NBER WORKING PAPER SERIES

INTERMEDIARIES IN DECENTRALIZED MARKETS:
EVIDENCE FROM USED-CAR TRANSACTIONS

Fei Li
Charles Murry
Can Tian
Yiyi Zhou

Working Paper 34136
http://www.nber.org/papers/w34136

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025

We have no relevant financial relationships or conflicts of interest to report. We thank Simon Anderson, Gaurab Aryal, Federico Ciliberto, Marco Duarte, Luca Flabbi, Christopher Flinn, Kenneth Gillingham, Alexander MacKay, Brian McManus, Espen R. Moen, Kevin Remmy, Tobias Salz, Katja Seim, Steven Stern, Valentin Verdier, Randall Wright, Jonathan Williams, audiences in 2025 Cowles Summer Conference on Models & Measurement, 25th CEPR-JIE Conference on Applied Industrial Organization, and seminar participants at Collegio Carlo Alberto, Fudan, RUC, StonyBrook, Tsinghua, UNC, and UVA for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Intermediaries in Decentralized Markets: Evidence from Used-Car Transactions
Fei Li, Charles Murry, Can Tian, and Yiyi Zhou
NBER Working Paper No. 34136
August 2025
JEL No. D83, L15, L62

ABSTRACT

We develop and estimate a spatial search-and-bargaining model to study the role of intermediaries and spatial frictions using data from the used-car market. We find that dealers earn price premiums by leveraging three key advantages: selective acquisition of higher-quality cars, superior matching efficiency, and greater bargaining power. Counterfactual simulations reveal that selective intermediation and spatial segmentation significantly affect market efficiency and consumer welfare. While dealers extract more surplus per transaction than sellers, policies reducing dealers' advantages increase search frictions and lower overall welfare. Counterintuitively, reducing spatial frictions harms consumers by shifting trade to less efficient private-seller channels.

Fei Li
University of North Carolina Chapel Hill
Department of Economics
lifei@email.unc.edu

Charles Murry
University of Michigan
Department of Economics
and NBER
ctmurry@umich.edu

Can Tian
University of North Carolina at Chapel Hill
Department of Economics
tiancan@email.unc.edu

Yiyi Zhou
Stony Brook University
Department of Economics
yiyi.zhou@stonybrook.edu

1 Introduction

The presence and role of intermediaries remain a central but controversial topic in economics. In decentralized markets, agents must search for trading partners, and once matched, they must bargain over the terms of trade. As a consequence, intermediaries naturally emerge to facilitate exchanges (Demsetz, 1968) by reducing search frictions, selecting products, and providing liquidity. Meanwhile, they may introduce distortions by capturing surplus, exercising market power, or crowding out direct trade. Public skepticism toward intermediaries is widespread, fueling calls for disintermediation with promises of lowered consumer price and enhanced market efficiency. In a recent book, Judge (2022) highlights that the rise of intermediaries creates both value and risks for the economy. To fully understand the welfare implications of intermediation, it is essential to analyze how realistic complexities such as product heterogeneity, dealer heterogeneity, and spatial friction shape market outcomes.

In this paper, we develop and estimate an equilibrium search model using detailed transaction-level data from the used-car market to examine when and how intermediaries enhance or diminish market efficiency and consumer welfare. The model incorporates two key features of the used car market. First, it captures fundamental trade-offs of intermediated trade. Dealers match with buyers more quickly than private sellers. They are motivated to trade high-quality cars, making their faster matches even more valuable. In the meantime, dealers have greater bargaining power and extract a larger share of the surplus, and they incur inventory-holding costs. In addition, dealers crowd out direct trade between private sellers and buyers, reducing the frequency of direct trade. Second, the model incorporates spatial friction, which is critical in this market. Buyers, dealers, and sellers are geographically dispersed, and search frictions of two trading parties increase in their distance. Within-location trade is easier than across-location trade, both for intermediated trade and direct trade. As a result, areas with higher population density face lower effective frictions, leading to spatial variation in market tightness, with peripheral or sparsely populated regions experiencing greater friction and weaker trading outcomes.

We model search and bargaining in the tradition of Rubinstein and Wolinsky (1987), Duffie et al. (2005) and Gavazza (2016). To capture the complexities of the used car market, we incorporate unobservable car quality, (vertical and horizontal) product differentiation, intermediary heterogeneity, and geographic frictions. We consider a single market with multiple locations. Each location contains sellers (current used car owners), buyers, and dealers. Each buyer has a unit demand for cars, incurs a cost of searching, randomly meets sellers/dealers from all locations, and bargains with them over the transaction price when trade occurs. There are multiple types of dealers, each with type-specific costs, bargaining ability, and matching technology. Dealers have a fixed inventory slot and must acquire cars from sellers. We solve the model for a steady-state distribution of agents.

The richness of the model leads to novel economics that help capture the complex welfare implications of intermediaries/dealers. Although dealers in the model help facilitate trade, their presence generates an equilibrium impact on decentralized trade between sellers and buyers through two mechanisms. First, it affects the two parties' outside options in their bilateral bargaining since they can always forgo the current transaction opportunity and look for trading partners in the future. Holding the product constant, to what extent the direct transaction price is affected due to the presence of dealers depends on the equilibrium impacts of dealers on the two parties' bargaining positions. Second, capacity constraints lead dealers to prioritize products with higher markups to better utilize their limited slots. This selection generates a different composition of products for dealers compared to individual sellers, which further leads to a complex equilibrium impact on price, allocation, and welfare. All these mechanisms exist in one framework that incorporates spatial frictions, where dealers are typically located in the center and are better at matching with buyers across space than individual sellers.

We bring our model to data to quantify the welfare implications of used car dealers. The primary data come from the Ohio Bureau of Motor Vehicles and contain detailed used car transactions, both directly between individual sellers and buyers and between dealers and buyers. We also supplement with dealer listing data obtained from a large online consumer-facing car-selling platform. We focus on a single metropolitan area, where there exists a large used-car chain store (CarMax), many franchised dealers that are mostly located in the center and specialize in younger vintages, as well as a large number of independent dealers that are dispersed in the area and sell cars of all vintages. In addition to confirming the data patterns on dealer price premium and car quality/vintage documented in (Biglaiser et al., 2020), we provide novel spatial features of used car transactions, that is, dealers' price premium and quality selection vary over geographic locations.

