NBER WORKING PAPER SERIES

INTERMEDIARIES IN BARGAINING:
EVIDENCE FROM BUSINESS-TO-BUSINESS USED-CAR INVENTORY NEGOTIATIONS

Bradley Larsen
Carol Hengheng Lu
Anthony Lee Zhang

Working Paper 29159
<http://www.nber.org/papers/w29159>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2021

An earlier version of this study was circulated under the title, “Mediation in Bargaining: Evidence from Large-Scale Field Data on Business-to-Business Negotiations.” We thank Rebecca Li for outstanding research assistance, as well as Jack Fanning, Matt Grennan, Nagore Iriberri, Pranav Jindal, Bruno Pellegrino, Joris Pinske, Renan Yoshida, Xingtang Zhang, and seminar and conference participants at the 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, IIOC 2021, and BEET 2021 (Bargaining: Experiments, Empirics, and Theory). We acknowledge support from the National Science Foundation through grants SES-1629060 and SES-1530632. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Bradley Larsen, Carol Hengheng Lu, and Anthony Lee Zhang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Intermediaries in Bargaining: Evidence from Business-to-Business Used-Car Inventory Negotiations
Bradley Larsen, Carol Hengheng Lu, and Anthony Lee Zhang
NBER Working Paper No. 29159
August 2021
JEL No. C7,D8,L1,L81

ABSTRACT

We analyze data on tens of thousands of alternating-offer, business-to-business negotiations in the wholesale used-car market, with each negotiation mediated (over the phone) by a third-party company. The data shows the identity of the employee mediating the negotiations. We find that who intermediates the negotiation matters: high-performing mediators are 22.03% more likely to close a deal than low performers. Effective mediators improve bargaining outcomes by helping buyers and sellers come to agreements faster, not by pushing disagreeing parties to persist, and have real effects on efficiency for some negotiations, overcoming some of the inefficiency inherent in incomplete-information settings.

Bradley Larsen
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305
and NBER
bjlarsen@stanford.edu

Anthony Lee Zhang
University of Chicago
Booth School of Business
5807 S Woodlawn Ave
Chicago, IL 60637
anthony.zhang@chicagobooth.edu

Carol Hengheng Lu
Department of Economics
Stanford University
579 Jane Stanford Way
Stanford, CA 94305
United States
hhlu@stanford.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 the role they play; consider, for example, investment banks handling firm acquisitions or lawyers mediating pre-trial settlement. To date, however, there is little quantitative evidence from real-world data on whether or how such mediators make a difference for negotiation outcomes and welfare in practice.

This paper provides an analysis of a large dataset containing 80,000 business-to-business negotiations from the wholesale used-car industry. This industry consists of hundreds of auction houses nationwide that facilitate trade of used cars between manufacturers, fleet companies, banks, leasing companies, and used and new car dealerships. More than \$80 billion worth of cars are traded through this industry each year. Each car is auctioned individually in a rapid process, and when the auction fails to yield a sufficiently high price, the auction house mediates a bilateral bargaining process between the highest bidder and the seller of the car.

The dataset we study is rich, containing information from six different auction houses that sold thousands of cars from 2006–2010. For each attempt to sell a car, the data records the auction price, the seller’s secret reserve price, and every action taken by each party in the negotiation process, including all back-and-forth offers. Importantly for our study, the data contains the identity of the auction house employee who mediated the negotiation over the phone. The data also contains detailed information on the characteristics of the vehicle, the timing and location of the transaction, and the identities of the negotiators (the buyer and seller on a given transaction). Such data is rare in the literature—only a handful of existing studies analyze information on offers and counteroffers within a real-world bargaining scenario, and we know of no other data setting containing information on the mediation as well. We view this as an unprecedented opportunity to study bargaining intermediaries in the field.

The buyers and sellers negotiating in these transactions can be considered “professional”

negotiators, and the auction house mediators “professional” mediators. They engage in these negotiations—with different parties—on a weekly basis, as each auction house sells hundreds to thousands of cars on a fixed day each week. For the buyers and sellers, the stakes are high, especially for small used-car dealers, where each transaction can make the difference between having the right amount of inventory on hand vs. not, and having the desired resale profit margin vs. not.

This data provides us with a setting in which mediators vary from transaction to transaction, allowing us to quantify the impact of mediators on negotiation outcomes. According to our conversations with industry participants and with the mediators themselves, these mediators (or *intermediaries*—we will use the two terms interchangeably throughout the paper) are largely random. We find evidence that the assignment is not actually random, and can depend on the features of a given negotiation that are determined before the 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 allows us to control for each of these features, and we demonstrate that our key findings are similar with or without these controls.

We first document that *who* is assigned as the mediator has a large effect on economic outcomes. We measure this by a regression of various outcomes on mediator fixed effects. We find that the 75th percentile-performing mediator leads to trade probability that is 22.03% higher than the 25th percentile mediator. We find that this large heterogeneity in performance persists as we include more rigorous controls for features of the negotiation. Some variation in mediator performance would be expected in any dataset simply due to statistical noise, even if true underlying performance was equal across mediators. To take this into account, we perform a number of placebo tests in which we randomly permute the assignment between bargaining outcomes to mediators in the data. These tests reveal that the observed performance heterogeneity we document is wider than can be explained by random statistical error, suggesting that mediator skill is a real phenomenon.

We then use this heterogeneity to study several questions that relate both to the negotiation literature in social psychology and organizational behavior, and to foundational

mechanism design concepts in economics. The classical theoretical mechanism design literature on bilateral trade (Myerson and Satterthwaite 1983) demonstrates that bargaining is generically inefficient: incomplete information leads to information rents in order to achieve incentive compatibility, and hence also leads to inefficiency. This inefficiency exists even when the mechanism designer is able to implement the theoretically most-efficient mechanism. In practice, however, the actual shortfall in efficiency may be far greater than would be explained by these traditional information constraints alone, suggesting that market designers in practice may either have different objective functions than posited by theory, or may simply be *incapable* of implementing the efficient mechanism.

In our setting, the objective function of the mechanism designer — the used-car clearinghouse — is well known: to maximize the probability of trade. This objective function is stated clearly in industry reports and industry blogs, and has been confirmed in personal interviews of many industry participants (Lacetera et al. 2016, Treece 2013).¹ We posit here that any efficiency shortfall in bilateral bargaining in this setting is therefore not due to a different designer objective function but rather to an *inability* of the designer to execute the theoretically preferred mechanism.

To study this inability, we exploit the variation in mediator performance. As in the theoretical discussion in Myerson (2008), we think of different mediators as constituting different *mechanisms*, achieving outcomes that differ in performance relative to the theoretically preferred mechanism. The goal of each mediator/mechanism is to maximize trade, but some mediators may be *unable to commit to* or *lack an understanding of how to execute* the optimal strategy, and may instead attempt to push bargaining agents to keep negotiating, even when trade is unlikely.

The theoretically optimal mechanism for maximizing trade can be implemented as a static mechanism—it does not require long, drawn-out negotiations between parties. A key feature of this mechanism is that the designer commits to *not* let certain trades occur, even if the buyer values the car more than the seller. This is also true of the surplus-

---

¹Our findings are consistent with this objective: We find that they have small and statistically insignificant effects on the prices at which trades occur. Thus, mediators affect trade probabilities without substantially changing how the pie is split between buyers and sellers.

maximizing mechanism studied in Myerson and Satterthwaite (1983) and Williams (1987). This commitment is necessary in order to achieve incentive compatibility: without some failed, efficient trades, the Myerson-Satterthwaite Theorem demonstrates that the mechanism cannot be incentive compatible and cannot be profitably operated by the mechanism designer. A mediator who cannot commit to the optimal mechanism will be tempted to do the following: after learning or believing that there are gains from trade, even if the optimal mechanism requires that trade fail for a given pair of buyer and seller valuations, the noncommittal mediator is tempted to *try again*, to relax the mechanism constraints and get this one trade to occur. Parties knowing this non-commitment, on the other hand, will exploit this knowledge and not truthfully report their valuations to the mediator upfront, resulting in a self-fulfilling prophecy that the game drags on into later periods. This dilemma faced by mediators is analogous to the Coase Conjecture—the dilemma faced by a durable good monopolist who cannot commit to *not* lower prices in the future.²

We construct a simple framework for decomposing bargaining success into two distinct components. We show that, in an alternating-offer bargaining game, the probability of bargaining success can be decomposed into two sets of probabilities for each bargaining round: the probability that participants agree and conclude bargaining, which we call an *agreement probability*; and the probability that participants disagree, but continue to the next round, which we call *continuation probability*. We find empirically that effective mediators improve bargaining outcomes by increasing agreement probabilities, not continuation probabilities. In fact, continuation probabilities appear to be somewhat lower for effective mediators.

