

Mediation in Bargaining: Evidence from Large-Scale Field Data on Business-to-Business Negotiations

Bradley J. Larsen, Carol Hengheng Lu, and Anthony Lee Zhang¹

March 30, 2020

Abstract

We analyze a dataset containing hundreds of thousands of full 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, and these mediators are quasi-randomly assigned to the bargaining pair. We find mediator's identities matter: high-performing mediators are 23.23% more likely to close a deal than low-performers. Experience is correlated with better mediator performance. Male and female mediators perform equally well, but mediate differently: female mediators close deals faster and at prices more favorable to buyers. Good mediators appear to respond to long-term company incentives rather than short-term incentives to close a given deal and they can do even better at reaching agreement for threads with ex-ante lower probability of trade. We provide a new decomposition of mediator effectiveness, demonstrating that effective mediators improve bargaining outcomes by causing buyers and sellers to come to agreements faster, not by causing buyers and sellers to be more persistent. We also show that better mediators appear less reliant on exploiting certain types of behavioral biases.

¹Larsen: Stanford University and NBER; bjlarsen@stanford.edu. Lu: Stanford University; hllu@stanford.edu. Zhang: University of Chicago Booth School of Business; anthony.zhang@chicagobooth.edu

1 Introduction

Many real-world bargaining situations—between nations, between businesses, between investors, between consumers, between disputing parties in a legal case—involve a third-party *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 mediating firm acquisitions or lawyers mediating pre-trial settlement). To date, however, there is little quantitative evidence from real-world data on whether and how mediators make a difference for negotiation outcomes.

This paper provides an analysis of a massive dataset containing hundreds of thousands of 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 and 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 bargaining process between the highest bidder and the seller of the car.

The data we study is rich, containing information from six different auction houses that sold hundreds of thousands of cars from 2007–2010. 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 and the timing and location of the 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 mediation as well. We view this as an unprecedented opportunity to study bargaining mediators in the field.

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) can be considered quasi-randomly assigned to the particular car and bargaining pair. We find evidence that this random assignment is especially strong for cars sold by used-car dealerships, which tend to be smaller clients of the auction house. For cars sold by larger clients (fleet/lease sellers), the seller in the negotiation often has a default assigned mediator who handles their negotiations. But even for these large sellers, the default mediator is sometimes not available and an alternative mediator takes on the case. This quasi-random assignment gives us a natural experiment—or thousands of little natural experiments—in which mediators vary from transaction to transaction, allowing us to make a first attempt at quantifying whether (and how) mediators matter for big ticket, real-world negotiations.

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.

Our first finding is 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 23.23% higher than a 25th percentile mediator. While mediators have large effects on trade probability, they have statistically insignificant effects on the prices at which trades occur. Our estimated confidence intervals show that we are able to detect effects on the order of 5% of average prices. In other words, mediators appear to affect trade probability, and thus total trade surplus, without substantially changing how the pie is split between buyers and sellers.

We perform a number of robustness checks for our results. Our results survive controlling for various features of bargaining sequences, such as the identity of buyers and sellers involved in bargaining, and various features of the car being sold. Given that some variation in mediator performance would be expected in any dataset simply due to statistical noise (even if true underlying performance were equal across mediators), 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 analyze what kinds of mediators achieve better outcomes. We find that experienced mediators are more effective: one extra year of experience is associated with an 8.36% higher probability of successful trade. We also find male and female mediators are equally likely

to achieve success in a trade, but they do so differently: male mediators achieve a higher price (favoring the seller) and take longer (more back-and-forth offers and more overall time negotiating) while female mediators finish the deal quickly at a lower price (favoring the buyer). Given that a majority of auction house fees comes from a fixed (non-commission) component paid by the buyer and seller whenever trade occurs, female and male mediators perform equally well from an platform-revenue perspective.

We find that better mediators tend to do even better at reaching agreement for threads with ex-ante lower probability of trade. These low-probability deals include cases in which the potential gains from trade (the difference between the buyer’s and seller’s valuation for the car) is much lower. We show that better mediators do not perform better on deals with larger short-term incentives (negotiations involving sellers who will pay a higher fee to the company if a deal is reached) and instead perform better on deals that involve important long-term relationships with the company (i.e. sellers who sell many cars through the auction house).

Bargaining parties in our data appear to have some behavioral biases, such as relying on offers that split the difference between the two most recent offers on the table. Interestingly, we find that negotiations with better mediators are less likely to exhibit this type of behavior, suggesting that better mediators may be less reliant on exploiting this type of bias.

Finally, we analyze how mediators are able to achieve better outcomes. We show that, in alternating-offer bargaining games, 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 a “continuation

probability”. We find that effective mediators improve bargaining outcomes by increasing agreement probabilities, not continuation probabilities: in fact, continuation probabilities are actually somewhat lower for effective mediators.

The effects we document in this paper are much larger than effects of mediators in trade in other 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 to influence outcomes; one could think of mediators as playing a kind of “equilibrium selection” role. We view the contribution of this paper as empirically demonstrating a situation in which the size of mediators’ effects on outcomes are precisely measurable and quantitatively very 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), Fanning (2019), and Horner et al. (2018). These studies suggest possible ways in which mediators can and cannot affect outcomes. For example, mediators who do not perfectly reveal the seller’s or buyer’s value to the opposing party, but who noisily reveal that value, can improve bargaining outcomes. Basak (2015) similarly 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.

A mediator might affect bargaining in a number of ways. A number of studies in organizational behavior and 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 increase 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 this analysis beyond the lab to a high-stakes negotiation in the field between professional negotiators.

We also relate to the economics literature on gender in negotiations. This literature has documented that women negotiators are less likely to enter negotiations than men (Leibbrandt and List 2015; Exley et al. 2016) and tend to ask for less than men (Hernandez-Arenaz and Iriberry 2018, 2019). Our analysis is unique relative to these studies in that we study the gender of the *mediator*, while previous studies have focused on the gender of one or both of the *negotiating parties*.

Our study relates to other empirical studies that document heterogeneous outcomes in sales situations in the field, such as Gilbukh and Goldsmith-Pinkham (2018), studying real estate agents; Lacetera et al. (2016), studying auctioneers in the same industry as we study; Bruno et al. (2018), studying art auctions; Jindal and Newberry (2019), studying heterogeneity across sellers in large-scale appliances, and Backus et al. (2018), studying heterogeneity across buyers and sellers in online e-commerce negotiations. Relative to much of 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.²

Finally, the data we use overlaps to some extent with the alternating-offer bargaining data studied in Larsen (2019) and Larsen and Zhang (2018), although these studies do not exploit information on the identity of the mediator. These studies provide structural welfare analyses of the performance of the bilateral bargaining (Larsen 2019) and the combined auction-plus-bargaining (Larsen and Zhang 2018) and show that the bargaining mechanism

²Several recent papers also study detailed data on back-and-forth actions in bargaining, including Keniston (2011), Bagwell et al. (2015), Larsen (2019), Larsen and Zhang (2018), and Backus et al. (2018).

falls short of first-best efficiency. These studies leave unanswered the question *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 probability of trade can be explained by mediator influence.