We empirically match the key moments of the model, including transaction price, time on market, trade volume, and the composition of transactions across car types, seller types, and geographic locations. Our estimates suggest that search frictions are not trivial in the market. On average, buyers incur more than \$200 and private sellers incur approximately \$400 to complete transactions. The existence of such substantial frictions helps to understand the roles of dealers in this market who facilitate trade, select higher-valued cars to trade, and command a higher price than private sellers. Specifically, our estimates indicate that compared to private sellers, high-type dealers (franchised dealers and CarMax) can match buyers three to six times more efficiently, sell two to five times more high-quality cars, and extract 20% to 50% more surplus from trade on average. Last but not least, our estimates suggest that dealers are especially more efficient in matching with distant buyers, compared to private sellers. Geographic frictions play an important role in determining market outcomes, and dealers' superior ability in reaching distant buyers is a crucial factor for understanding the role of intermediaries.

We perform two sets of counterfactual simulations. In the first set of counterfactuals, we quantify the total welfare implications of used-car dealers. Specially, we reduce the capacity of dealers of all types in the same proportion. As a response, trade shifts from the more efficient intermediated channel to the less efficient direct channel. Consequently, trade speed slows down, trade prices rise, buyers are worse off, and private sellers and remaining dealers are better off. According to our simulation results, the quantitative impact is sizable. For example, a 10% reduction in dealer capacity results in a \$1,432 welfare loss for an average buyer, a \$1,554 welfare gain for an average seller and a \$3,943 value gain for an average remaining dealer slot.

Furthermore, we disentangle the effects of dealers by examining the contribution of their advantages over private sellers separately: greater bargaining ability and superior matching technology. We find that undermining dealers' advantage on either dimension often discourages them to participate and shifts trade from the more efficient intermediation channel to the less efficient direct channel. This reallocation changes the relative steady-state mass of sellers and dealers, extends agents' search time, raises price, and ultimately redistributes welfare among buyers and sellers. The quantitative effect of undermining each advantage is substantial. For example, when dealers are forced to have the same bargaining power as private sellers, the equilibrium number of occupied dealer slots decreases by 21%. This significant reduction in dealer participation results in \$2,043 welfare loss for an average buyer, \$2,923 welfare gain for an average seller, and \$2,581 value loss for an average dealer slot.

The second set of counterfactual exercises quantifies the importance of spatial search frictions in determining the equilibrium outcomes. Specifically, we narrow the spatial discrepancy in trading friction by increasing the efficiency of matching across locations. Lowering trading frictions implies more intensified competition, driving down the retail price and benefiting buyers. However, when sellers' matching efficiency is improved, they are more likely to bypass dealers and sell directly to buyers. As trade shifts from the more efficient intermediary channel to the less efficient direct channel, the retail price can go up, and buyers become worse off. Our simulation results suggest that improving the matching efficiencies of retail trade between locations of all types with that of counterparts within locations by 20%, the retail price increases by \$1,832 for intermediated trade and \$1,113 for direct trade. As a result, on average, buyers lose

$1,334, sellers gain $1,121, occupied dealers slightly gain while unoccupied dealers lose $1,707.

This result stands in stark contrast to conventional wisdom (see, e.g., Burdett and Judd (1983)) in the consumer search literature, which holds that reducing search frictions intensifies competition and lowers prices. In our model, competition manifests itself in agents' equilibrium continuation values: fiercer competition elevates buyers' continuation values and strengthens their bargaining position. When multiple trading channels of different efficiency levels coexist, such as the more efficient intermediated channel and the less efficient direct channel in our context, the reduction of search frictions may also shift the trading from the efficient channel to the less efficient channel, which countervail the competition effect. This analysis highlights the necessity to distinguish *search frictions in the environment* and *search frictions in the equilibrium* when multiple trading channels coexist. In addition, our results suggest that the role of dealers as intermediaries is crucially dependent on the spatial characteristics of their advantage in matching compared to private sellers. Our estimates suggest that dealers' visibility is less restricted by geography, allowing them to trade with distant buyers more efficiently than private sellers. When it becomes easier for private sellers to trade with distant buyers, misallocation due to shifting trade away from efficient intermediaries tends to dominate the competition effect and eventually leads to higher price and harm buyers.

Literature Our paper develops an empirical framework where intermediaries facilitate trade, select high-quality products, and extract more surplus. The framework extends Duffie et al. (2005) and Gavazza (2016) by adding rich heterogeneities to products, intermediaries, and geography to understand how these heterogeneities interact and shape market outcomes. We bridge two strands of literature on intermediaries. The first strand, initiated by Rubinstein and Wolinsky (1987), theoretically emphasizes how intermediaries reduce trading frictions by helping agents identify more trading opportunities. Recent empirical studies have confirmed the role of intermediaries in reducing trading frictions in various markets, including (Gavazza, 2016) on used aircrafts, (Salz, 2022) on waste disposals, (Allen et al., 2014, 2023) on mortgages, and (Brancaccio et al., 2017; Brancaccio and Kang, 2022) on municipal bonds. The other strand, pioneered by Biglaiser (1993), highlights the role of intermediaries in the selective trading of high-quality products. Biglaiser et al. (2020) provides empirical evidence using used car transaction data. Gavazza and Lizzeri (2021) provides a comprehensive review of these two strands of literature.

Closest to our paper is Gavazza (2016), who develops and estimates an empirical search-and-bargaining model to study how intermediaries affect asset allocation, prices, and welfare in decentralized markets using business-aircraft data. While Gavazza (2016) quantifies the overall welfare effects of intermediaries, our paper addresses a distinct and fundamental question:

How do realistic features of used-car market, including (vertical and horizontal) product differentiation, spatial frictions, and dealer heterogeneity, jointly shape the role of intermediaries? We explicitly model and quantify these interactions and reveal novel channels through which intermediaries affect market efficiency and consumer outcomes. Our results show that intermediaries' product selection and spatial market segmentation play key roles in shaping consumer welfare and market efficiency.