We then offer a structural exercise that estimates the actual trade mechanism corresponding to mediators of different abilities. Following the Revelation Principle, a bilateral trade mechanism can be summarized by a probability of trade assigned to each pair of agent types (a buyer valuation and a seller valuation)—the allocation function in a direct revelation mechanism. We estimate this direct revelation mechanism separately for mediators of differing skills. We combine these estimated mechanisms with estimates of the distributions

---

²The Coase Conjecture in fact has a long and intimate relationship with the bilateral negotiation literature. See Gul et al. (1986), as well as Ausubel and Deneckere (1992) and Cramton (1992).

of buyer and seller valuations in this marketplace from Larsen (2021). With these objects, we then compute surplus in the trade-maximizing (efficient) mechanism as well as the actual surplus achieved by mediators of different skill levels.³ We find that mediators who achieve a higher probability of trade are not doing so simply by capturing 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*—thus creating significant value. This is particularly the case for car transactions in which the seller is a large fleet or lease institution.

The effects we document in this paper are much larger than effects of mediators in trade in non-bargaining context, such as auctions (Lacetera et al. 2016). We believe this is because bargaining games are complex, theoretically less determinate, and therefore richer mechanisms than auctions. In negotiation settings, game-theoretic concerns, such as dominant-strategy solvability, as in the second-price or VCG auctions, or even Nash equilibrium, have little bite or predictive power. This, in theory, suggests that there is a large space for soft factors, such as mediators’ ability, to influence outcomes; one could think of mediators as playing a kind of equilibrium selection role, with different mediators corresponding to less or more efficient equilibria. We view a contribution of this paper as empirically demonstrating a situation in which the size of these mediators’ effects on outcomes are precisely measurable and quantitatively large.

2 Related Literature

There is a growing theoretical literature in economics and political science on the influence of bargaining mediators, such as Goltsman et al. (2009), Hörner et al. (2015), and Fanning (2021). These studies suggest possible ways in which mediators can and cannot affect outcomes. For example, Fanning (2021) demonstrates that mediators can improve bargaining

---

³Note that the first-best mechanism maximizes the gains from trade as well as the probability of trade. With incentive constraints, the second-best mechanism that maximizes gains from trade and the second-best mechanism that maximizes the volume of trade are not exactly equivalent. In our paper, we evaluate the first-best mechanism as well as the surplus-maximizing second-best.

outcomes by committing to, with some probability, withhold information from each negotiating agent about her opponent’s willingness to trade. Basak (2015) finds that if agents have sufficiently close bargaining strengths, mediation can strictly improve efficiency. Kydd (2003) provides a model in which a bargaining party will only follow the advice of a mediator who she views as sufficiently “on her side.” Other theoretical studies suggesting that mediation in bargaining can improve efficiency include Copic and Ponsati (2008), Glode and Opp (2016), and Kim (2017). More broadly, a mediator might be viewed as a mechanism for helping agents execute something closer to the efficient direct mechanism (e.g. Myerson and Satterthwaite 1983) by withholding some trades between players—even when the mediator knows the buyer values the good more than the seller—in order to keep agents’ reporting incentive compatible. Zhang (2019) demonstrates theoretically that more informed mediators can indeed achieve greater efficiency.⁴

A number of studies in at the intersection of economics and organizational behavior/psychology gain insights on the impact of mediation via controlled laboratory experiments, finding evidence that a third-party mediator decreases the probability of trade and increases the price, depending on the mediator’s incentives and what the mediator knows (Bazerman et al. 1992, Valley et al. 1992). Experiments, and also studies of international conflict, have shown that bargaining parties tend to prefer mediation (over un-mediated bargaining) in cases with uncertainty about negotiation outcomes or unequal bargaining power (Neale 1984, Bercovitch and Jackson 2001). A mediator might also be able to help by convincing one side or the other that the market demand or supply is different than that agent initially expected. For example, if the agent has overly optimistic beliefs or other biases, the mediator may be able to help the agent correctly adjust these beliefs (Babcock and Loewenstein 1997). Other experimental work along these lines includes Yavas et al. (2001), and Eisenkopf and Bachtiger (2013). Our work takes the analysis of intermediaries beyond the lab to a business-to-business setting: a

---

⁴Lang (2020) describes a model in which agents’ ability to tell and interpret narratives can improve efficiency in bilateral bargaining, and Kim (2020) shows that an intermediary with the objective of maximizing trade (as opposed to revenue) is more likely to truthfully reveal agents values to one another. In Gottardi and Mezzetti (2019), the mediator is more informed than the negotiating agents about their own values, and the mediator finds it optimal to limit the flow of information to the agents.

high-stakes negotiation in the field between professional negotiators.⁵

Our study relates to other empirical studies that document heterogeneous outcomes in sales situations in the field, such as Barwick and Pathak (2015), Kim (2020), Robles-Garcia (2020), and Gilbukh and Goldsmith-Pinkham (2021), 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 (2020), studying heterogeneity across sellers in large-scale appliances. Relative to this literature, our core contribution is to analyze large, detailed data on the inner workings of the negotiations (most studies only observe the final negotiated price, and only for consummated sales, unlike our data) and to demonstrate that mediator effects are surprisingly quantitatively large.

Finally, our study relates to the empirical literature studying bargaining under two-sided incomplete information (Keniston 2011, Larsen and Zhang 2018, Larsen 2021). The data we use overlaps to some extent with the alternating-offer bargaining data studied in Larsen (2021), although that study does not exploit information on the identity of the mediator. That study provides structural welfare analyses of the performance of the bilateral bargaining but leaves unanswered the question of *why* bargaining is inefficient at all, and what explains variation in outcomes across negotiating pairs. The current paper takes a first step in this direction, demonstrating that a large fraction of the variation in the outcomes can be explained by mediator influence. In laboratory experiments, Valley et al. (2002) and Bochet et al. (2020) study two-sided incomplete-information bargaining, as in our paper, and demonstrate that other factors contributing to efficiency are agents' ability to communicate with one another and their ability to bundle together multiple issues in their price offers.

---

⁵In legal settings, McEwen and Maiman (1981) find that small claims suits resolved through mediation result in higher compliance and satisfaction rates than claims resolved through adjudication, and Emery et al. (1991) finds that mediation improved satisfaction in child custody cases. Dixon (1996) finds that, in the setting of international military disputes, various forms of third-party intermediation are associated with higher settlement rates.

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 many other parts of the world. Millions of used-cars arrive each year to used-car lots as trade-in vehicles and then are never sold on those lots, but are instead brought to a wholesale used-car auction house, where the inventory is sold to other car dealerships. We refer to these cars as cars sold by *dealers*. Millions of company-fleet vehicles, rental cars, repossessed vehicles owned by banks, or off-lease or lease-buy-back vehicles are also offered for sale at these auction houses. We refer to these as cars sold by *fleet/lease* sellers. Total revenue at these auction houses is more than \$80 billion annually.

At the auction houses we study, for each car brought to an auction house, the auction house runs a rapid (approximately 90-second) ascending auction, and if the auction price fails to reach the seller’s secret reserve price, the auction house facilitates a bilateral negotiation between the high bidder and the seller. This negotiation proceeds over the phone. The auction house employee—the mediator—first calls up the seller and reports the outcome of the auction (the auction price). The seller can choose to accept this price, give a counteroffer, or quit (ending the negotiation). If the seller gives a counteroffer, the mediator calls up the buyer and the buyer is given the same choice. This process continues until one party accepts or quits. The auction house records all actions taken by either party, as well as the identity of each party and the identity of the mediator.

The data consists of several hundred thousand realizations of bargaining sequences at six different auction house locations between 2006 and 2010. We take several steps to clean the data. We first drop observations for which the following variables lie outside their respective 0.01 and 0.99 percentiles: the auction price, the reserve price, and the auction house’s estimated market value of the car (which we will refer to as the *blue book price*). We will use the term *thread* to refer to a bargaining sequence. The auction house creates a new record for each action taken during a given bargaining thread, allowing us to see the timing of each action and also allowing us to see that some threads involve several different mediators facilitating different stages of the negotiation. Among all the bargaining threads, 68.76% of

Table 1: Descriptive Statistics at Bargaining-Thread Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	.589	.492	0	1
Final price (\$)	5,659	4,948	900	12,800
Bluebook price (\$)	6,969	5,261	1,550	14,475
Auction price (\$)	5,575	4,932	800	12,700
Reserve Price (\$)	7,361	5,412	1,900	15,000
# Offers in a thread	1.37	.687	1	2
Length of a thread (hours)	5.92	15.3	.33	19.8
Fleet/lease car	.473	.499	0	1
Car age (years)	6.35	3.59	2	11
Mileage	93,265	50,734	30,257	157,766
Engine displacement (liters)	3.61	1.52	2	5.7
No. Threads	78,950			

Notes: This table presents the descriptive statistics at the thread level. Final price is conditioned on trade occurring. See Appendix Tables A1–A2 for descriptive statistics of cars from the fleet/lease and dealer subsamples.

them are handled by one mediator per thread, 28.84% are handled by two mediators, and the remaining by more than two. We exclude threads involving multiple mediators. Because we want a sufficient number of observations per mediator to estimate trade probabilities, we restrict our sample to mediators that we observe participating in at least 50 separate bargaining threads. In the end, we are left with 119 mediators and 78,950 bargaining threads.