3 Institutional background and data

The wholesale used-car auction industry is the backbone of the supply-side of the used-car market (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 cars are sold to other dealerships. 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. Total revenue from these sales 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) 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 by the auction house employee—the mediator—first calling up the seller and reporting 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.

To better understand the industry, we have spent time in these auction houses, observed the bargaining process, and interviewed mediators, buyers, and sellers. The following true dialogue, related to us by one of the mediators who appears in our data, illustrates one of the tricks a mediator might employ. In each statement below, the seller and buyer speak only to the mediator, and we clarify to whom the mediator is speaking:

Seller: “I will let the car go for \$3,200.”

Mediator (to buyer): [Misrepresenting the seller offer] “The seller is asking for *\$3,600*.”

Buyer: “No way. Tell the seller I will only pay \$3,200.”

Mediator (to buyer): “I will have to check.”

Mediator (to seller): “Ok, deal. The buyer is good with your \$3,200 offer.”

Mediator (to buyer): “Ok, deal. The seller is good with your \$3,200 offer.”

A number of insights can be gleaned from this example. For instance, the mediator did not truthfully reveal to the buyer what the seller had said. The mediator withheld information until both parties had actually (unknowingly) come to an agreement. But the above interaction also suggests that behavioral (non-rational) elements may be at play in the mediator’s actions. In particular, the mediator also made each party *feel* as if he or she had made the final offer—as if he or she were the one setting the terms. This example suggests that a mediator who is better at employing these practices may be able to achieve agreement more often.

Not all of this detailed dialogue is contained in the data we analyze, but the auction house records all actions taken by either party, as well as the identity of each party and the

identity of the mediator, and any notes taken by the mediator during the bargaining. The data consists of several hundred thousand realizations of bargaining sequences recorded by the auction houses (six different auction house locations) between 2007 and 2010. We will also use the term *thread* to refer to a given bargaining sequence.

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: auction price, reserve price, and blue book price. The database creates a new record for each action taken during a given bargaining sequence, allowing us to see the timing of each action and also allowing us to see that some sequences involve several different mediators recording different stages of the negotiation. Among all the bargaining threads, 68.76% of them are handled by a single mediator, 28.84% by two mediators and the remaining by more than two. We exclude threads involving multiple mediators for our main analysis. Because we want a fairly large number of observations per mediator to accurately 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 120 mediators and 80,285 bargaining threads.

Our primary measure of mediator performance is the rate at which they achieve agreement. The auction house makes it clear that its main goal is to facilitate as many trades as possible. We also look at some secondary performance measures such as final price of a successful deal and the speed of negotiation. 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.

Tables 1 and 2 show summary statistics of our primary estimation sample, at the level

of mediators and threads respectively. Table 1 shows that 47.3% of mediators are female. Mediators vary in their employment length from half a year to more than nine years. On average, mediators handle 669 bargaining threads during the sample period and successfully facilitate 55.4% of them. The final price is close to the auction price, around 20% below the reserve price, and around 12% below the blue book price of the car. Dispersion in average trade probabilities among mediators is quite high, with a standard deviation of 0.229, and dispersion in prices is much lower. In terms of speed of mediation, bargaining threads end within 1.35 offers and 323 minutes (about 5 hours) on average.

Table 2 shows that the average final price for successful sales is \$5,558, somewhat below the average bluebook price of \$6,854. The average final price is between the average auction price of \$5,490 and the average secret reserve price of \$7,284, though it is much closer to the auction price. Cars in our sample are on average around six years old, with mileages around 95,000, though there is substantial variation in car age and mileage in our sample. 47.6% of the cars are from fleet/lease sellers and the rest come from dealerships. When a trade happens, the auction house can earn on average \$118 from the seller.

4 The effect of mediators on bargaining outcomes

4.1 Mediator fixed effects

Our baseline analysis consists regressions of the following form. Let k index the mediator, i index a given bargaining thread. We estimate

Table 1: Descriptive Statistics, Mediators

	Mean	Std. Dev.	10% Quantile	90% Quantile
Agreement reached	.554	.229	.312	.882
# Offers in a thread	1.35	.255	1	1.71
Length of a thread (min)	323	210	121	582
Price/bluebook price	.881	.097	.771	.997
Price/reserve price	.79	.0575	.728	.858
Price/auction price	1.02	.0119	1	1.03
Gender (F=1)	.473	.502	0	1
# Threads mediated	669	719	97	1,532
Length of employment (day)	1,501	1,906	145	3,446
No. Mediators	120			

Table 2: Descriptive Statistics, Bargaining Threads

	Mean	Std. Dev.	10% Quantile	90% Quantile
Final price	5,558	4,939	800	12,650
Bluebook price	6,854	5,284	1,425	14,400
Auction price	5,490	4,924	775	12,600
Reserve price	7,284	5,411	1,800	15,000
Fleet/lease car	.476	.499	0	1
Car age	6.42	3.62	2	12
Odometer reading	94,496	51,792	30,469	160,502
Sale fee	118	53	80	185
No. Threads	80,285			

$$Y_{ik} = \alpha + \beta_k + X_i' \gamma + \epsilon_{ik}, \quad (1)$$

where Y_{ik} is an indicator for whether the bargaining thread is successful, or some other performance metrics in other specifications, including the price, the number of offers, or the length of the bargaining in time. The vector X_i includes characteristics of the bargaining thread that vary by specification in our analysis below, but can include characteristics of the car (make, model, year fixed effects), buyer and seller fixed effects, day-of-week controls, and so on. We are interested in the mediator effects, $\hat{\beta}_k$. Following Lacetera et al. (2016), we re-position these mediator fixed effects so that they have mean zero, as follows:

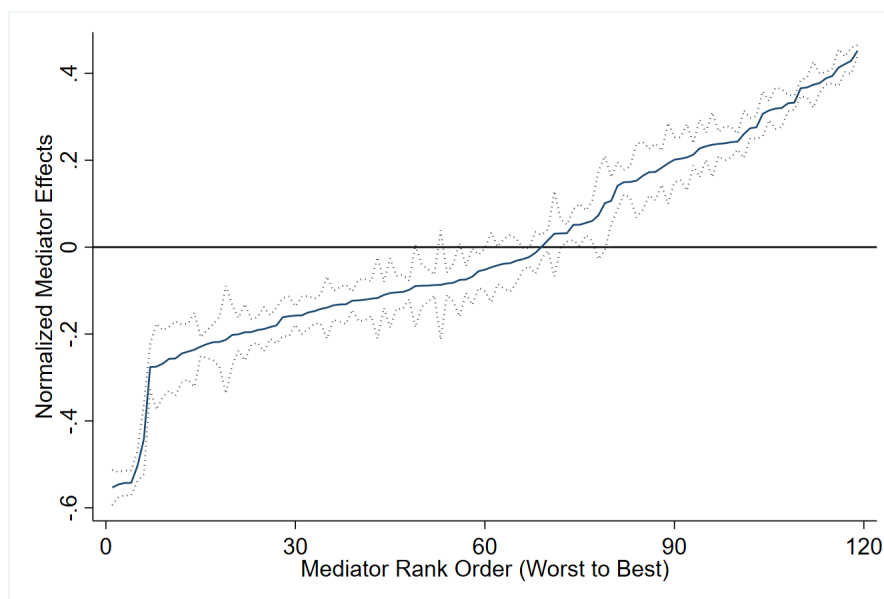
$$\hat{\beta}_{norm,k} = \begin{cases} \hat{\beta}_k - \frac{1}{M} \sum_{j=2}^M \hat{\beta}_j & \text{for } k \neq 1 \\ 0 - \frac{1}{M} \sum_{j=2}^M \hat{\beta}_j & \text{for } k = 1 \end{cases}$$

where 1 denotes the omitted mediator in specification (1).