This paper joins the growing literature on used car markets and the roles of dealers. Murry and Schneider (2016) provides a survey of early works on car markets. Gavazza et al. (2014) and Gillingham et al. (2022) study the interaction between new-car and used-car markets in a setting with transaction cost. Larsen and Zhang (2021) finds that the type of seller is an important factor in determining their bargaining position in the wholesale used car markets.

Also related is the literature on decentralized transport markets, which integrates search and matching frictions in an economy with multiple locations. To our knowledge, Lagos (2000) is the first to propose such a spatial search theoretical framework. Recent papers apply and extend the framework to empirically study the market for New York City taxis (see, e.g., Lagos (2003), Frechette et al. (2019), and Buchholz (2022)) and the endogenous trade costs in the bulk shipping industry (see, e.g., Brancaccio et al. (2020, 2023).)

Organization The rest of the paper is organized as follows. Section 2 presents the data and key empirical facts. Section 3 develops the equilibrium model. Section 4 describes the empirical specifications, discusses the identification, reports the estimation results, and assesses the fit of the model. Section 5 performs a series of counterfactual analyses to quantify the roles of intermediaries in the used car market. Section 6 concludes. Additional Figures and Tables are relegated to the Appendix.

2 Used-Car Market and Data

2.1 Used-Car Market

The used car market is a suitable setting to study the role of intermediaries due to its unique features.[1] First, used cars are heterogeneous, and the market is highly frictional. For example, it typically takes more than a month to sell a used car (see our data section). Prices

[1] Our general understanding of the industry is based on conversations with dealers and various industry reports, including Edmunds' "Used Vehicle Market Report," Manheim's "Used Car Market Report," and Murry and Schneider (2016). For industry reports, see https://dealers.edmunds.com/static/assets/articles/2017_Feb_Used_Market_Report.pdf and https://publish.manheim.com/content/dam/consulting/2017-Manheim-Used-Car-Market-Report.pdf

for most transactions are determined through bilateral bargaining. Second, dealers are very active participants in the used car market. In the United States, about two thirds of used car sales are made through dealers, and the other one-third occur directly between individual sellers and buyers. The coexistence of direct and intermediated transactions helps us examine their differences and empirically understand the roles of intermediaries. Third, the distributions of car dealers and transactions are highly uneven in geography. Finally, the used car market is large, with total retail sales of more than 500 billion dollars annually in the United States.[2] In 2016, 38.5 million used cars were sold in the United States, more than twice the number of new cars sold.

2.2 Data Description

DMV Transaction Data Our primary source of transaction data includes all administrative transaction records for used cars in the state of Ohio in 2017. This data set is obtained from the Ohio Bureau of Motor Vehicles (henceforth "DMV transaction data"). For each transaction record, we observe the transaction date and price, the buyer's zip code, the seller's identity if it is a dealer or the seller's zip code if it is a private seller, and the basic information of the traded car including its mileage, make, model, model year, and Vehicle Identification Number (VIN) which is a unique identifier of the car.

Dealer Listing Data We obtain dealer listing data from `cars.com`, one of the largest online consumer-facing platforms for cars. For each listing, we observe the date the car was initially listed, the date it was taken off the platform, the list price, the identity of the dealer, and the basic information of the listed car, including its mileage, make, model, model year, and VIN.

According to our conversations with `cars.com`, dealers typically pay a fixed fee for using the platform, and the variable cost for dealers to list an additional car is close to zero. Hence, it is reasonable to assume that they list their entire inventory as long as they have paid the fixed fee. We use these data to measure dealers' inventory of used cars. In addition, because our data from `cars.com` is at the daily level, we effectively see when a car is listed initially and when it is removed from the platform. We use this information to measure how long a car has been on the market.

Sample Selection To make our study tractable, we focus our attention on a single geographical area, the Columbus Census Metropolitan Statistical Area (MSA). This metropolitan area is in central Ohio, so it does not share borders with other states where we do not have data.

[2]This number, constructed from Edmunds' and Manheim's yearly reports, represents revenues from franchised and independent dealers only, so it is a conservative reflection of the size of the industry.

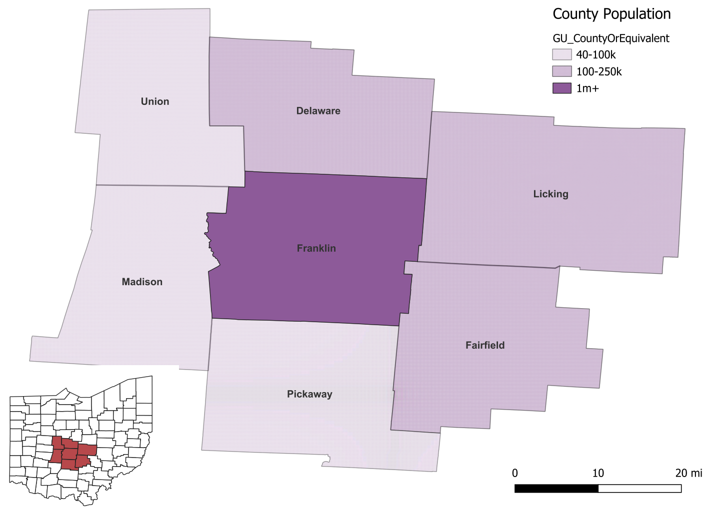

Figure 1: Seven Counties in Columbus Area, Ohio

This MSA has the typical layout of metropolitan areas in the U.S., with an urban core city, a suburban ring, and outlying rural areas. The area we study includes seven counties: Franklin County (the core county that includes the city of Columbus) and six less populated outlying counties[3] The geographic area we study has a distance from the core county to each surrounding county of no more than 30 miles, and the counties farther from the core are sparsely populated. See Figure 1 for the map of the seven counties and their population.

We focus on a specific class of cars, the non-luxury midsize sedan with a vintage of 4-13 years old.[4] We exclude cars older than 14 years because trades of those extremely old cars are few and occur mostly between buyers and private sellers. We exclude cars aged three years or younger mainly for two reasons. First, it can alleviate the concern that the transaction price could be confounded by those non-vehicle components such as Certified Pre-Owned (CPO) or trade-in discounts. Nevertheless, cars older than three years are usually not covered by CPO and only a small proportion of buyers of these relatively old cars have trade-in cars.[5] The other reason why we exclude cars aged three years or younger is that a large proportion of these young cars are off-lease cars. Empirically, the inventory of franchised dealers spikes for cars aged two or three, which is the typical length of leasing contracts. However, this pattern is not the one that this paper is intended to explain.