The primary measure of mediators’ performance in this industry is the probability with which they achieve agreement. The auction house makes it clear that its main goal is to facilitate as many trades as possible, offering a liquid two-sided marketplace that can attract both buyers and sellers (Lacetera et al. 2016).⁶ We will evaluate differences in mediator performance by this metric. We also look at heterogeneity in other outcomes across mediators, such as the final price of a successful deal. To make final prices more comparable across various cars, for much of our analysis we normalize prices by the auction house’s blue book estimate for the car.

⁶In contrast, an objective of maximizing prices—the target in many auction-design problems—would be attractive to sellers but unattractive to buyers.

Table 2: Descriptive Statistics at Mediator Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	.558	.225	.309	.896
Final price/bluebook price	.836	.0757	.753	.933
Final price/reserve price	.797	.0612	.729	.869
Final price/auction price	1.02	.012	1	1.03
Female	.468	.501	0	1
# Threads mediated	663	707	97	1,542
Years of employment	4.1	5.25	.396	9.44
No. Mediators	119			

Notes: This table presents the summary statistics of our main sample at the mediator level. The selected 119 mediators each handled at least 50 threads during the sample period.

Tables 1 and 2 show summary statistics of our primary estimation sample, at the level of threads and mediators respectively. Table 1 shows that the average probability of trade is 58.9%. The average final price for a successful trade is \$5,659, lower than the average blue book price of \$6,969. The average final price is also in between the average auction price of \$5,575 and average secret reserve price of \$7,361, though it is much closer to the former. A bargaining thread ends after 1.37 rounds of offers on average, and within six hours. An average car is 6.35 years old, has an odometer reading of 93,265 miles and 3.61 liters of engine displacement. There is substantial variation in car characteristics. 47.3% of the cars in our sample are from fleet/lease sellers and the rest come from dealers. Fleet/lease cars are on average newer and have higher trade probability and negotiated price. Appendix Tables A1–A2 show descriptive statistics of cars from both types of sellers separately, where we limit to mediators who handled at least 25 threads of fleet/lease cars or 25 threads of dealer cars, retrospectively. We focus on the full sample for most of our results in the body of the paper, but use these subsamples in Section 6.

Table 2 shows mediators’ characteristics and their negotiation outcomes. Among the 119 mediators, 46.8% of them are female. An average mediator has worked in the auction house for four years.⁷ During the sample period, she handles 663 bargaining threads and

⁷We obtain years of employment for a mediator 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 threads

successfully facilitates 55.8% of them. The achieved final price is, on average, 80% of the reserve price and 2% higher than the auction price. Dispersion in average trade probability is quite large across mediators, with a standard deviation of 0.23, while the dispersion in prices is much smaller.

4 Heterogeneity in Outcomes Across Mediators

In this section, we provide empirical evidence on mediators’ heterogeneous performance as measured by their probability of achieving agreement rate and by the final price of their mediated bargaining threads.

4.1 Mediator Effects on Agreement Rates

We begin by evaluating the raw differences across the 119 mediators in our sample in terms of their probability of achieving agreement. Let Y_i be an indicator variable representing whether trade/agreement occurs for thread i , which is mediated by intermediary $k(i)$. Let $\mathcal{J}_k$ be the set of all threads mediated by k , and $|\mathcal{J}_k|$ the number of elements in this set. We estimate the average of Y_i over the subset of threads mediated by intermediary k ,

$$\bar{Y}_k \equiv \frac{1}{|\mathcal{J}_k|} \sum_{i \in \mathcal{J}_k} Y_i. \quad (1)$$

These probabilities are shown in Figure 1, sorted from smallest to largest. This figure illustrates a core stylized fact of our study: different mediators have very different probabilities of bargaining success. The best to worst mediators span a difference of 98.4% in the probability of trade. The 90–10 percentile spread is 58.6%. The 25th percentile mediator achieves agreement with probability 40.6%, while the 75th percentile mediator probability is 75.4% (these are shown with the horizontal dotted lines in Figure 1). We will refer to a mediator with a higher probability of trade as a *higher-skilled* mediator, as the probability of trade is the primary metric upon which individual mediators are evaluated by the auction house.

---

she handles during the sample period.

Figure 1: Mediator Differences in Trade Probability

Notes: This figure shows the average trade probability for each of the 119 mediators in the data. We rank mediators from low to high based on this performance metric. Dashed lines represent the trade probability achieved by the 25th and 75th percentile mediators.

To better understand the industry, we spent time in these auction houses, observing the bargaining/mediation process and interviewing mediators, buyers, and sellers. Through this effort, we learned about potential drivers of mediator heterogeneity and the mediator assignment process. These conversations reveal 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 is assigned to mediate negotiations involving that seller. This default mediator is often unavailable, however, and in such cases an alternative mediator handles the negotiation. This is particularly the case for dealer-sold cars, where, for several exciting hours each week, the sales room in which the mediating phone calls occur is alive with activity, akin to a miniature stock market trading floor, with no pre-specified protocol determining which mediator ends up handling a given thread. With fleet/lease cars, it is more often the case that the pre-assigned mediator is available, but even there this is not always the case.

Mediator assignment for a given negotiation thus involves an element of randomness, but is by no means completely random. We evaluate the extent of this non-random assignment

here, and identify a number of features of the negotiation that we can control for to help account for non-random assignment. We first construct a leave-one-out estimate of a mediator’s agreement rate—the average probability of trades for all threads other than i that are mediated by intermediary k —given by

$$\bar{Y}_{-i,k(i)} \equiv \frac{1}{|\mathcal{J}_k|-1} \sum_{j:j \neq i, j \in \mathcal{J}_k} Y_j.$$

We then estimate the following regression:

$$\bar{Y}_{-i,k(i)} = Z_i' \gamma + \epsilon_i, \quad (2)$$

where Z_i contains characteristics of negotiation thread i that are determined before the mediator is assigned, and hence before the negotiation begins. These include fixed effects for the auction house location, as well as controls for vehicle mileage, age, engine displacement, the auction house blue book estimate, and the high bid and reserve price from the auction. We also include in Z_i a leave-one-out average of the agreement probability of the agent who is the seller on thread i . This seller leave-one-out average is constructed similarly to the mediator leave-one-out rate, taking the empirical frequency with which seller $s(i)$ (the seller on thread i) comes to agreement, using only threads in which seller s was involved but omitting thread i . We construct a similar leave-one-out agreement probability for the buyer.⁸

We first plot the estimated coefficients from estimating (2) with the outcome being the main agreement indicator; that is, we regress Y_i (rather than the leave-one-out agreement rate $\bar{Y}_{-i,k(i)}$) on thread characteristics Z_i . These estimates are shown in panel A of Figure 2. 95% confidence intervals are shown around each estimated coefficient. Panel A shows that each of these thread features 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

---

⁸In similar notation to $\bar{Y}_{-i,k(i)}$, we can write the seller leave-one-out agreement rate as $\bar{Y}_{-i,s(i)}^{seller} \equiv \frac{1}{|\mathcal{J}_s|-1} \sum_{j:j \neq i, j \in \mathcal{J}_s} Y_j$, 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 $\bar{Y}_{-i,b(i)}^{buyer} \equiv \frac{1}{|\mathcal{J}_b|-1} \sum_{j:j \neq i, j \in \mathcal{J}_b} Y_j$, where $\mathcal{J}_b$ is the set of all threads with b as the buyer.

reserve price or higher auction price, if the car is a less-expensive car (lower blue 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).

Nonrandom assignment is potentially a concern if good mediators are systematically assigned to threads with characteristics associated with higher probabilities of trade: for example, if effective mediators are systematically assigned to less expensive cars, or cars with lower reserve prices. To test whether this holds in our data, we regress the leave-one-out estimates of mediator fixed effects on characteristics: formally, we estimate (2), using the leave-one-out agreement rate $\bar{Y}_{-i,k(i)}$ as the outcome variable. If good mediators are not systematically assigned to better threads, then the coefficients on this regression should be 0.

