure 8, but separately using the dealers vs. fleet/lease samples. Units = \$1,000.

Figure D.9: Dealer vs. Fleet/lease Sales: Surplus Differences Across Mediators

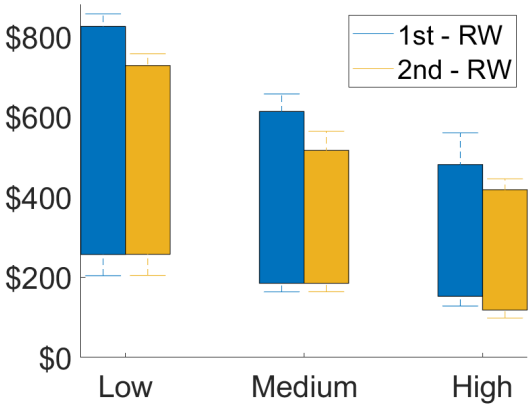
(A) RW vs. 1st/2nd Best, Dealers



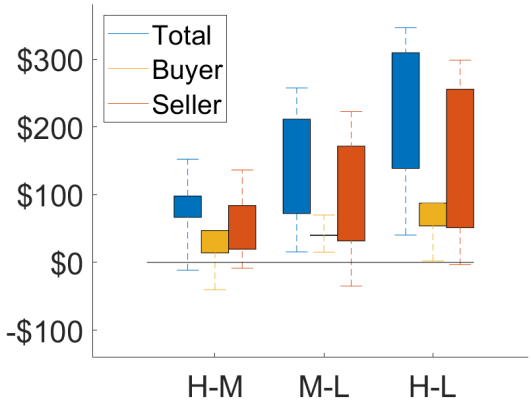
(B) Diff. Across Mediators, Dealers



(C) RW vs. 1st/2nd Best, Fleet



(D) RW vs. 1st/2nd Best, Fleet



Notes: Estimates as in Figure 9 but estimated separately using the dealers vs. fleet/lease samples.