[3]The city of Columbus has a population of about 900,000 and is the 14th most populous city in the U.S..

[4]Non-luxury midsize sedan models include Chevrolet Malibu, Dodge Avenger, Ford Taurus, Honda Accord, Hyundai Sonata, Kia Optima, Mazda MX-6, Nissan Altima, Subaru Legacy, Toyota Camry, and Volkswagen Passat.

[5]https://www.edmunds.com/car-news/used-car-prices-2025.html

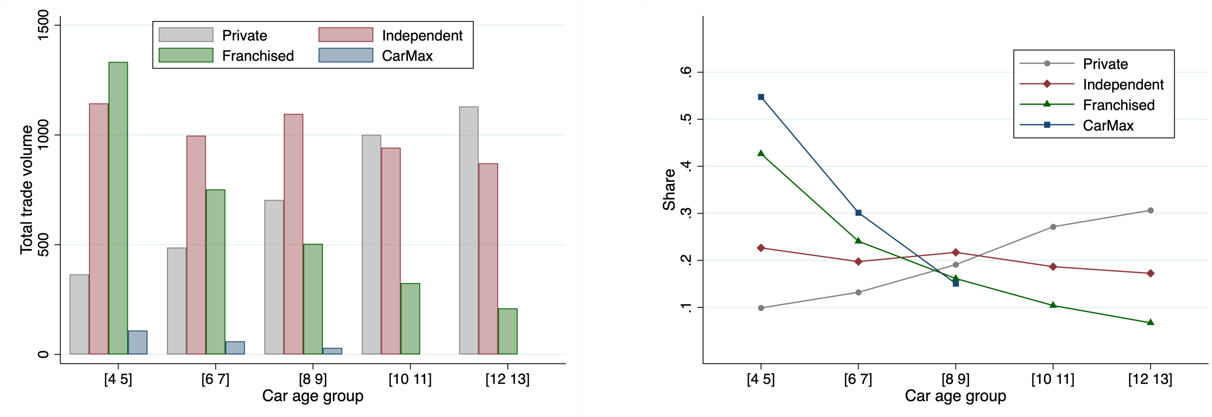

(a) Volume

(b) Fractions by Car Vintage

Figure 2: Transaction Volume

In the end, we are left with 14,785 transactions that satisfy the geographic boundaries and the product restrictions discussed above. In total, 11,954 cars of this particular class are listed by dealers located in the Columbus MSA throughout the year of 2017. We classify all dealers into three types: (i) independent dealers that sell used cars only, (ii) franchised dealers that sell both new cars and used cars, and (ii) CarMax, which is the largest national chain store of used cars.[6] There is one CarMax store in this area, located in Franklin (core county).

2.3 Motivating Facts

In this section, we document several empirical facts that motivate our focus on used car dealers, support our modeling decisions, and motivate our counterfactual analyses.[7]

Fact 1: Dealers predominantly trade younger cars than those exchanged in private transactions.

Figure 2a presents the total volume of transactions sold by private owners, independent dealers, franchised dealers and CarMax in different vintages of cars, and Figure 2b presents the share of each vintage of cars between all cars transacted for each type of seller. Overall,

[6] We do not include online used car dealers such as Carvana and Vroom in our analysis for two reasons. First, the market share of online dealers in 2017 is less than 0.5% in our DMV transaction data. Second, the process of trading with online dealers is quite different from that of trading with private sellers or conventional physical dealers.

[7] These empirical patterns are similar for the full sample that includes all used car models with all vintage, and the results are upon request.

Table 1: Transaction Price and Weeks on Market

	Private Seller		Independent Dealer		Franchised Dealer		CarMax	
Car Age (years)	price ($1,000)	weeks on market	price ($1,000)	weeks on market	price ($1,000)	weeks on market	price ($1,000)	weeks on market
4-5	9.028	-	11.872	6.621	13.367	4.774	15.724	3.144
6-7	6.199	-	9.250	6.071	10.191	4.170	12.421	2.986
8-9	4.455	-	6.767	6.455	8.230	3.547	10.998	2.581
10-11	2.945	-	4.656	5.450	6.293	2.962	-	-
12-13	2.093	-	3.098	5.372	4.910	2.663	-	-

cars sold by dealers are substantially younger than those sold directly by private owners. In addition, the volume of direct transactions increases with age of the car, whereas the volumes of transaction of dealer cars decrease with age of the car.

Fact 2: The average price in intermediated transactions is substantially higher than in direct transactions, with this price premium declining as car age increases.

Table 1 reports the average transaction price for each age group, comparing transactions by private sellers and three types of dealers. Across all age groups and dealer types, dealer prices are higher than private seller prices. Independent dealers have the lowest price premium while CarMax has the highest premium. In addition, the dealer price premium decreases in car age for all dealer types.

The impact of heterogeneous characteristics of cars, such as car models, on the price gap is limited. To tease out these impacts, we estimate a hedonic price regression on a rich set of variables related to transaction, including indicators for independent dealer, franchised dealer, and CarMax, indicators of different car age groups, log of mileage, monthly fixed effects, car model fixed effects, and seller county - buyer county pair fixed effects. The coefficient before the indicator of a particular dealer type captures the average price premium of that dealer type. We find that the average price premium for each dealer type is at a similar level. The estimation results are reported in the Appendix.

One concern is that the prices of direct transactions may be underreported for the sake of paying a lower transfer tax, which is around 7% of the gross price in this region. To alleviate this concern, we perform two sanity checks. First, Biglaiser et al. (2020) reports a similar dealer price premium using Virginia DMV data where cars are taxed as property, and hence under-reporting for tax avoidance is less of a concern. In addition, we manually check the market price of cars sold by dealers and private owners on *Kelley Blue Book* (KBB) for the car models and car vintages we focus on and find that the price differences are of similar magnitude as ours.