In Figure 1, we plot our baseline estimates of these fixed effects, in which we control for various car-level features: car age, reserve price, auction price, bluebook value, mileage, a dummy for whether the car is in the fleet/lease or dealer subsamples, and the car's sale fee. We plot these estimated fixed effects sorted from smallest to largest.

Figure 1 shows the core stylized fact of this paper: different mediators have very different average probabilities of bargaining success, controlling for various features of the bargaining thread. Quantitatively, the 75th percentile mediator has a fixed effect 35.96 percentage points higher than the 25th percentile mediators; that is, one standard deviation increase in mediator

Figure 1: Mediator heterogeneity in trade probability



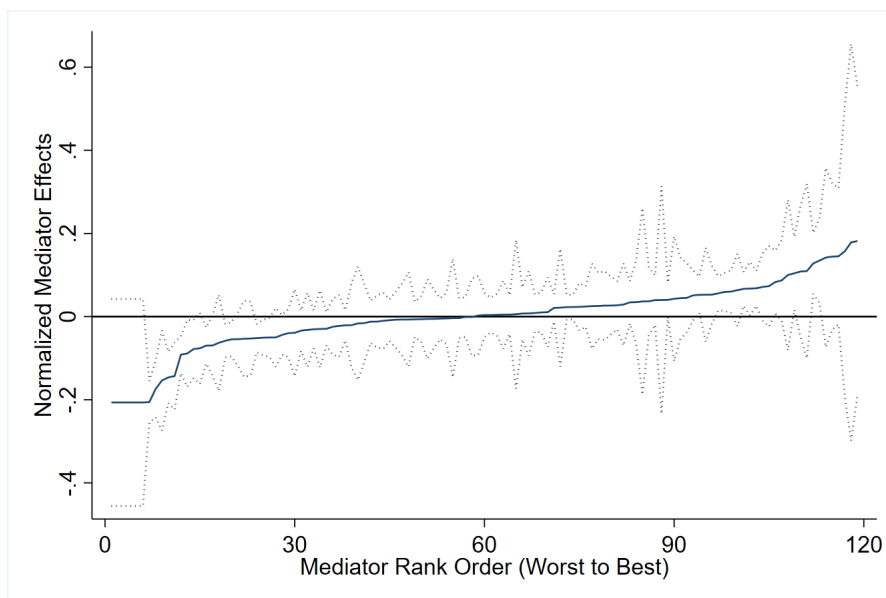
Notes: Baseline mediator fixed effect estimates for trade probability. Dotted lines represent 95% confidence intervals for the fixed effects, computed for each fixed effect based on the OLS standard error of the fixed effect.

ranking is associated with a 23.64 percentage point increase in trade probability.

Figure 2 shows the results of the baseline specification estimated using final trade prices divided by bluebook prices as the dependent variable. The estimated fixed effects are smaller in magnitude, and are mostly not statistically significant. One standard deviation increase in mediator ranking is only associated with a 6.11 percentage point increase in the final price. Confidence intervals for each estimated effect are also wider in terms of price measure: 95 out of 120 mediator effects are not significantly different from zero.

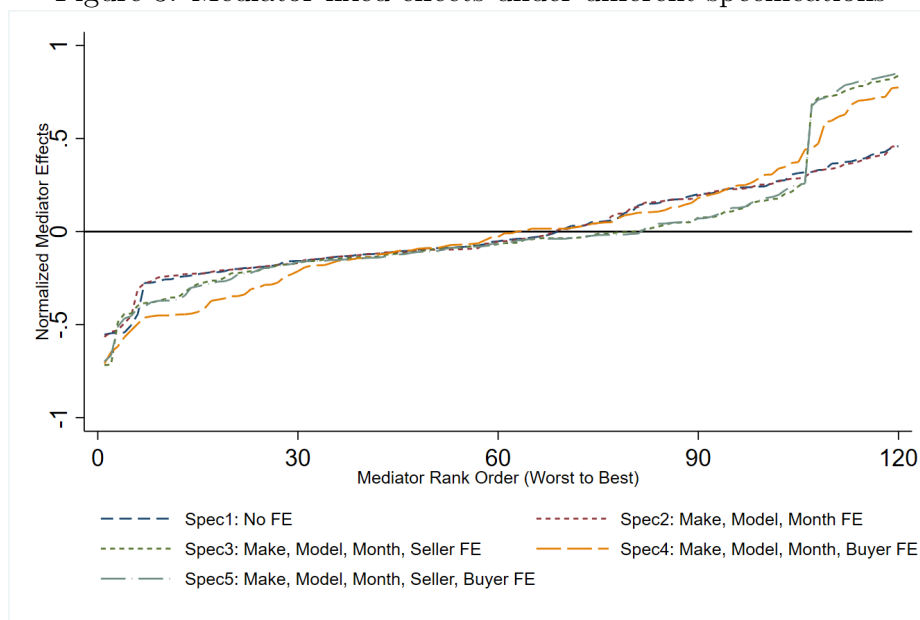
We then sequentially add a number of fixed effects to specification (1) (the baseline model). In specification (2) we add fixed effects for the make and model of the car and the year-month of the bargaining. In specification (3), we further add seller fixed effects. In specification

Figure 2: Mediator heterogeneity in final price (normalized by blue book value)



Notes: Baseline mediator fixed effect estimates for prices normalized by blue book value. Dotted lines represent 95% confidence intervals, computed for each fixed effect based on the OLS standard error of the fixed effect.