We show these estimated coefficients in panel B of Figure 2, again with 95% confidence intervals surrounding each point. The results suggest that assignment is not completely random: better mediators—those with a higher leave-one-out average trade probability—tend to be assigned to threads with lower reserve prices and higher auction prices. We also find that these better mediators are more likely to be assigned to bargaining pairs in which the buyer or seller has a higher probability of coming to agreement.⁹ However, these effects are small relative to the magnitudes in panel A: the F-statistic decreases from 453 to 22, and the coefficients are about five times smaller. This implies that thread-level features have much less power to predict leave-one-out agreement rates, compared to realized agreements. These results suggest that, while there is some element of randomness in mediator assignment, part of the raw heterogeneity in mediator performance from Figure 1 is not specifically due to the mediator, but instead arises from characteristics of the thread. This highlights the importance of conditioning on these factors as we pursue estimating the effect of the mediator on bargaining outcomes.

To measure mediator performance incorporating these controls, we expand (1) to estimate

---

⁹In Appendix A, we offer an alternative test of how mediators are assigned to bargaining threads. We find that mediators are fairly randomly assigned to car types and buyers. However, the assignment to sellers is less likely to be random, especially for the fleet/lease subsample.

Figure 2: Mediator Assignment Test

Notes: This figure shows coefficients from (2). Panel A shows coefficients when the outcome is the main agreement indicator Y_i . Panel B shows estimates from the same regression but with the leave-one-out agreement rate $\bar{Y}_{-i,k(i)}$ as the dependent variable. We show 95% confidence intervals, computed by clustering at the mediator level. See Appendix Tables A4–A5 for the assignment test results using fleet/lease and dealer subsamples separately.

a fixed effect for each mediator conditional on thread-level characteristics, as follows:

$$Y_i = \beta_{k(i)} + X_i' \gamma + \epsilon_i. \quad (3)$$

In (3), $\beta_{k(i)}$ is the effect for mediator k and X_i is a vector of controls varying from specification to specification. For each specification, we follow Lacetera et al. (2016) and re-position the mediator fixed effects so they have mean zero, as follows:

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

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

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

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 our base case, specification 1, X_i includes various thread-level features: the car's age, blue book value, engine displacement and mileage; a dummy for whether the car is sold by a fleet/lease seller; and the reserve price and auction price. These latter two variables are among those that Figure 2 demonstrates are particularly important to include to control for mediator assignment.

We then sequentially add a number of other fixed effects to regression (3). In specification 2, we add fixed effects for the make and model of the car and the date of the negotiation. In specification 3, we add buyer fixed effects. Specification 4 replaces buyer fixed effects with seller fixed effects. These seller and buyer fixed effects help control for the possibility (highlighted in Figure 2) that mediator assignment may depend on the buyer or seller. In

the most-saturated version of our model—specification 5—we include all of these controls together. Figure 3 shows results comparing these different specifications, ranking mediator fixed effects from smallest to largest as in Figure 1. While adding these various controls changes the shape of the mediators’ fixed effect curve to some extent, the heterogeneity remains relatively consistent and large across specifications. The inner-quartile range is large in each specification: the minimum value of the 75th to 25th percentile gap is 22.03%.

Figure 4 illustrates how mediator fixed effects are correlated across these specifications. In each panel, we plot, on the vertical axis, a mediator’s rank (from 1 to 119) as measured by her fixed effect from the most-saturated specification (specification 5 from Figure 3). On the horizontal axes we plot the mediator’s rank from other specifications (1–4). We find that our ranking of mediator performance is highly correlated across specifications. Regardless of which controls we include, we find strong evidence for mediators having heterogeneous effects on the probability of trade, and the ranking of mediator performance are estimated relatively consistent across our different regression specifications.

4.2 Mediator Effects on Prices

We repeat the estimation of (3) but using final negotiated prices (normalized by the auction house bluebook price) as the outcome of interest rather than the agreement indicator. Panel B of Figure 5 shows the estimated mediator fixed effects under the saturated model (specification 5), along with pointwise 95% confidence intervals surround these estimated mediator effects. Panel A shows the estimates of the trade-probability effects (also under specification 5) for comparison. The estimated mediator fixed effects on normalized final prices are smaller in magnitude than those on trade probability outcomes. A one-standard-deviation increase in mediator performance is only associated with a 3.8 percentage point increase in the final price, but with a 32.2 percentage point increase in the trade probability. Confidence intervals for each estimated effect are also wider in terms of price measure: 118 out of 119 mediator effects are not significantly different from zero. None of these results are surprising given the objective of the mediators, which is clearly delineated as a goal to increase the probability

Figure 4: Mediator Rank Across Specifications

Notes: This figure presents mediators' rank in trade probabilities across different specifications. Each data point represents a mediator. The y-axis shows mediator fixed effects from specification 5, and the x-axis shows fixed effects from other specifications. We add a 45 degree dashed line as a reference.

of the trade irrespective of the price. Our results here demonstrate that indeed mediators do not differ in terms of the final prices at which they get parties to agree, and do differ drastically in their agreement probabilities.

Figure 5: Mediator Fixed Effects Under Most-Saturated Specification

Notes: Panel A shows mediator fixed effect estimates for trade probability along with 95% confidence intervals, under fully specified model (specification 5 of Equation (3)). Panel B shows results from the same exercise but using final prices normalized by bluebook value as the outcome variable. Dashed lines represent the estimated fixed effect for the 25th and 75th percentile mediators.

4.3 Sampling Error

A concern with any study of heterogeneity is that at least part (and potentially all) of the observed heterogeneity may be driven by sampling error: even if true mediator performance is constant across mediators, some variance in outcomes would arise in any finite sample. To examine this possibility, we conduct a placebo test by repeatedly shuffling bargaining outcomes, assigning the outcome from each observation to a randomly chosen mediator, and then re-estimating specification 5 of regression (3). We repeat this 100 times, and construct a 95% pointwise confidence band for the CDFs of estimated fixed effects under random assignment. We plot CDF of the median placebo estimates, as well as these 95% confidence bands, together with the CDF of the real mediator fixed effect estimates from specification 5, both for trade probabilities and for prices.

The results are shown in Figure 6. The estimated mediator fixed effects for prices fall within the 95% confidence band of the placebo estimates. This implies that we cannot reject the null hypothesis that mediators have no effect on prices (and that any apparent effect observed in panel B of Figure 5 is due to sampling error). As a robustness check, we run the same placebo test for the final prices normalized by the auction price and reserve price rather than the blue book price. We do not find that mediators differ by these measures either (see Appendix Figure A8).

In contrast, the distribution of trade probability fixed effects is quite dispersed, lying well outside the 95% confidence bands of the placebo distribution, suggesting that the heterogeneity in trade probability is not simply an artifact of sampling noise. Our analysis suggests that, if heterogeneity were purely driven by sampling error, we would expect the difference in trade probability between the 75th percentile and 25th percentile mediator to be about 6.1%; the actual inner quartile range is much larger than this, and thus cannot simply be explained by finite-sample error.¹⁰

---

¹⁰A Kolmogorov-Smirnov test comparing the actual CDF of trade probability fixed effects to the CDF of the median placebo effect yields a p-value of 0.00. A similar test for price fixed effects yields a p-value of 0.13, implying that we cannot reject the null hypothesis that the price heterogeneity is driven by sampling error. Note that this placebo analysis captures an idea analogous to what would be captured by applying a method of Bayes' shrinkage to the mediator fixed effects.

Figure 6: Placebo Test of Heterogeneity from Sampling Error

Notes: Panel A presents placebo test results for trade probability. Panel B presents results for final price normalized by bluebook value. To obtain the placebo estimates and confidence intervals, we first shuffle bargaining outcomes, assign the outcome from each observation to a randomly chosen mediator, and then re-estimate specification 5 of regression (3). We repeat this process 100 times. Estimated distribution of fixed effects from the actual data are shown in blue; placebo median and 95% pointwise confidence bands are shown in black. See Appendix Figures A6–A7 for placebo test results from fleet/lease and dealer samples separately.

5 Decomposing Bargaining Success

Having established that mediators differ substantially in their ability to help negotiating parties reach an agreement, we now turn to the question of *how* they do so. Because we have detailed data on the back-and-forth offers that buyers and sellers make in the process of bargaining, we can study how the bargaining threads of effective mediators differ from those of less-effective mediators. In this section, we propose a simple way to decompose the probability of bargaining success into two sets of probabilities: *agreement probabilities* and *continuation probabilities*. We show that effective mediators improve bargaining outcomes entirely by increasing the former set of probabilities.