Another concern is that dealer sales may come along with buyers' trade-ins or valuable manufacturer warranties, potentially confounding the observed price recorded by the DMV. As discussed earlier, limiting the sample to vehicles at least three years old mitigates these issues, making it unlikely that they have a meaningful impact on our results.

Fact 3: Time on the market varies by car age and dealer type. Older cars sell faster than younger cars. CarMax has the shortest selling time, followed by franchised dealers, then independent dealers.

Table 1 also reports the average time on the market of cars listed by dealers for each age group of cars and each type of dealer. First, dealer cars stay on market for a few weeks on average, suggesting the existence of nontrivial search frictions in the market. Second, cars listed by high-type dealers (CarMax and franchised dealers) sell substantially faster than those listed by independent dealers. For example, the average selling time of CarMax is three quarters of that of franchised dealers and only half that of independent dealers. Lastly, older cars stay on the market shorter than younger cars.[8]

The model we present in Section 3 can reconcile these rich time-on-market patterns in the data by incorporating seller heterogeneity and product differentiation, specifically the empirical fact that high-type dealers sell more quickly at higher prices than low-type dealers, while young cars sell more slowly at higher prices than old cars.

Fact 4: Used car transactions exhibit a core-periphery spatial pattern.

We group all transactions into four groups based on whether the seller is from the core county or any of the periphery counties, and whether the buyer is from the core county or any of the periphery counties. Table 2 reports the volume of the trading, the average price of the transaction, the average age of the vehicle transacted, and the share of direct transactions for each group of transactions. Clearly, used car transactions are spatially unbalanced and exhibit a *core-periphery* pattern. About 80% of all transactions originate in the core county, of which 85% are within-county trades, while the remaining 15% flow to peripheral counties.

Among the cars sold from the six periphery counties, 29% of them flow into the core county.[9] In addition, cars flowing out of the core county are younger and more expensive on average, and more likely to be intermediated than cars flowing into the core county. The sharp difference

[8]The only exception is that the time on market of 8-9 years old cars is longer than 6-7 years old cars for independent dealers.

[9]This pattern remains even if we focus on direct transactions between buyers and sellers.

Table 2: Geographic Distribution of Transactions

	Trade Volume	Price ($1,000)	Car Age (years)	Share of Direct Transactions
Core → core	9,982	7.534	8.335	0.403
		(4.684)	(2.918)	
Core → periphery	1,714	8.522	7.883	0.343
		(4.620)	(2.867)	
Periphery → core	881	7.161	8.508	0.616
		(4.466)	(2.923)	
Periphery → periphery	2,208	7.561	8.335	0.568
		(4.354)	(2.936)	
All	14,785	7.645	8.311	0.433
		(4.637)	(2.919)	

Note: The sample includes all transactions of selected used car models in 2017. The sample selection is described in the text. Standard deviations are reported in parenthesis.

between core-to-periphery and periphery-to-core trade volume, price, age, and the share of direct transactions suggests that the transaction network is directed. Moreover, as we argue in the estimation section, this asymmetry is an important source for separately identifying different economic forces.

Fact 5: Cars listed by dealers in the core area sell faster than those listed by dealers in peripheral areas.

Table 3 reports the dealer capacity and the average weeks on the market of used cars listed by dealers in the core area and by dealers in the periphery. First, dealers in the core area have more slots. Second, cars listed by dealers located in the core area stay on the market shorter.

Fact 4 and **Fact 5** highlight the importance of spatial heterogeneity in the used car market. The fact that location is important is not surprising given the nature of used car transactions. Unlike new cars, used cars differ substantially in condition, history, and maintenance, potentially increasing the necessity for in-person inspections. As a result, location heterogeneity crucially influences who trades with whom, at what price, and how fast those trades occur.

3 Model

In this section, we present a spatial search and bargaining model to investigate the roles of intermediaries in decentralized markets. We build on canonical models of frictional inter-

Table 3: Capacity and Weeks on Market of Dealer Cars

County	Independent Dealers		Franchised Dealers		CarMax	
	capacity	weeks on market	capacity	weeks on market	capacity	weeks on market
Core	591	5.951	223	4.009	18	3.018
Periphery	261	6.901	120	4.430	-	-

Note: Inventory is the average number of listings of the car class that we focus on in this study by each dealer type in each county average over all weeks in 2017. Capacity is the maximal number of weekly inventory by each dealer type in each county over all weeks in 2017.

mediation, but allows for rich product differentiation, heterogeneity among intermediaries, and spatially distributed agents, rationalizing the empirical facts that we present in Section 2.3.

3.1 Environment

Locations and agents We consider a stationary economy with multiple locations $l \in L$, where L is finite. Time is continuous and lasts forever. Each location is populated by three groups of agents: buyers, dealers, and sellers. They are risk-neutral and discount the future at a rate of r. They trade heterogeneous cars with multidimensional *characteristics*, denoted as $x \in X$. The characteristic of the car summarizes its payoff-relevant features, including those that econometricians do not observe. It captures the *vertical* product differentiation among cars. In the remainder of the paper, we use x-car to denote a car with characteristics x. We build the model toward steady-state equilibrium. We suppress the time index to ease the notation.

Each location has an endogenous continuum of sellers, each of whom owns a car. At each instant, there is a constant flow of sellers with x-cars into each location l, denoted as $\nu_S(x, l)$, $\forall x, l$. A seller's instantaneous payoff for owning the car is normalized to be zero. Whenever he sells his car, he leaves the economy. Until leaving the economy, every seller at the location l faces a flow search cost of $\kappa_s(l)$. Let $\mu_S(x, l)$ denote the mass of location l's sellers with x-cars at steady-state equilibrium.

Each location also has an endogenous continuum of buyers, each of which has a unit demand for cars. At each instant, a constant flow $\nu_B(l)$ of buyers arrives at location l, $\forall l$. If a buyer purchases an x-car at price p, she receives a lump-sum payoff $u(x) - p + \epsilon$, and leaves the economy, where the deterministic component $u(x)$ depends on the car's characteristic x, and the random component, $\epsilon \overset{i.i.d.}{\sim} \text{Logistic}(0, \sigma(x))$, captures the idiosyncratic utility shock or *horizontal* product differentiation. The distribution assumption ensures that (i) a buyer and a seller trade with logistic probability once they meet, and (ii) both the mean and the dispersion of transaction prices can depend on the characteristics of the car. Until leaving the economy, every buyer at the location l faces a flow search cost of $\kappa_b(l)$. Let $\mu_B(l)$ denote the steady-state equilibrium mass of the buyers in location l'.