Figure 3: Mediator fixed effects under different specifications



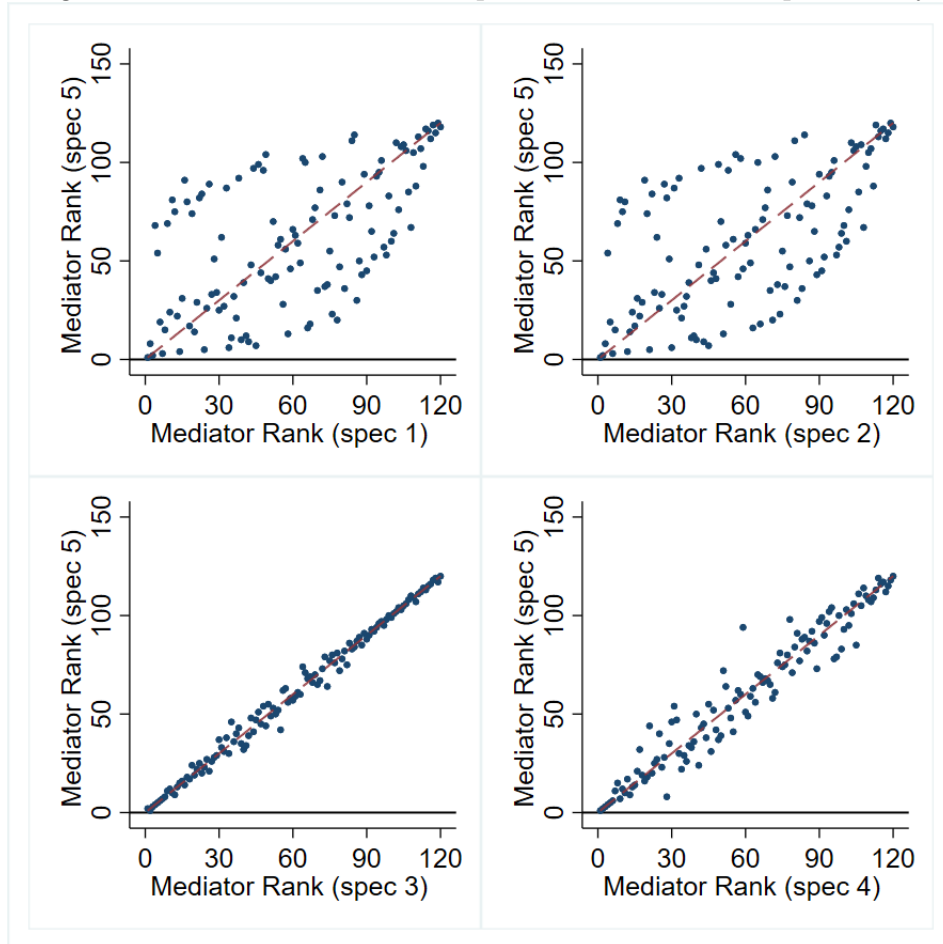
Notes: Mediator fixed effect estimates for probability of trade under different specifications with different sets of controls.

(4), we add buyer fixed effects instead. In the most saturated version – specification (5) we add all the fixed effects mentioned above. Figure 3 shows results. While adding fixed effects changes the shape of the fixed effect curve somewhat, the inner-quartile range is still large; the minimum value of the 75th to 25th percentile gap is 23.23%.

Figure 4 shows scatter plots of mediators’ fixed effect ranks across specifications. We plot the ranking from most saturated specification (specification 5 from Figure 3) against that from the other specifications (1–4). We find that our ranking of mediator performance is highly correlated across all specifications. This suggests that mediators do have heterogeneous effects on the probability of trade, which are estimated relatively consistently across our different regression specifications.

Another concern is that part of the variation we measure in estimated fixed effects is

Figure 4: Mediator rank across specifications for trade probability



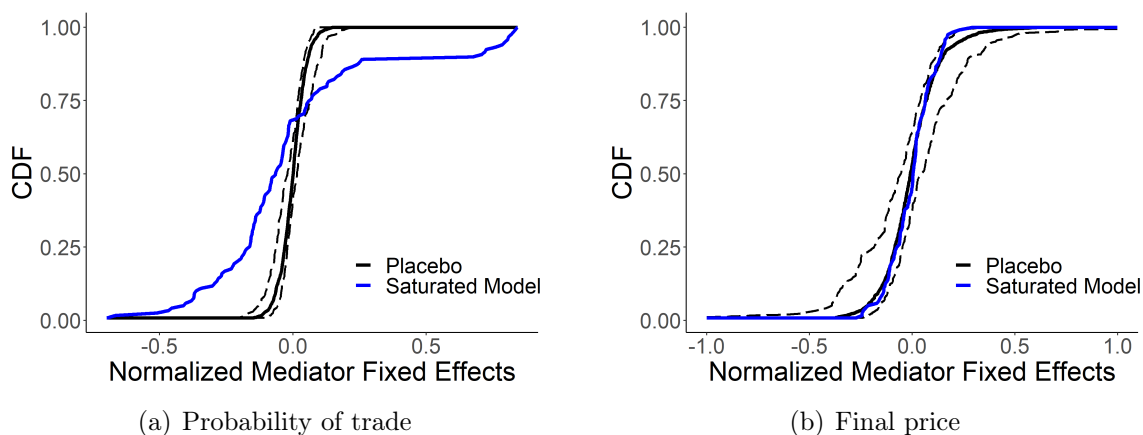
Notes: Mediator fixed effect estimates for different specifications. Each data point represents 1 mediator. The y-axis shows fixed effects from specification 5, and the x-axis shows fixed effects from other specifications.

entirely due to 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 the fixed effects. We repeat this 50 times, and construct a 95% pointwise confidence band for the CDF of estimated fixed effects under random assignment. We plot these confidence bands together with the fixed effect estimates from specification (5), for both trade probabilities and prices, in Figure 5.

The distribution of trade probability fixed effects is very disperse, lying well outside the 95% confidence bands. In words, if threads were randomly assigned to mediators, so that the probability of trade success were independent of mediator identities, we would expect the difference in trade probability fixed effects between 75th percentile and 25th percentile mediators to be around 5.97%; our estimates are much larger than this, and thus cannot simply be explained by finite-sample error. In contrast, our estimated fixed effects for prices falls within the 95% confidence band from our placebo test. This implies that we cannot reject the null hypothesis that mediators have no effect on prices. The point estimates are large, however: we cannot rule out the possibility that mediators differ in their price effects by 17% (which is one standard deviation of the estimated price fixed effects).³

³A Kolmogorov-Smirnov test comparing the actual CDF of trade probability fixed effects to the mean simulated effect yields a p-value of 0.000, further evidence that we can reject the hypothesis that the effects are generated by sampling error. A similar test for price fixed effects yields a p-value of 0.260, implying that we cannot reject the null hypothesis.

Figure 5: Placebo tests



Notes: Estimated distribution of fixed effects from specification (5) in blue, and placebo median and 95% pointwise confidence bands in black, for trade probabilities (left) and prices (right).

4.2 Testing random assignment of mediators to bargaining sequences

We now offer an empirical test of whether mediators are randomly assigned to bargaining threads. As highlighted in Section 1, industry participants argue that this assignment is as good as random for cars sold by smaller sellers; for cars sold by larger sellers, a default mediator is assigned when possible, but this default mediator is not always available to mediate a given thread, and so larger sellers are often also assigned a non-default mediator.