Formally, for any given bargaining thread, we index rounds by t , where $t = 1$ is the first round. Let T represent the max number of rounds observed in data. If a bargaining thread reaches round t , the round t agent has three mutually exclusive choices: agree to trade and end bargaining, an event which we denote $Agree_t$; disagree and leave bargaining,

which we denote $Leave_t$; or disagree but continue to the next round, which we denote $Cont_t$. Let $P_t(Agree_t)$, $P_t(Leave_t)$, and $P_t(Cont_t)$ denote the probabilities of each event occurring, conditional on bargaining reaching round t ; hence, these three probabilities always sum to 1. Let $Disagree_t \equiv Leave_t \cup Cont_t$ represent disagreement in round t , which results in either leaving ($Leave_t$) or continuing to the next round ($Cont_t$). Define

$$P_t(Cont_t | Disagree_t) \equiv \frac{P_t(Cont_t)}{P_t(Cont_t) + P_t(Leave_t)}$$

as the probability of continuing to round $t + 1$ conditional on disagreeing in round t .

Let $P(Success)$ represent the probability that bargaining ultimately ends in agreement, and let $P_t(Success)$, with a t subscript, denote the probability that bargaining ultimately ends in agreement *conditional* on bargaining reaching round t . Bargaining can succeed either if agents agree in round t , or if agents disagree but continue to round $t + 1$ and agree in some subsequent round; that is, we can inductively define $P_t(Success)$ as

$$P_t(Success) = P_t(Agree_t) + (1 - P_t(Agree_t)) P_t(Cont_t | Disagree_t) P_{t+1}(Success), \quad (4)$$

with the terminal condition $P_t(Success) = P_T(Agree_T)$. Applying (4) to $t = 1$, we can represent the unconditional probability of bargaining success as:

$$P(Success) = P_1(Agree_1) + (1 - P_1(Agree_1)) P_1(Cont_1 | Disagree_1) P_2(Success). \quad (5)$$

Expressions (4) and (5) show that bargaining outcomes in period t can be summarized by two numbers: $P_t(Agree_t)$, the probability that agents agree in period t , and $P_t(Cont_t | Disagree_t)$, the probability that agents do not agree and continue on to period $t + 1$. Moreover, the probability of success is an increasing function of all of these terms: increasing either $P_t(Agree_t)$ or $P_t(Cont_t | Disagree_t)$ for any round, holding all other terms fixed, increases $P(Success)$. The terms $P_t(Agree_t)$ and $P_t(Cont_t | Disagree_t)$ thus allow us to decompose the ultimate probability of bargaining success into two conditional probabilities for each bargaining round.

Table 3: Bargaining Decomposition and Counterfactuals

A. Probabilities	Agreement Prob.			Continuation Prob.		
	Low	Medium	High	Low	Medium	High
Round 1	0.341	0.509	0.667	0.643	0.593	0.404
Round 2	0.129	0.131	0.148	0.298	0.229	0.199
Round 3	0.503	0.538	0.614	0.627	0.713	0.693
Round 4	0.401	0.477	0.485			
B. Counterfactuals	Change $P_t(agree_t)$			Change Only $P_1(agree_1)$		
	Low	Medium	High	Low	Medium	High
Actual Probability	0.465	0.587	0.704	0.465	0.587	0.704
Counterfactual	0.738	0.725	0.704	0.729	0.720	0.704
Ratio	1.588	1.234	1.000	1.569	1.226	1.000

Notes: The first three columns in panel A shows agreement probabilities across rounds for 1st, 2nd, and 3rd terciles of mediators, ranked by their agreement fixed effects. The last three columns in panel A shows continuation probabilities across rounds for 1st, 2nd, and 3rd terciles of mediators, ranked by their agreement fixed effects. The first three columns in panel B shows counterfactual trade probabilities, assuming all mediators' agreement probabilities are equal to agreement probabilities for the top tercile of mediators. The last three columns in panel B shows counterfactual trade probabilities, assuming all mediators' first-round agreement probabilities are equal to first-round agreement probabilities for the top tercile of mediators.

Formally, this decomposition is an accounting identity; it does not rely on any economic or statistical model of the world. Intuitively, however, the terms $P_t(agree_t)$ and $P_t(cont_t | disagree_t)$ can be thought of as representing “soft” versus “hard” behavior in bargaining. Increases in $P_t(agree_t)$ are increases in “soft” behavior, in the sense that agents back down and simply accept offers. Increases in $P_t(cont_t | disagree_t)$ are increases in “persistence”, in the sense that agents do not agree more, but persist to round $t + 1$ rather than giving up and leaving bargaining.

In panel A of Table 3, we first show $P_t(agree_t)$ and $P_t(cont_t | disagree_t)$ separately for the top, middle, and bottom terciles of mediators (titled *low*, *medium*, and *high*), ranked by their estimated fixed effects for the probability of trade. We see that $P_t(agree_t)$ is higher for better mediators. In round 1, low-skilled mediators achieve agreement 34% of the time, medium-skilled mediators 51% of the time, and high-skilled mediators 67% of the time. This increasing pattern holds true in each round of the game. In contrast, continuation probabilities, $P_t(cont_t | disagree_t)$, are not monotonically related to mediator skill, and in

most rounds are actually somewhat lower for higher-skilled mediators. In round 1, disagreeing negotiating pairs continue on to the next round 64% of the time when the negotiation is handled by a low-skilled mediator, 59% of the time under a medium-skilled mediator, and only 40% of the time under a high-skilled mediator.

Using our decomposition, we can also quantitatively measure the relative contributions of agreement and continuation probabilities to increased bargaining success rates as follows. Holding all continuation probabilities fixed at the level of low-skilled mediators, we change agreement probabilities $P_t(Agree_t)$ for each round to the level corresponding to high-skilled mediators and then calculate the counterfactual bargaining success probability using expressions (4) and (5). The results are shown in the first three columns of panel B of Table 3. We find that $P_t(Agree_t)$ explains more than 100% of the effect of good mediators; that is, effective mediators improve outcomes because they increase agreement probabilities, not continuation probabilities.

Finally, most of the effect can be explained purely using $P_1(Agree_1)$, the probability of period 1 agreement. The last three columns of panel B of Table 3 shows counterfactual trade probabilities, assuming we increase $P_1(Agree_1)$ to the round-1 agreement probability of high-skilled mediators. The results suggest that the probability of agreement in the first round explains over 100% of the effect of middle and top tercile probabilities, again consistent with quicker agreement—not persistence—being the key explanatory factor in eventual success.

6 The Effect of Mediators on Welfare

In this section, we analyze the effects of mediator skill on welfare. In particular, we ask whether mediators who achieve a higher probability of trade also achieve higher overall efficiency as measured by their expected gains from trade, or whether instead these high-trade-probability mediators are simply capturing low-surplus trades (i.e., cases where the buyer values the car only slightly more than the seller, which adds 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).

For a given bargaining pair, we treat the buyer as having a private value, $B \sim F_B$ (with density f_B), and the seller as having a private value, $S \sim F_S$ (with density f_S), where B and S are assumed to be independent. We also assume that the actions observed in bargaining sequences mediated by a given mediator represent play from a Bayes Nash equilibrium (BNE). By the Revelation Principle (Myerson 1979), such an equilibrium has an equivalent *direct revelation mechanism* in which agents truthfully report their types to a mechanism designer who ensures outcomes occur as they would have had the original game played out. The key function for summarizing any such mechanism is the *allocation function*, which specifies the probability with which a given seller and buyer pair should trade.

In the post-auction bargaining setting we study, a direct mechanism can be thought of as $x(s, b; p^A) = 1\{b \geq g(s; p^A)\}$, where $g(\cdot)$ is some function of the seller's value that, conditional on the auction price P^A being equal to p^A , defines the region in (s, b) space where trade occurs and where it does not.¹¹ Larsen (2021) demonstrates that, in any BNE of this game, the allocation function does indeed take this step-function form, and that the secret reserve price is a strictly increasing function of s , which we denote $r = \rho(s)$. Hence, we can re-write the allocation function as a function of r (which is observable to the econometrician) rather than s , which is privately known only to the seller. The expected gains from trade in the bargaining stage of the game under any mechanism $x(r, b; p^A)$ can be computed as

$$\int_{\underline{b}}^{\bar{b}} \left[\int_{\underline{b}(p^A)}^{\bar{b}} \int_{\underline{s}(p^A)}^{\bar{s}} (b - s) x(\rho(s), b; p^A) f_S(s|p^A) f_B(b|p^A) ds db \right] f_{p^A}(p^A) dp^A, \quad (6)$$

where $f_S(s|p^A) = \frac{f_S(s)}{1 - F_S(\underline{s}(p^A))}$ and $f_B(b|p^A) = \frac{f_B(b)}{1 - F_B(\underline{b}(p^A))}$ are the densities of seller and buyer values conditional on the auction price P^A and on bargaining occurring. The objects $\underline{s}(p^A)$ and $\underline{b}(p^A)$ are the lower bound on seller and buyer types, respectively, conditional on bargaining occurring. The unconditional support of B is given by $[\underline{b}, \bar{b}]$, and $\bar{s}$ is the upper bound of the support of S . To estimate (6), we import estimates of these conditional densities and supports from Larsen (2021). We depart from Larsen (2021), however, in that we

---

¹¹Throughout this section, we use uppercase letters to denote random variables and lowercase letters to denote realizations.

estimate the mechanism $x(r, b; p^A)$ separately for low-, medium-, and high-skilled mediators. In essence, we allow for each of these three mediator skill types to be implementing a distinct *mechanism*, allowing these mechanisms to differ in the gains from trade they achieve (i.e. how *efficient* they are).