There are multiple types of dealers, and there is a fixed mass $m(d,l)$ of type-d dealers at location l. Let $\mathcal{D}$ denote the finite set of all dealer types. Dealers are permanently present. Their search cost is normalized to be zero. Each dealer has one slot and can hold *at most one unit of inventory.*[10] Therefore, a dealer slot is either occupied (when the inventory is 1) or vacant (when the inventory is 0). An occupied dealer slot incurs a flow inventory cost $c(d,l)$, which depends on the dealer's type and location. We use $\mu_0(d,l)$ and $\mu_1(x,d,l)$ to denote the respective steady-state equilibrium masses of vacant and occupied dealers of type d at location l with x-cars. Obviously, $m(d,l) = \mu_0(d,l) + \sum_x \mu_1(x,d,l)$. When a dealer sells a car at price p, it receives a lump-sum payoff p, and when a dealer buys a car at price p, it receives a lump-sum payoff $-p+\varepsilon$, where $\varepsilon \overset{i.i.d.}{\sim} \text{Logistic}(0, \sigma_w(x))$. This assumption ensures that a dealer and a seller trade with a logistic probability once they meet. Dealers maximize their expected discounted lifetime profits.

In our setting, cars ultimately enter the economy with sellers' inflow. Therefore, sellers and dealers compete for buyers' scarce demand in the retail market; whereas buyers and dealers also compete for sellers' limited supply in the wholesale market.

Matching and trade We explicitly model search frictions in the decentralized market using a matching function as in Duffie et al. (2005). Agents, within their locations or across locations, can meet bilaterally, and when they meet, they decide whether to trade.

Specifically, two matching technologies provide opportunities for buyers to pursue a purchase, either directly or via dealers. First, buyers and sellers contact each other directly and randomly. A seller at location l_S with a car of characteristic x and a buyer at location l_B meet *pairwise independently* at a Poisson rate. The total mass of meetings between the two groups is given by

$$\lambda_{SB}(l_S, l_B)\mu_S(x, l_S)\mu_B(l_B), \tag{1}$$

where $\lambda_{SB}(l_S, l_B)$ is the matching coefficient, and the subscript SB represents a direct trade between a seller (S) and a buyer (B). The expression (1) imposes an upper bound on the instantaneous transaction volume. Notice that the population mass of each group captures the effect of the "market size" and the coefficient $\lambda_{SB}(l_S, l_B)$ captures the geographic impact on matching, such as the distance between locations l_S and l_B. The trade flow has a direction,

[10] In section 3.3, we provide a discussion of this assumption.

i.e., for $l \neq l'$, it is possible that $\lambda_{SB}(l,l')\mu_S(x,l)\mu_B(l') \neq \lambda_{SB}(l',l)\mu_S(x,l')\mu_B(l)$. The mass of meetings is uniformly distributed within each group, so a location l_B buyer meets a location l_s seller with an x-car at a rate $\lambda_{SB}(l_S,l_B)\mu_S(x,l_S)$, and a location l_S seller meets a location l_B buyer at a rate $\lambda_{SB}(l_S,l_B)\mu_B(l_B)$.

Second, trade takes place through intermediation through similar matching technologies. At each instant, the total mass of meetings between location l_D type d dealers with an x-car and location l_B buyers is given by $\lambda_{DB}(d,l_D,l_B)\mu_1(x,d,l_D)\mu_B(l_B)$, and the total mass of contact between location l_S sellers with x-cars and location l_D vacant dealers is $\lambda_{SD}(d,l_S,l_D)\mu_S(x,l_S)\mu_0(d,l_D)$. These meetings are also allocated uniformly within a group.

Two agents in a meeting share symmetric information and observe each other's locations, the car's characteristics, and the random payoff shock ϵ.[11] They trade if and only if the joint surplus of two agents is positive. If they trade, we use the generalized Nash bargaining solution to pin down the transaction price. In a buyer-seller meeting, the seller's bargaining power is $\theta_{SB} \in [0,1]$; in a dealer-buyer (or dealer-seller) meeting, the type d dealer's *bargaining power* is $\theta_{DB}(d) \in [0,1]$ (or $\theta_{SD}(d)$). We allow $\theta_{SB} \neq \theta_{DB}(d) \neq \theta_{SD}(d)$ for each d, and we allow $\theta_{DB}(d)$ and $\theta_{SD}(d)$ to vary across d to capture dealers' heterogeneous (dis)advantages in negotiation.

3.2 Steady-State Equilibrium

We look for *steady-state equilibria* where (i) given the masses of buyers $\mu_B(l)$, sellers $\mu_S(x,l)$, occupied dealers $\mu_1(x,d,l)$ for each (x,d,l), each agent maximizes their lifetime discounted payoff, and (ii) for each (x,l,d), $\mu_B(l)$, $\mu_S(x,l)$, and $\mu_1(x,d,l)$ are time-invariant and consistent with agents' optimal choices. In the remainder of this section, we first formulate the transaction outcome of each meeting pair of agents given their steady-state value functions and distributions. We then establish the conditions that agents' value functions and distributions must obey in a steady state.