To test for random assignment, we randomly shuffling a car-type (make-by-model) identifier across observations within a given auction house location, year, and month. We repeat this exercise 500 times. We perform this same exercise but instead of car types we randomly shuffle seller, and then we repeat the exercise a third time shuffling buyers. In Figures 6, 7, and 8 respectively, we plot the number of unique buyers, sellers, and car makes that a mediator interacts with in the real data, against the mean number of unique sellers, buyers,

and cars respectively that the mediator interacts with in the simulated data. Each data point represents one mediator. If mediators are indeed randomly assigned to buyers, 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 6 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 sold by large fleet/lease seller (“fleet sellers”) and cars sold by used-car dealers (“dealers sellers”).

In Figures 6 and 8, most data points lie close to the $y = x$ line, suggesting that the assignment of mediators to cars and buyers appears to be fairly random. On the other hand, Figure 7 shows that mediators interact with fewer sellers in the real data than they do in the shuffled dataset, suggesting that the assignment of mediators to sellers is non-random. However, Figure 3 shows that the estimated dispersion in mediator fixed effects is fairly large even when controlling for seller fixed effects; that is, effective mediators achieve higher trade probabilities controlling for the identities of sellers they interact with. This implies that the dispersion in estimated effectiveness of mediators cannot fully be explained by nonrandom assignment of mediators to sellers.


Together, our estimates in this section suggest that mediators have economically large effects on the probability that bargaining success, and that these effects are statistically significant and relatively consistent across regression specifications with a variety of controls.

Figure 6: Buyer random assignment test



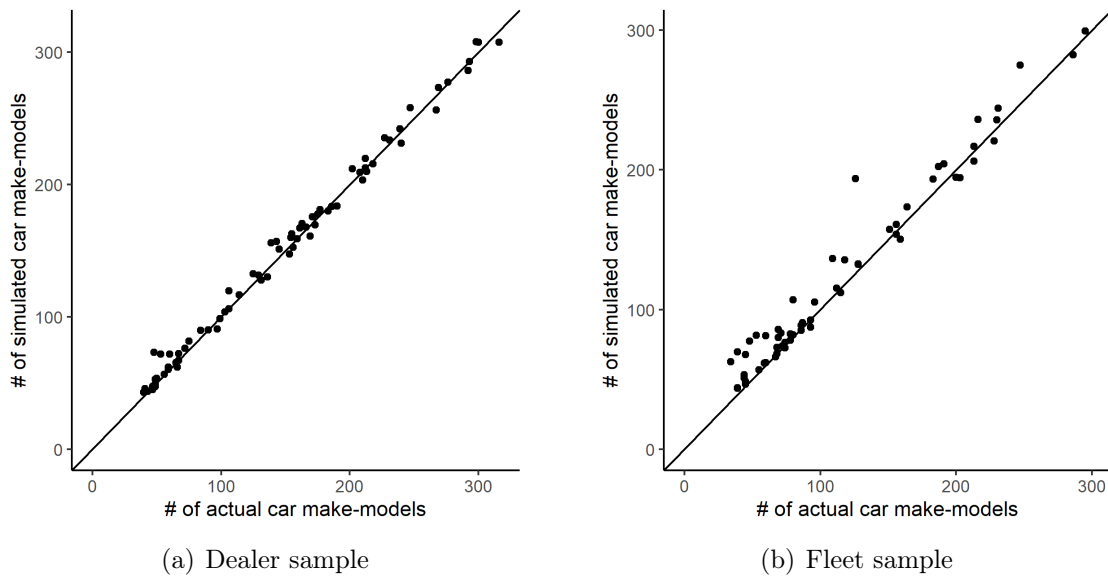
Notes: Each data point is a mediator. The x-axis shows the number of buyers the mediator interacts with in the data. The y-axis shows the mean number of buyers the mediator interacts with when assignments of buyers to mediators are randomly reshuffled.

Figure 7: 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.

Figure 8: 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.

5 Which mediators do better, and how?

5.1 What kinds of mediators do better?

We now turn to the question of what—if anything—in our data can help explain differential performance of mediators. We address this question in several different ways.

First, we regress the estimated mediator fixed effects on various characteristics of the mediators. Formally, we estimate:

$$\hat{\beta}_k = \alpha_0 + \alpha_1 \text{Female}_k + \alpha_2 \text{Log}(\text{employ})_k + \varepsilon_k, \quad (2)$$

These regressions do not include the full 120 mediators because we only observe intermediary characteristics for a subset of 82 mediators. Our two key characteristics are mediator gender and mediator employment length at the auction house. We use a baby names database to infer gender from mediators' first names. We then run simple linear regressions at the intermediary level on these two characteristics. The results are shown Table 3. In column 1, the dependent variable is the intermediary's fixed effect from specification 5. We also estimate regressions where the dependent variable is the intermediary's fixed effect for other bargaining outcomes: the number of offers per thread (column 2), the time in minutes to complete the thread (column 3), the price divided by the blue book value (column 4), and the price divided by the reserve price (column 5). In each column, an observation in the regression is a single intermediary and the dependent variable is a fixed effect from estimating specification 5 of (1) with the corresponding bargaining outcome.

A large number of previous studies have shown women negotiate differently than men. These studies tend to focus on the gender on one side of the negotiation or the other, not the gender of the mediator. We find that female mediators and male mediators are equally likely to reach an agreement (the first column), but female mediators reach agreement faster than males. This speed effect is significant when measured by the number of offers in the bargaining (column 2—females tend to have 0.087 fewer offers per thread) and insignificant, but still in the same direction, when measured in actual minutes (the point estimate in column 3 suggests that female mediators wrap up the bargaining about two hours faster than males). We find no significant effect on our main price measure, the final price divided by the blue book price (column 4), but in column 5 we do find significant differences for males and females in their fixed effects for the final price divided by the reserve price. We find that females tend to reach final prices that are farther from the seller’s reserve price, thus favoring the buyer and not the seller. Given that the primary revenue-generator the auction house company is the fee paid when trade occurs, the company cares less about the actual price (which is mainly a transfer between the buyer and seller). In this light, our results suggest that male and female mediators perform equally well.

Table 3 also shows that mediators with longer tenure tend to do better: an extra year of employment is associated with a 8.36% increase in our estimated trade probability fixed effect. These mediators also tend to end the negotiation more quickly (with fewer offers and less total negotiation time).

Table 3: Mediator Differences by Gender and Employment

	(1)	(2)	(3)	(4)	(5)
	Prob. Trade	No. Offers	Time	Price/BB	Price/Reserve
Female	0.0292 (0.0675)	-0.0870** (0.0415)	-130.3 (96.02)	0.00995 (0.0236)	-0.0178*** (0.00379)
Log(employment)	0.0836*** (0.0247)	-0.0306** (0.0141)	-97.78*** (35.25)	-0.000669 (0.00879)	-0.00218 (0.00150)
Observations	82	82	82	82	82
R-squared	0.112	0.098	0.102	0.002	0.244

Notes: Coefficient estimates from regressing intermediary fixed effects on trade probability, offer count, thread time, and two measures of prices on a gender dummy and the log of employment. Each data point is a mediator.