To simplify exposition in the body of the paper, we ignore game-level heterogeneity in explaining the identification and estimation of the function g . Let $\mathcal{A}$ be the event that trade occurs. The function g in $1\{b \geq g(r; p^A)\}$ is related to the conditional probability of agreement as follows:

$$\Pr(\mathcal{A}|R = r, P^A = p^A) = \frac{1 - F_B(g(r, p^A))}{1 - F_B(p^A)}. \quad (7)$$

Given the objects R , P^A , and $\mathcal{A}$ observed in data, the left-hand side of (7) can be non-parametrically estimated; we use a tensor product of cubic b-splines. The object $g(\cdot)$ can then be estimated given knowledge of F_B (the distribution of buyer values) by inverting the right-hand side. We perform estimation of $g(\cdot)$ separately for bargaining sequences in the data handled by low-, medium-, and high-skilled mediators, and denote these estimates g^L , g^M , and g^H .

With these estimates in hand, we can evaluate total surplus, i.e. the gains from trade, given by (6). To perform this step, we import estimated upper and lower bounds on F_S and point estimates of F_B from Larsen (2021), who estimates these objects separately for fleet/lease vs. dealer cars. We therefore estimate the mediator-specific allocation functions — g^L , g^M , and g^H — separately for fleet/lease vs. dealer cars. This requires merging our identifiers of mediator ability with the data used in Larsen (2021), which only partially overlaps with our data because some observations used in his study do not have mediator identifiers, and some observations used in our paper do not have variables recorded that are used in his analysis.¹²

In practice, we implement a much more sophisticated estimation process than that de-

---

¹²This merge yields the following samples of low-, medium-, and high-skilled mediators, respectively: 5,879, 9,113, and 6,483 in the dealers sample, and 1,006, 4,851, and 6,470 in the fleet/lease sample.

scribed by (7), allowing for heterogeneity at the game level that is observable to agents and to the econometrician (such as the make, model, year, trim of the vehicle, as well as a rich set of other characteristics), and also heterogeneity that is observable to agents but unobserved to the econometrician (capturing features such as dents or odors in the car). We describe the technical details for these estimation steps in Appendix B. An underlying assumption of this analysis is that, after controlling for observed and unobserved game-level heterogeneity, the distributions of seller and buyer valuations (F_B and F_S) do not vary with the assigned mediator, and thus these objects can be taken directly from Larsen (2021). It is only the mechanism — summarized by g^L , g^M , and g^H — that we allow to vary by mediator.¹³

In addition to evaluating the gains from trade achieved in practice by each mediator type, we also estimate, for comparison, the gains from trade that would be achieved if mediators could implement the theoretically (ex-ante) efficient mechanism. Myerson and Satterthwaite (1983) derived this mechanism under the condition that both bargaining parties (the buyer and seller) have equal welfare weights. Williams (1987) extended this analysis, deriving the theoretical Pareto frontier—a continuum of mechanisms that maximize the gains from trade under different welfare weights for the buyer and seller.¹⁴ The first-best, ex-post-efficient mechanism for bilateral trade is simply an indicator function for whether the buyer values the good more than the seller. The seminal Myerson-Satterthwaite Theorem shows that this first-best mechanism is infeasible whenever the supports of buyer and seller values overlap.¹⁵ Our estimates of mediator-specific mechanisms allow us to study whether better mediators, as measured by their ability to achieve trade, are also able to move significantly closer to the Pareto frontier, potentially overcoming some of the inefficiency inherent in incomplete-

---

¹³As described in the Appendix B, the estimation of bounds on F_S relies on the seller’s decision to accept or reject the auction price. For most sales, this decision occurs immediately following the auction, without any mediated phone call taking place, but for about 20% of sales this decision could potentially be influenced by the mediator, and our choice to hold F_S fixed regardless of which mediator handles the negotiation abstracts away from this possibility.

¹⁴Williams (1987) shows that the Pareto frontier is a function of F_S and F_B , and demonstrates how to compute the frontier in the case where these distributions satisfy regularity (Myerson 1981). Larsen (2021) demonstrates how to compute the frontier in a way that is robust to nonregularity, and we follow this approach here. Loertscher and Marx (2021) extends the theoretical analysis of Williams (1987) to the case of multiple buyers/sellers.

¹⁵Specifically, their impossibility result states that no individually rational, incentive-compatible mechanism achieves ex-post efficiency while maintaining ex-ante budget balance.

information bargaining.

Figure 7: Bargaining Relative to Second-Best and First-Best Frontiers

Notes: Each panel displays estimated expected seller and buyer surplus on the ex-post efficient frontier (dashed line), on ex-ante efficient frontier (solid line), and in real-world bargaining under mediator skills (solid dots, with low-skill in blue, medium-skill in red, and high-skill in yellow). Top panels use seller distribution lower bound and bottom panels use seller distribution upper bound. Panels on left use dealers sample and on right use fleet/lease sample. Units = \$1,000.

In Figure 7 we display the expected gains under each of the three mediator skill levels. The dashed line shows the first-best (infeasible) frontier, corresponding to the outcome if trade were to occur whenever the buyer values the car more than the seller. The solid line shows the second-best frontier derived by Williams (1987). In each panel, the horizontal axis

represents the expected gains from trade of the seller and the vertical axis the expected gains from trade of the buyer. We display these plots separately using the upper bound on F_S and lower bound on F_S . The units are \$1,000.

We observe larger differences across skill groups for the fleet/lease sample (panels B and D) than for the dealers sample (panels A and C). Figure 7 also suggests that efficiency gains across mediator skill levels are primarily achieved by increasing the expected surplus of the *seller* rather than the *buyer*. In the fleet/lease sample, we find that high-skilled mediators achieve a level of efficiency that lies approximately on the second-best frontier (and close to the first-best as well). This suggests that the most-effective mediators are able to overcome much, if not all, of the inefficiency described in Myerson and Satterthwaite (1983) and documented empirically in Larsen (2021).

Figure 8: Surplus Lost Relative to Ex-Post Efficient Bargaining

Notes: Each panel displays bounds on the difference in total welfare in the real-world mechanism and the first-best (ex-post efficient bargaining), separately for negotiation sequences mediated by low, medium, or high-skilled mediators. Dashed lines show 95% confidence bounds. Panel A uses the dealers sample and panel B uses the fleet/lease sample. Units = \$1,000.

In Figure 8, we illustrate bounds on the gap between the first-best and the mechanism implemented by each mediator type, along with 95% confidence intervals.¹⁶ These represent

¹⁶These confidence bounds are constructed by taking 40 bootstrap replications of our estimation procedure and reporting the 0.025 quantile of the estimated lower bound on the gap between the first-best and mediator mechanism and the 0.975 quantile of the estimated upper bound on this gap.

bounds on the deadweight loss in bargaining: gains that real-world bargaining fails to achieve 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.

In both panels, confidence intervals overlap between different mediator skill levels. This overlap is consistent with the possibility that, even though some mediators systematically achieve a higher *probability of trade*, they achieve similar outcomes as low-performing mediators in terms of welfare; under this interpretation, their higher probability of trade would come from realizing additional low-surplus trades missed by lower-skilled mediators.

However, the estimated bounds also offer a story of better mediators improving efficiency. In particular, the upper bound on the loss relative to first best is lower for high-skilled than for low skilled mediators in both panels A and B (and drastically so in panel B, the fleet/lease sample). For example, in panel A, we cannot rule out the possibility that low-skilled mediators miss out on \$2,000 of trade gains, but we can rule out this possibility to medium- or high-skilled mediators. In panel B, the results are much starker: we can rule out that high-skilled mediators lose any more than about \$350 of surplus. This is consistent with better-skilled mediators potentially significantly increasing overall welfare by realizing high-surplus trades that low-skilled mediators fail to consummate.

7 Conclusion

In this paper, we have shown that mediators have statistically significant and economically large effects on bargaining outcomes. Quantitatively, we find that the 75th percentile intermediary is 22% more likely to close a deal than the 25th percentile intermediary. Our estimated mediator effects are robust to a variety of different controls, and are not driven by sample noise.