3.2.1 Pairwise Transaction

Buyer-Seller Trade Suppose that a location l_B buyer and a location l_S seller who owns an x-car meet. The buyer draws a payoff shock $\epsilon \overset{i.i.d.}{\sim} \text{Logistic}(0,\sigma(x))$. The two agents' *joint trade*

[11] While some papers in the literature (e.g., Biglaiser (1993) and Biglaiser and Li (2018)) relies on asymmetric information to generate dealer selection, arguing that dealers certify and selectively sell high-quality cars, in our setting, a similar selection arises endogenously as a consequence of search frictions and dealers' limited inventory capacity. It is difficult to empirically distinguish them, so we maintain a symmetric information assumption for parsimony and tractability. See Section 3.3 for more discussion.

surplus is denoted as

$$\underbrace{u(x) - V_B(l_B) - V_S(x, l_S)}_{S_{SB}(x, l_S, l_B)} + \epsilon, \tag{2}$$

where $V_B(l_B)$ and $V_S(x, l_S)$ represent the buyer's and the seller's endogenous continuation values if the trade fails to occur, and $u(x) + \epsilon$ is their joint payoff if the trade occurs. Hence, they trade if and only if their joint trade surplus is non-negative, that is, $S_{SB}(x, l_S, l_B) + \epsilon \geq 0$. Let $\alpha_{SB}(x, l_S, l_B) = \Pr[S_{SB}(x, l_S, l_B) + \epsilon \geq 0]$ denote the probability that such a trade occurs. Because ϵ is drawn from a logistic distribution, this probability satisfies

$$\alpha_{SB}(x, l_S, l_B) = \frac{\exp\Big(S_{SB}(x, l_S, l_B)/\sigma(x)\Big)}{1 + \exp\Big(S_{SB}(x, l_S, l_B)/\sigma(x)\Big)}. \tag{3}$$

The instantaneous volume of transactions between location l_S sellers with x-cars and location l_B buyers is $\lambda_{SB}(l_S, l_B)\mu_S(x, l_S)\mu_B(l_B)\alpha_{SB}(x, l_S, l_B)$.

Whenever a buyer-seller pair trades, the transaction price is set to split the joint surplus in fixed proportions according to $\theta_{SB} \in [0, 1]$, i.e.,

$$p_{SB}(\epsilon, x, l_S, l_B) - V_S(x, l_S) = \theta_{SB}[S_{SB}(x, l_S, l_B) + \epsilon]. \tag{4}$$

The left-hand side of equation (4) is the seller's surplus, which is the difference between the transaction price and the value of keeping the car and searching for alternative buyers. The right-hand side corresponds to the θ_{SB} fraction of the total surplus. The rest $(1-\theta_{SB})$ fraction of the total surplus goes to the buyer. An important observation is that the *minimum transaction price* given (x, l_S, l_B) is $V_S(x, l_S)$, which is observed when the joint trade surplus is exactly zero. Fixing (x, l_S, l_B), the dispersion of transaction prices is solely driven by the random utility component through the term $\theta_{SB}\epsilon$.

Combining equations (2) and (4) further reveals how the transaction price is determined:

$$p_{SB}(\epsilon, x, l_S, l_B) = (1 - \theta_{SB})V_S(x, l_S) + \theta_{SB}[u(x) - V_B(l_B) + \epsilon], \tag{5}$$

which is a convex combination of the seller's and the buyer's trade surplus. Even when the bargaining power, θ_{SB}, is fixed, an agent's outside option, reflected in the continuation value of searching, still influences their *bargaining position* through its effect on surplus. Throughout the paper, we say that an agent's bargaining position is weakened if their outside option decreases and strengthened if it increases.

An agent's bargaining position is determined in the equilibrium, influenced by competition. For example, retail market competition among selling parties determines the buyer's equilibrium continuation value: When competition intensifies, buyers expect higher payoffs from continued search, which strengthens their bargaining position and results in lower retail prices.

Following McFadden (1977), one can write the expected surplus as

$$\mathbb{E}\Big[\max\big\{S_{SB}(x,l_S,l_B)+\epsilon,0\big\}\Big]=\sigma(x)\ln\Big[1+\exp\Big(S_{SB}(x,l_S,l_B)/\sigma(x)\Big)\Big].$$

It follows that the *expected transaction price* can be written as

$$\begin{aligned} p_{SB}(x,l_S,l_B) \equiv & \mathbb{E}\Big[p_{SB}(\epsilon,x,l_S,l_B)|S_{SB}(x,l_S,l_B)+\epsilon\geq 0\Big] \\ = & V_S(x,l_S)+\theta_{SB}\frac{\sigma(x)\ln\Big[1+\exp\Big(S_{SB}(x,l_S,l_B)/\sigma(x)\Big)\Big]}{\alpha_{SB}(x,l_S,l_B)}, \end{aligned} \tag{6}$$

where $V_S(x,l_S)$ is the seller's opportunity cost of trade (also the minimum transaction price for an x-car between location l_S sellers and location l_B buyers), $\frac{\sigma(x)\ln\left[1+\exp\left(S_{SB}(x,l_S,l_B)/\sigma(x)\right)\right]}{\alpha_{SB}(x,l_S,l_B)}$ is the expected joint surplus of the trade *conditional on trade occurring*, and θ_{SB} is the fraction of the share received by the seller.

Equation (6) suggests an expression for the "mean-min price difference," i.e., the difference between the average transaction price and the lowest transaction price. As Hornstein et al. (2011) has argued, the mean-min difference (or ratio) is an informative measure of frictional price dispersion generated by search models. Specifically, the "mean-min price difference" depends on the seller's bargaining power θ_{SB} and the dispersion of buyer's idiosyncratic utility component $\sigma(x)$.

Dealer-Buyer Trade When a location l_D type d dealer with an x-car and a location l_B buyer meet, they trade if and only if the joint trade surplus is positive, i.e.,

$$\underbrace{u(x)-V_B(l_B)+W_0(d,l_D)-W_1(x,d,l_D)}_{S_{DB}(x,d,l_D,l_B)}+\epsilon\geq 0.$$

where $V_B(l_B)$ is the value of a location l_B buyer who continues to search on the market, $W_0(d,l_D)$ is the value of a location l_D type d vacant dealer, and $W_1(x,d,l_D)$ is the value of a location l_D type d dealer with an x-car. If the trade takes place, the pair receives a joint continuation payoff $u(x)+\epsilon+W_0(d,l_D)$, but if the trade does not occur, the joint continuation payoff will be $V_B(l_B)+W_1(x,d,l_D)$.