5.2 When Do Good Mediators Do Better?

Regression (1) assumes that effective mediators have a constant effect on the probability of successful trade; a natural way to relax this assumption is to allow mediators to have heterogeneous effects that depend on the characteristics of bargaining threads, allowing us to see whether good mediators are relatively more effective in certain kinds of negotiations. To study this question, we modify our mediator fixed effect regressions by adding mediator dummies interacted with a variety of features of the bargaining thread, while still including the main mediator fixed effects. Formally, we run regressions of the following form:

$$Y_{ik} = \alpha + \beta_k + z_i\eta + \delta_k z_i + X_i'\gamma + \epsilon_{ik}, \quad (3)$$

In these regressions, we focus on agreement as the outcome of interest Y_{ik} . The term z_i represents some feature of bargaining thread i . We consider several different possibilities for

z_i , described below. As before, the set of β_k are the main mediator fixed effects. These can be thought of as mediator-specific *constant* terms. Here the set of δ_k can be thought of as mediator-specific *slope* terms. The correlation between β_k and δ_k is informative about how sensitive better mediators are to different features of bargaining threads.

Figure 9 shows scatter plots of δ_k against β_k under various specifications for z_i . For all the plots in Figure 9, the horizontal axis shows the mediator fixed effect in reaching agreement (β_k) and the vertical axis shows the interaction effect δ_k plus the mean effect of z_i ($z_i\hat{\eta}$).

The first bargaining thread characteristic z_i with which we interact mediator dummies is a measure of the surplus, or gains from trade. The precise value of the surplus is unobserved in the data, as it depends on agents' private valuations. Our measure is constructed as the auction price divided by the reserve price. Under the structural model of Larsen (2019), the reserve price is strictly increasing in the seller's private valuation, and the expectation of the buyer's valuation is strictly increasing in the auction price (because the auction price is the second-highest-value bidders' valuation, and the buyer in the bargaining is the first-highest-value bidder). The auction price divided by the reserve price is always less than 1, because otherwise the car would have sold through the auction rather than entering the bargaining stage. As this ratio approaches 1, the expected surplus of a negotiating pair is larger.

Figure 9, panel (a), demonstrates the results of regression (3) using this surplus ratio as our measure of z_i . We find that better mediators (indicated by a higher probability-of-trade fixed effect on the horizontal axis) are better at getting trade to occur for threads with smaller surplus (indicated by a less-positive surplus effect on the vertical axis). These better mediators appear to not need a larger surplus in order to achieve a trade.

The second thread characteristic we examine is the ex-ante probability of given type of car selling—a measure of how easy this car is to sell. This probability is obtained in a separate first step by replacing mediator dummies in regression (1) with car-type (make-by-model) dummies. We set z_i to be the estimated fixed effect for the car type corresponding to the make and model of thread i . Figure 9, panel (b), shows that bad mediators are relatively more effective at selling easy-to-sell cars, whereas trade probabilities of better mediators are less sensitive to the average trade probability for a given car type. We also conduct this analysis where z_i is seller or buyer fixed effect for thread i , which measure the ex-ante probability of coming to agreement for a given buyer or seller. Similar to car agreement probabilities, panel (d) shows that better mediators are less sensitive to this ex-ante measure of a buyer’s likelihood of agreeing. Panel (c) shows the opposite relationship for sellers, though this relationship is insignificant. This may potentially be driven by the feature highlighted in Section 4.2 that mediator assignment based on seller identity does not appear as random as it does based on car type or buyer identity.

In panel (e), z_i is formed by regressing log of the car price on car type dummy (make-model effects), age and mileage, and taking the standard deviation of the residuals from this regression within a given make-by-model cell. In panel (e), we find that good mediators perform better when the pricing of car is more standard (possibly newer cars). For these cars, there may be less room to negotiate, and better mediators perform relatively better on these cars.

The last thread characteristic we examine is the expected fee the auction house will receive from the seller. Auction house fees consist of a fee paid by both the buyer and and the seller,

paid only when and if agreement is reached. Buyer fees consist mostly of a fixed fee, but also include a commission of about 1% of the final price. Seller fees, on the other hand, are comprised entirely of a fixed fee, and this fee varies depending on the seller. As buyer fees are constant across buyers, an intermediary will primarily make more immediate revenue for the auction house by getting trades to occur that involve large seller fees.


In panel (f) of Figure 9, we find that better mediators are more likely to get trades to occur on threads that involve *low* seller fees. This relationship can potentially be explained by the way in which these seller fees are determined. Large sellers (such as Ford factory, Bank of America, Hertz Rental Car) are considered important clients, and have special low sales fee agreements with the auction house. Good mediators, therefore, appear to be getting trades to occur especially well for important clients. In this sense, better mediators appear to respond to longer-term incentives of the company (keeping the market liquid for big clients) rather than short-term gains (getting only high-fee trades to occur today).

We report an alternative version of this analysis in Table 4. Here we replace the mediator dummies in regression (3) with their agreement fixed effect rankings obtained through the main specification (1). Our regression model is


$$Y_{ik} = \alpha + z_i\eta + \widehat{mediator_rank}_{i(k)}\beta + \widehat{mediator_rank}_{i(k)}z_i\delta + X_i'\gamma + \epsilon_{ik}, \quad (4)$$

The dependent variable for all columns in Table 4 is the agreement indicator for the thread. The coefficients of all z_i are positive, meaning that trade is more likely to occur when the surplus is higher; when the car, buyer, or seller has a higher ex-ante probability of agreement;


Figure 9: Interaction at mediator level



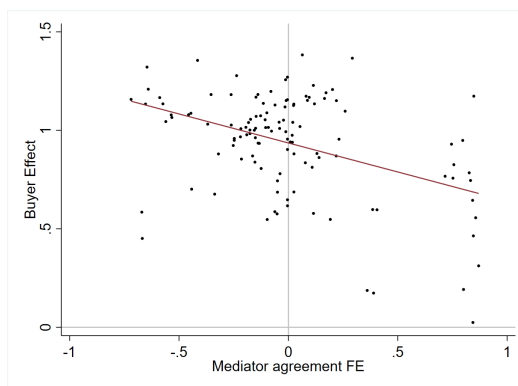
(a) Surplus interaction




(b) Car fixed effect interaction




(c) Seller fixed effect interaction



(d) Buyer fixed effect interaction



(e) Car price residual standard deviation interaction



(f) Expected seller fee interaction

Notes: Each data point is a mediator. The x -axis shows β_k , the mediator fixed effect on trade probability, and the y -axis shows the interaction effect δ_k plus the mean effect of z_i ($z_i \hat{\eta}$).