We proposed a way to decompose bargaining success into acceptance and continuation probabilities, and show that effective mediators improve outcomes entirely through improving acceptance probabilities. This result suggests that, in our context, a higher likelihood of

bargaining success comes from getting negotiating parties to agree more, not bargain longer or harder; indeed, most of the effects can be explained by good mediators increasing the probability of agreement *soon* or even *immediately*. These results suggest that better mediators may be those who understand, early on in a negotiation, the propensity of a given pair to agree.

Our structural exercise demonstrated that better mediators can not only improve the probability that agreement is reached, but can also have real effects on realized gains from trade. Myerson and Satterthwaite (1983) demonstrated that incomplete-information bilateral bargaining 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 that are left on the table. Our findings in this paper demonstrate that, relative to low-performing mediators, higher-performing mediators are able to execute a mechanism that lies much closer to the efficient outcome. Implementing something close to the Myerson and Satterthwaite (1983) second-best-efficient mechanism requires commitment on the part of the mediator—the discipline to let some trades fail even when gains from trade exist in order to preserve incentive compatibility.¹⁷ We see each of these key findings of our study as starting points for future experimental, theoretical, and empirical work in incomplete-information bargaining settings.

References

- Ausubel, L. M. and Deneckere, R. J. (1992). Durable goods monopoly with incomplete information. *Review of Economic Studies*, 59(4):795–812.
- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives*, 11(1):109–126.
- Barwick, P. J. and Pathak, P. A. (2015). The costs of free entry: an empirical study of real estate agents in greater boston. *RAND Journal of Economics*, 46(1):103–145.
- Basak, D. (2015). Transparency and delay in bargaining. Working paper, New York University.

---

¹⁷A similar force arises in the theoretical model of Fanning (2021), where a mediator improving efficiency requires being able to commitment to not always let the parties trade even if when the mediator knows both parties are privately in favor of it.

- 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. *Organizational Behavior and Human Decision Processes*, 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. *Negotiation Journal*, 17(1):59–77.
- Bochet, O., Khanna, M., and Siegenthaler, S. (2020). Beyond dividing the pie: An experimental study on bargaining over multiple issues. Available at SSRN 3414288.
- Bruno, B., Garcia-Appendini, E., and Nocera, G. (2018). Experience and brokerage in asset markets: Evidence from art auctions. *Financial Management*, 47(4):833–864.
- Copic, J. and Ponsati, C. (2008). Robust bilateral trade and mediated bargaining. *Journal of the European Economic Association*, 6(2-3):570–580.
- Cramton, P. C. (1992). Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59(1):205–225.
- Dixon, W. J. (1996). Third-party techniques for preventing conflict escalation and promoting peaceful settlement. *International Organization*, 50(4):653–681.
- Eisenkopf, G. and Bachtiger, A. (2013). Mediation and conflict prevention. *Journal of Conflict Resolution*, 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. *Journal of consulting and clinical psychology*, 59(3):410.
- Fanning, J. (2021). Mediation in reputational bargaining. *American Economic Review*, 111(8):2444–72.
- Gallant, A. and Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica*, 55(2):363–390.
- Gilbukh, S. and Goldsmith-Pinkham, P. (2021). Heterogeneous real estate agents and the housing cycle. Working paper, City University of New York.
- Glode, V. and Opp, C. (2016). Asymmetric information and intermediation chains. *American Economic Review*, 106(9):2699–2721.
- Goltsman, M., Hörner, J., Pavlov, G., and Squintani, F. (2009). Mediation, arbitration and negotiation. *Journal of Economic Theory*, 144(4):1397–1420.
- Gottardi, P. and Mezzetti, C. (2019). Mediation design. Working paper, University of Queensland.

- Gul, F., Sonnenschein, H., and Wilson, R. (1986). Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39(1):155–190.
- Haile, P., Hong, H., and Shum, M. (2003). Nonparametric tests for common values at first-price sealed-bid auctions. Working paper, Yale University.
- Hörner, J., Morelli, M., and Squintani, F. (2015). Mediation and peace. *Review of Economic Studies*, 82(4):1483–1501.
- Jindal, P. and Newberry, P. (2020). The profitability of revenue-based quotas under price negotiation,. Working paper, Penn State. Forthcoming, *Management Science*.
- Keniston, D. (2011). Bargaining and welfare: A dynamic structural analysis. Working Paper, Yale University.
- Kim, J. Y. (2017). Interim third-party selection in bargaining. *Games and Economic Behavior*, 102:645–665.
- Kim, S. (2020). Informed agent’s advice: Theory and evidence. Working paper, University of Buffalo.
- Kotlarski, I. (1967). On characterizing the gamma and normal distribution. *Pacific Journal of Mathematics*, 20(1):69–76.
- Kydd, A. (2003). Which side are you on? Bias, credibility, and mediation. *American Journal of Political Science*, 47(4):597–611.
- Lacetera, N., Larsen, B. J., Pope, D. G., and Sydnor, J. R. (2016). Bid takers or market makers? The effect of auctioneers on auction outcomes. *American Economic Journal: Microeconomics*, 8(4):195–229.
- Lang, M. (2020). Mechanism design with narratives. CESifo Working Paper 8502.
- Larsen, B. and Zhang, A. L. (2018). A mechanism design approach to identification and estimation. NBER Working Paper 24837.
- Larsen, B. J. (2021). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. *Review of Economic Studies*, 88(2):851–882.
- Loertscher, S. and Marx, L. M. (2021). Incomplete information bargaining with applications to mergers, investment, and vertical integration. Working paper, University of Melbourne.
- 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. *Journal of Economic Theory*, 29(2):265–281.

- Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica: journal of the Econometric Society*, pages 61–73.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Myerson, R. B. (2008). Perspectives on mechanism design in economic theory. *American Economic Review*, 98(3):586–603.
- Neale, M. A. (1984). The effects of negotiation and arbitration cost salience on bargainer behavior: The role of the arbitrator and constituency on negotiator judgment. *Organizational Behavior and Human Performance*, 34(1):97–111.
- Robles-Garcia, C. (2020). Competition and incentives in mortgage markets: The role of brokers. Working paper, Stanford University.
- 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.
- Valley, K., Thompson, L., Gibbons, R., and Bazerman, M. H. (2002). How communication improves efficiency in bargaining games. *Games and Economic Behavior*, 38(1):127–155.
- Valley, K. L., White, S. B., Neale, M. A., and Bazerman, M. H. (1992). Agents as information brokers: The effects of information disclosure on negotiated outcomes. *Organizational Behavior and Human Decision Processes*, 51(2):220–236.
- Williams, S. (1987). Efficient performance in two-agent bargaining. *Journal of Economic Theory*, 41(1):154–172.
- Yavas, A., Miceli, T. J., and Sirmans, C. (2001). An experimental analysis of the impact of intermediaries on the outcome of bargaining games. *Real Estate Economics*, 29(2):251–276.
- Zhang, X. (2019). Efficient bargaining through a broker. Available at SSRN 2753074.

A An Alternative Test of Assignment

Here we offer an alternative test of how mediators are assigned to bargaining threads. We begin by randomly shuffling 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 this same exercise but instead of car types we randomly shuffle buyers, and then we repeat the exercise a third time shuffling sellers.

In Figures A1, A2, and A3 respectively, we plot the number of unique car makes, buyers, and sellers that a mediator interacts with in the real data, against the mean number of unique car makes, buyers, and sellers that the mediator interacts with in the simulated 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 Figure A2 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. In this analysis, we separately analyze assignment for cars in the fleet/lease vs. dealers subsamples.

In Figures A1 and A2, most data points lie close to the $y = x$ line, suggesting that the assignment of mediators to car types and to buyers appears to be fairly random. On the other hand, Figure A3 shows that mediators interact with fewer sellers in the real data than they do in the shuffled dataset (particularly for fleet/lease cars), suggesting that the assignment of mediators to sellers is less likely to be random, and highlighting the importance of the seller fixed effects we include in our preferred specification (specification 5) of regression (3).

Figure A1: Car Random Assignment Test

Notes: Each data point is a mediator. The x-axis shows the number of car make-models the mediator interacts with in the data. The y-axis shows the mean number of car make-models the mediator interacts with when assignments of car make-models to mediators are randomly reshuffled.

Figure A2: Buyer Random Assignment Test

Notes: Each data point is a mediator. The x-axis shows the number of buyers the mediator interacts with in the data. The y-axis shows the mean number of buyers the mediator interacts with when assignments of buyers to mediators are randomly reshuffled.

Figure A3: Seller Random Assignment Test

Notes: Each data point is a mediator. The x-axis shows the number of sellers the mediator interacts with in the data. The y-axis shows the mean number of sellers the mediator interacts with when assignments of sellers to mediators are randomly reshuffled.