Table 4: Mediators' heterogeneous response to thread characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
$z_i =$	Surplus	Car Prob	Buyer Prob	Seller Prob	Std Dev Price	Sell fee
z_i	1.413*** (0.0281)	1.298*** (0.0674)	1.223*** (0.0169)	1.130*** (0.00905)	0.0562** (0.0226)	0.000434*** (0.0000811)
Mediator Rank	0.00877*** (0.000316)	0.00697*** (0.000137)	0.00773*** (0.000135)	0.00852*** (0.0000821)	0.00702*** (0.000190)	0.0103*** (0.000172)
Mediator Rank * z_i	-0.00414*** (0.000367)	-0.00481*** (0.00093)	-0.00387*** (0.000243)	-0.00447*** (0.000094)	-0.0000951 (0.000292)	-0.0000200*** (0.00000130)
N	55,300	76,477	78,197	78,530	74,564	77,265

Notes: Results from estimating regression (4).

when the car's price variance is larger; or when the seller fee is higher. The coefficient on the interaction term, δ represents how much a mediator's ability to get agreement depends on the thread characteristics. For example, column 1 shows that if a mediator is one rank higher, she will be 0.4% less responsive to an increase in surplus. All interaction terms in the table are negative (and all are significant except column 5), suggesting again that good mediators do not need sales to be particularly easy in order to succeed. Rather, better mediators perform even better when the situation is harder to deal with, and they do not appear to react as strongly as poor mediators to the short-term gains for the company (the seller fee).

5.3 Decomposing bargaining success

Because we have detailed data on the back-and-forth offers that buyers and sellers make in the process of bargaining, we can study in detail 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, and 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 end bargaining, which we will call A_t , disagree and leave bargaining, L_t , or disagree but continue to the next round, which we will call C_t . We will use $P_t(A_t)$, $P_t(L_t)$, and $P_t(C_t)$ to denote the probabilities of each event occurring, conditional on bargaining reaching round t ; hence, these three probabilities always sum to 1. Let $D_t \equiv L_t \cup C_t$ represent disagreement in round t , which results in either leaving L_t or continuing (to the next round) C_t . Then define

$$P_t(C_t | D_t) \equiv \frac{P_t(C_t)}{P_t(C_t) + P_t(L_t)}$$

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

Let S represent the event that bargaining ultimately succeeds, and define $P_t(S)$ as the probability of ultimate success 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 bargaining succeeds in some subsequent round; that is, we can inductively define $P_t(S)$ as:

$$P_t(S) = P_t(A_t) + (1 - P_t(A_t)) P_t(C_t | D_t) P_{t+1}(S) \tag{5}$$

with the terminal condition: $P_T(S) = P_T(A_T)$. Applying (5) to $t = 1$, we can represent the

unconditional probability of bargaining success as:

$$P(S) = P_1(S) = P_1(A_1) + (1 - P_1(A_1)) P_1(C_1 | D_1) P_2(S) \quad (6)$$

Expressions (5) and (6) show that bargaining outcomes in period t can be summarized by two numbers: $P_t(A_t)$, the probability that agents agree in period t , or $P_t(C_t | D_t)$, the probability that agents do not agree, but continue on to period $t + 1$. Specifying $P_t(A_t)$ and $P_t(C_t | D_t)$ allows us to calculate probabilities, and thus the ultimate probability of bargaining success $P(S)$. Moreover, the probability of success is an increasing function of all these terms: increasing either $P_t(A_t)$ or $P_t(C_t | D_t)$ for any round, holding all other terms fixed, increases $P(S)$. The terms $P_t(A_t)$ and $P_t(C_t | D_t)$ thus allow us to decompose the ultimate probability of bargaining success into two conditional probabilities for each bargaining round.

Formally, this decomposition is an accounting identity, it does not rely on any economic or statistical model of the world. Intuitively, however, the terms $P_t(A_t)$ and $P_t(C_t | D_t)$ can be thought of as representing “soft” versus “hard” behavior in bargaining. Increases in $P_t(A_t)$ are increases in “soft” behavior, in the sense that agents back down and simply accept offers. Increases in $P_t(C_t | D_t)$ 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.

Expressions (5) and (6) allow us to quantitatively decompose effective mediators’ effect on sale probability into the probabilities $P_t(A_t)$ and $P_t(C_t | D_t)$. In Panel A of Table 5, we first show $P_t(A_t)$ and $P_t(C_t | D_t)$ separately for the top, middle, and bottom terciles of mediators (titled low, medium, and high), ranked by their estimated fixed effects. We see that $P_t(A_t)$ is

higher for better mediators, but $P_t(C_t | D_t)$ is actually somewhat lower: the probability that any given round concludes in agreement is higher with effective mediators, but conditional on disagreement, buyers and sellers are more likely to walk away than to continue to the next round.

Using our decomposition, we can also quantitatively measure the relative contributions of agreement and continuation probabilities to increased bargaining success rates as follows. We hold all terms $P_t(C_t | D_t)$ at their level for the bottom tercile of mediators, change $P_t(A_t)$ to their values for the middle and top terciles, and calculate the counterfactual bargaining success probability using expressions (5) and (6). We find that $P_t(A_t)$ explains more than 100% of the effect of good mediators in Panel B of Table 5; that is, effective mediators improve outcomes because they increase agreement probabilities, not continuation probabilities.

Finally, most of the effect can be explained purely using $P_1(A_1)$, the probability of period 1 agreement. In Table 5, the “Counterfactual prob” row shows counterfactual trade probabilities, assuming we increase $P_1(A_1)$ from its value for the lowest tercile to its value for the middle and top terciles, holding fixed all other agreement and continuation probabilities at their values for the bottom tercile. Once again, the right table in Panel B shows that the probability of agreement in the first round explains over 100% of the effect of middle and top tercile probabilities.

Intuitively, the fact that effective mediators influence outcomes through agreement probabilities rather than continuation probabilities suggests that, in our context, higher bargaining success probabilities come from agreeing more, not bargaining harder; most of the effects can be explained by good mediators increasing the probability of agreement *soon* or even

Table 5: Bargaining success decomposition and counterfactuals

A. Probabilities	Agreement			Continuation		
	Low	Medium	High	Low	Medium	High
Round 1	0.334	0.526	0.671	0.635	0.594	0.405
Round 2	0.124	0.140	0.147	0.284	0.245	0.193
Round 3	0.499	0.560	0.623	0.633	0.692	0.690
Round 4	0.402	0.506	0.488			

B. Counterfactuals	Change $P_t(A_t)$			Change Only $P_1(A_1)$		
	Low	Medium	High	Low	Medium	High
Actual Probability	0.452	0.608	0.707	0.452	0.608	0.707
Counterfactual	0.452	0.619	0.739	0.452	0.610	0.730
Ratio	1.000	1.020	1.045	1.000	1.004	1.032

Notes: The left table in Panel A shows agreement probabilities across rounds for 1st, 2nd, and 3rd terciles of mediators, ranked by their agreement fixed effects. The right table in Panel A shows continuation probabilities across rounds for 1st, 2nd, and 3rd terciles of mediators, ranked by their agreement fixed effects. The left table in Panel B shows counterfactual trade probabilities, assuming all mediators’ agreement probabilities are equal to agreement probabilities for the top tercile of mediators. The right table 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.

immediately.