B Technical Details of Model Estimation

In this section, we describe the key pieces for estimating a distinct mechanism for low, medium, and high-skilled mediators. This description largely follows Larsen (2021), and we refer the reader to that paper for details on other estimation steps, including the distributions of private valuations; we import these estimates directly from Larsen (2021) as they are not specific to mediator skill.

Let R^{raw} and PA^{raw} be random variables representing the reserve price and auction price in the raw data, prior to any adjustments for heterogeneity. Let W be a random variable representing game-level heterogeneity that is observable to players but not the econometrician. Let X be a random variable representing game-level heterogeneity that is instead observed by the econometrician and players, with X independent of W , S , and B . Let realizations of R^{raw} , PA^{raw} , X , and W for game j be denoted by lowercase letters with subscript j . We specify the total game-level heterogeneity (observed plus unobserved) for observation j to be $x_j' \gamma + w_j$, where γ is a vector of parameters to be estimated.

We use a standard *homogenization* approach (Haile et al. 2003) to account for observable

heterogeneity by estimating the following joint regression of reserve prices and auction prices on observables:

$$\begin{bmatrix} r_j^{raw} \\ p_j^{A,raw} \end{bmatrix} = \begin{bmatrix} x_j' \gamma \\ x_j' \gamma \end{bmatrix} + \begin{bmatrix} \tilde{r}_j \\ \tilde{p}_j^A \end{bmatrix}, \quad (8)$$

where $\tilde{r}_j = r_j + w_j$, $\tilde{p}_j^A = p_j^A + w_j$. The vector x_j includes a rich set of controls.¹⁸ An estimate of $\tilde{r}_j$ is then given by subtracting $x_j' \hat{\gamma}$ from r_j^{raw} , and similarly for $\tilde{p}_j^A$. Variation in these two quantities can then be attributed to unobserved game-level heterogeneity and to players' private values.

To account for heterogeneity W in the game that is observed by the players but not by the econometrician, 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 the densities f_W , f_R , and f_{P^A} . We estimate these densities using a flexible maximum likelihood approach, where the likelihood of the joint density of $(\tilde{R}, \tilde{P}^A)$ is given by

$$\mathcal{L}(f_{P^A}, f_R, f_W) = \prod_j \left[\int f_{P^A}(\tilde{p}_j^A - w) f_R(\tilde{r}_j - w) f_W(w) dw \right] \quad (9)$$

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

---

¹⁸We take our estimate of $\hat{\gamma}$ directly from Larsen (2021), who includes the following in x_j : fifth-order polynomials in the auction houses' blue-book estimate, the odometer reading, run number within an auction-house-by-day combination, and the 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 online; a dummy for whether or not the odometer reading is considered accurate, and the interaction of this dummy with the odometer 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. 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 in consideration. The six categories we consider are make, make-by-model, make-by-age, make-by-model-by-age, age, or seller identity.

(as in Gallant and Nychka 1987).

Accounting for these heterogeneity components, the conditional probability of trade described in (7) in the body of the paper becomes

$$\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A) = \int \frac{1 - F_B(g(\tilde{r} - w, \tilde{p}^A - w))}{1 - F_B(\tilde{p}^A - w)} \left(\frac{M_g(\tilde{r}, \tilde{p}^A, w)}{\int M_g(\tilde{r}, \tilde{p}^A, z) dz} \right) dw \quad (10)$$

where $M_g(\tilde{r}, \tilde{p}^A, w) \equiv f_R(\tilde{r} - w)f_{P^A}(\tilde{p}^A - w)f_W(w)$ is the joint density of (R, P^A, W) . To exploit the relationship in (10), we approximate $h_g(r, p^A) \equiv \frac{1 - F_B(g(r, p^A))}{1 - F_B(p^A)}$ using a flexible bilinear spline parameterized by θ^g , with 25 knots in each dimension, uniformly spaced between the 0.001 and 0.999 quantiles of $\tilde{R}$ and $\tilde{P}^A$. We estimate θ^g using constrained least squares by evaluating the left-hand side and right-hand side of (10) on a fixed grid of points and searching for the parameters θ^g to minimize the distance between the left- and right-hand sides. This requires first estimating the conditional probability $\Pr(\mathcal{A}|\tilde{R} = \tilde{r}, \tilde{P}^A = \tilde{p}_A)$, which we do using a tensor product of cubic b-spline functions with fifteen uniformly spaced knots in each dimension. With estimates of $\hat{\theta}^g$, we then obtain $g(r, p^A) = F_B^{-1}(1 - (1 - F_B(p^A))h_g(r, p^A; \hat{\theta}^g))$. We perform this exercise separately using observations handled by low-, medium-, or high-skilled mediators, to obtain estimates of g^L , g^M , and g^H .

The other objects required to evaluate the expected gains from trade in (6) are the distributions F_B and F_S . Estimates of F_B are obtained using a standard order statistics inversion, and bounds on F_S are obtained using rationality assumptions combined with observations of the seller's decision to accept or reject the auction price. Bayes updating of agents' beliefs translates these marginal distributions to distributions conditional on the auction price P^A and on the agents having arrived at the bargaining stage of the game. We refer the reader to Larsen (2021) for details on these estimation steps, as well as nonparametric identification proofs for each of these objects.

Table A1: Descriptive Statistics at Bargaining-Thread Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	.747	.435	0	1
Final price (\$)	5,678	4,747	1,000	12,300
Bluebook price (\$)	7,157	5,189	1,650	14,375
Auction price (\$)	5,604	4,720	900	12,100
Reserve Price (\$)	7,566	5,174	2,000	14,700
# Offers in a thread	1.15	.462	1	2
Length of a thread (hours)	5.31	13.5	.369	7.73
Fleet/lease car	1	0	1	1
Car age (years)	5.6	3.31	2	10
Mileage	88,616	50,092	28,722	151,486
Engine displacement (liters)	3.59	1.56	2	5.7
No. Threads	37,266			

Notes: This table shows the descriptive statistics of the fleet/lease sample at the thread level. Final price is conditioned on trade occurring.

Table A2: Descriptive Statistics at Bargaining-Thread Level

	Mean	Std. Dev.	0.1 Quantile	0.9 Quantile
Agreement reached	.448	.497	0	1
Final price (\$)	5,636	5,227	800	13,400
Bluebook price (\$)	6,803	5,317	1,500	14,600
Auction price (\$)	5,552	5,112	800	13,100
Reserve Price (\$)	7,211	5,577	1,800	15,308
# Offers in a thread	1.57	.789	1	2
Length of a thread (hours)	6.42	16.6	.286	21.5
Fleet/lease car	0	0	0	0
Car age (years)	7.03	3.69	2	12
Mileage	97,439	50,877	31,884	162,438
Engine displacement (liters)	3.62	1.49	2	5.7
No. Threads	41,641			

Notes: This table shows the descriptive statistics of the dealers sample at the thread level. Final price is conditioned on trade occurring.

Figure A4: Assignment Test, Fleet/lease sample

Notes: This figure shows coefficients from (2) estimated only on the fleet/lease subsample. Panel A shows coefficients when the outcome is the main agreement indicator Y_i . Panel B shows estimates from the same regression but with the leave-one-out agreement rate $\bar{Y}_{-i,k(i)}$ as the dependent variable. We show 95% confidence intervals, computed by clustering at the mediator level.

Figure A5: Assignment Test, Dealer sample

Notes: This figure shows coefficients from (2) estimated only on the dealers subsample. Panel A shows coefficients when the outcome is the main agreement indicator Y_i . Panel B shows estimates from the same regression but with the leave-one-out agreement rate $\bar{Y}_{-i,k(i)}$ as the dependent variable. We show 95% confidence intervals, computed by clustering at the mediator level.

Figure A6: Placebo Test of Heterogeneity from Sampling Error, Fleet/lease Sample

Notes: This figure presents placebo test results for the fleet/lease sample. Estimated distribution of fixed effects from the actual data are shown in blue; placebo median and 95% pointwise confidence bands are shown in black.

Figure A7: Placebo Test of Heterogeneity from Sampling Error, Dealer Sample

Notes: This figure presents placebo test results for the dealer sample. Estimated distribution of fixed effects from the actual data are shown in blue; placebo median and 95% pointwise confidence bands are shown in black.

Figure A8: Placebo Test of Heterogeneity from Sampling Error, Other Price Measures

Notes: This figure presents placebo test results for the main sample, where the outcome of interest is the final price divided by the reserve price in panel A and the final price divided by the auction price in panel B. Estimated distribution of fixed effects from the actual data are shown in blue; placebo median and 95% pointwise confidence bands are shown in black.