5.4 Behavioral Biases: Splitting the Difference

We now examine whether mediators differ in their exploitation of negotiators’ behavioral biases. The particular type of bias on which we focus is a preference of negotiating parties to “split the difference”—to propose counteroffers that lie halfway between the two most recent offers. This type of behavior may be justified by a rational model in certain complete-information environments, where parties are splitting a known surplus, but in an environment such as ours, where parties likely have privately known valuations for the item, proposing a

split-the-difference offer is likely driven by biases toward fairness.


We measure split-the-difference behavior by noting that the offer in round t of the bargaining game (p_t) by a player should be convex combination of most recent opponent offer (p_{t-1}) and most recent own offer (p_{t-2}). We denote the weight on this convex combination by γ_t , where $p_t = \gamma_t p_{t-1} + (1 - \gamma_t) p_{t-2}$. A value of γ_t closer to 1 indicates greater concession by the offering party, and closer to zero indicates greater stubbornness. Figure 10 shows histograms for γ_3 , γ_4 , and γ_5 , each of which show clear modes at $1/2$ (splitting the difference) and at 0 (stubbornness).

Similar patterns have been documented elsewhere in the literature in alternating-offer bargaining data (e.g. Backus et al. 2018). A distinctive feature of our environment is that we find that the likelihood of these split-the-difference offers is associated with the *skill* of the mediator. We show this result in Figure 11. The first panel shows kernel density estimates of γ_3 separately for the low, medium, and high performing mediators based on agreement rates. We find that higher-skilled mediators have less split-the-difference behavior in their negotiations. The second and third panels show the same result for γ_4 and γ_5 . This is suggestive evidence that these better mediators may be less reliant on exploiting this particular type of behavioral bias.

6 Discussion

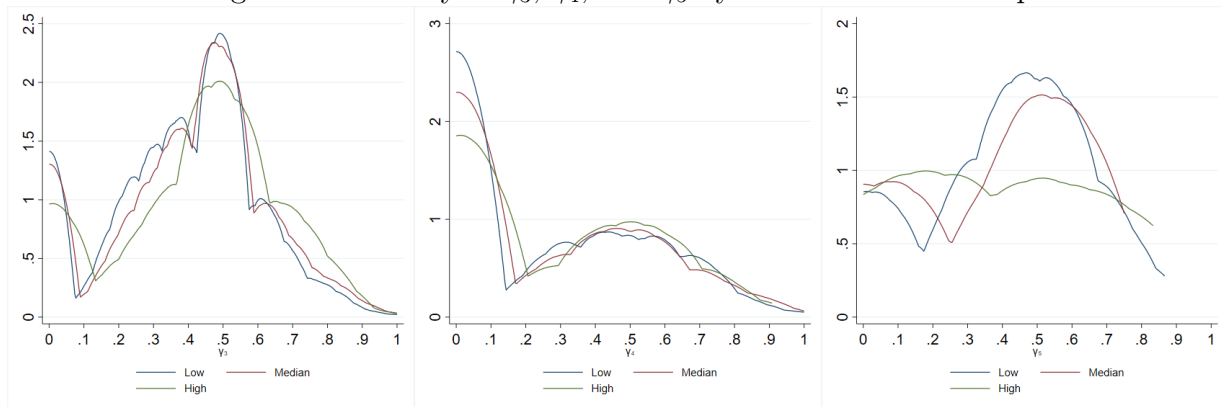
In this paper, we have shown that mediators have statistically significant and economically large effects on bargaining outcomes. Quantitatively, we find that a 75th percentile intermediary

Figure 10: Histogram of γ_3 , γ_4 , and γ_5 Across Threads



Notes: Histograms of γ_3 , γ_4 , and γ_5 across threads. Each data point is a thread.

Figure 11: Density of γ_3 , γ_4 , and γ_5 by Mediator Skill Group



Notes: Distribution of γ_3 , γ_4 , γ_5 for the bottom, middle, and top terciles of mediators, ranked by agreement probability fixed effects. Each data point is a thread.

is 23.23% more likely to close a deal than a 25th percentile intermediary. Our estimated mediator effects are robust to using a variety of different regression specifications. We have found that more experienced mediators tend to perform better, and that female and male mediators perform equally well at closing a deal, but they do so differently: female mediators end negotiations more quickly and at prices more favorable to buyers than to male mediators. Good mediators care more about long-term company incentives and they are better at dealing with harder situations, i.e., threads where the gains from trade are lower or that are less likely to reach agreement ex-ante. 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. We also found that better mediators appear less reliant on split-the-difference behavior of bargaining parties. Our analysis relied on large-scale field data on business-to-business transactions that is new to the literature.

References

- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives*, 11(1):109–126.
- Backus, M., Blake, T., Larsen, B., and Tadelis, S. (2018). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. Forthcoming, *Quarterly Journal of Economics*.
- Bagwell, K., Staiger, R. W., and Yurukoglu, A. (2015). Multilateral trade bargaining: A first look at the GATT bargaining records. NBER Working Paper 21488.
- Basak, D. (2015). Transparency and delay in bargaining. Working paper, New York University.
- 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.
- 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.
- Eisenkopf, G. and Bachtiger, A. (2013). Mediation and conflict prevention. *Journal of Conflict Resolution*, 57(4):570–597.
- Exley, C. L., Niederle, M., and Vesterlund, L. (2016). Knowing when to ask: The cost of leaning-in. Forthcoming, *Journal of Political Economy*.
- Fanning, J. (2019). Mediation in reputational bargaining. Working Paper, Brown University.
- Gilbukh, S. and Goldsmith-Pinkham, P. (2018). 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., Horner, J., Pavlov, G., and Squintani, F. (2009). Mediation, arbitration and negotiation. *Journal of Economic Theory*, 144(4):1397–1420.
- Hernandez-Arenaz, I. and Iriberry, N. (2018). Women ask for less (only from men): Evidence from bargaining in the field. *Journal of Economic Behavior and Organization*, 152:192–214.
- Hernandez-Arenaz, I. and Iriberry, N. (2019). Gender differences in alternating-offer bargaining: An experimental study. Working paper, University of the Balearic Islands.
- Horner, J., Morelli, M., and Squintani, F. (2018). Mediation and peace. Working paper, Yale University.
- Jindal, P. and Newberry, P. (2019). Within and across salesperson heterogeneity in negotiated prices. *Available at SSRN 3001846*.
- 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.
- 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.
- Larsen, B. (2019). The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions. Forthcoming, *Review of Economic Studies*.
- Larsen, B. and Zhang, A. L. (2018). A mechanism design approach to identification and estimation. NBER Working Paper 24837.
- Leibbrandt, A. and List, J. A. (2015). Do women avoid salary negotiations? Evidence from a large-scale natural field experiment. *Management Science*, 61(9):2016–2024.
- Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.
- 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.
- 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.
- 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.