The total transactions of an x-car between location l_D type d dealers and location l_B buyers is $\lambda_{DB}(d,l_D,l_B)\mu_1(x,d,l_D)\mu_B(l_B)\alpha_{DB}(x,d,l_D,l_B)$, where $\alpha_{DB}(x,d,l_D,l_B) = \Pr[S_{DB}(x,d,l_D,l_B)+\epsilon \geq 0]$ is the probability of trading between such a pair, satisfying

$$\alpha_{DB}(x,d,l_D,l_B) = \frac{\exp\Big(S_{DB}(x,d,l_D,l_B)/\sigma(x)\Big)}{1+\exp\Big(S_{DB}(x,d,l_D,l_B)/\sigma(x)\Big)}.$$

When a trade occurs, the transaction price is given by

$$p_{DB}(\epsilon,x,d,l_D,l_B) + W_0(d,l_D) - W_1(x,d,l_D) = \theta_{DB}(d)[S_{DB}(x,d,l_D,l_B)+\epsilon].$$

That is, $\theta_{DB}(d) \in [0,1]$ is the fraction of the surplus the dealer receives and $(1-\theta_{DB}(d))$ is the fraction of the surplus the buyer receives. The expected transaction price can be written as

$$\begin{aligned} p_{DB}(x,d,l_D,l_B) \equiv & \mathbb{E}\Big[p_{DB}(\epsilon,x,d,l_D,l_B)|S_{DB}(x,d,l_D,l_B)+\epsilon \geq 0\Big] \\ = & W_1(x,d,l_D) - W_0(d,l_D) + \theta_{DB}(d)\frac{\sigma(x)\ln\Big[1+\exp\Big(S_{DB}(x,d,l_D,l_B)/\sigma(x)\Big)\Big]}{\alpha_{DB}(x,d,l_D,l_B)}, \end{aligned} \tag{7}$$

On the right-hand side of equation (7), the first term $W_1(x,d,l_D) - W_0(d,l_D)$ corresponds to the minimum transaction price for an x-car between type-d dealers at location l_D and buyers at location l_B, and the second term is the surplus extracted by the dealer in transactions.

It is interesting to compare the prices in equations (6) and (7) to understand the difference between dealers' and sellers' positions in negotiating with buyers. Even if the characteristics of the car being traded are identical, the expected transaction price a buyer will pay to a dealer may differ from the one she pays to a seller in the same location, that is, $p_{DB}(\epsilon,x,d,l_D,l_B) \neq p_{SB}(\epsilon,x,l_S,l_B)$ for $l_D = l_S$. The difference is present because sellers and dealers may have different (i) opportunity costs, (ii) bargaining powers, or (iii) surplus from the trade.

Seller-Dealer Trade When a location l_D type-d dealer and location l_S seller with an x-car meet, they trade if and only if their joint surplus is positive, that is,

$$\underbrace{W_1(x,d,l_D) - W_0(d,l_D) - V_s(x,l_S)}_{S_{SD}(x,d,l_S,l_D)} + \varepsilon \geq 0.$$

If the trade takes place, the joint continuation payoff between two agents is $W_1(x,d,l_D)+\varepsilon$; otherwise it is $W_0(d,l_D) + V_s(x,l_S)$.

The total transactions of x-cars between location l_S sellers and location l_D type d dealers can be written as $\lambda_{SD}(d, l_S, l_D)\mu_s(x, l_S)\mu_0(d, l_D)\alpha_{SD}(x, d, l_S, l_D)$, where $\alpha_{SD}(x, d, l_S, l_D) = \Pr(S_{SD}(x, d, l_S, l_D) + \varepsilon \geq 0)$ is the probability of trading between such a pair, satisfying

$$\alpha_{SD}(x, d, l_S, l_D) = \frac{\exp\Big(S_{SD}(x, d, l_S, l_D)/\sigma_w(x)\Big)}{1 + \exp\Big(S_{SD}(x, d, l_S, l_D)/\sigma_w(x)\Big)}. \tag{8}$$

Equation (8) reveals the dealer's selection of the characteristics of the car. An x-car is more likely to be purchased by a dealer if $S_{SD}(x, d, l_S, l_D)$ is higher. Since $S_{SD}(x, d, l_S, l_D)$ varies across car characteristics x, the above inequality essentially generates a car characteristics-based *selection criteria* for dealers. This selection mechanism implies that in a steady-state equilibrium, the characteristics distribution of cars held by dealers can differ substantially from that of cars held by sellers.

The economics behind the product selection warrants further discussion. From the dealer's perspective, purchasing an x-car comes with an opportunity cost, that is, the temporary loss of the option to buy other cars, captured by $W_0(d, l_D)$. This term is independent of x and acts as a hurdle to the transaction. The remaining component of $S_{SD}(x, d, l_S, l_D)$ is $W_1(x, d, l_D) - V_S(x, l_S)$. It represents the change in continuation value when ownership transfers from the seller to the dealer. Ultimately, the value of holding a car for either party is derived from the buyer's utility. Both sellers and dealers find it more beneficial to trade cars that are highly valued by buyers. However, when dealers can match with buyers more efficiently in the retail market (higher $\lambda_{DB}(d)$) and extract more surplus from transactions (higher $\theta_{DB}(d)$), $W_1(x, d, l_D)$ grows more rapidly with $u(x)$ than with $V_S(x, l_S)$. Consequently, as $u(x)$ increases, $S_{SD}(x, d, l_S, l_D)$ increases, making it more likely that the car will be sold to a dealer. Finally, dealer selection criteria also depend on the dealer's bargaining power and inventory costs. Holding agents' continuation value constant, an increase in the dealer's bargaining power allows them to extract more surplus in the retail market, while a decrease in inventory costs lowers the expense of holding cars. In both cases, the dealer finds it more profitable to acquire inventory.

When the trade occurs, the transaction price is given by

$$W_1(x, d, l_D) - W_0(d, l_D) - p_{SD}(\epsilon, x, d, l_S, l_D) = \theta_{SD}(d)[S_{SD}(x, d, l_S, l_D) + \varepsilon].$$

That is, $\theta_{SD} \in [0, 1]$ is the fraction of the surplus the dealer receives and $(1-\theta_{SD})$ fraction of the surplus that the seller receives. The expected price $\mathbb{E}[p_{SD}(\varepsilon, x, d, l_S, l_D)|S_{SD}(x, d, l_S, l_D)+\epsilon \geq 0]$ can be written as a function of $S_{SD}(x, d, l_S, l_D)$ in a way similar to the previous cases.

3.2.2 Agents' Value Functions and Distributions

Seller's Value Function First, consider a location l_S seller who owns an x-car. His value function $V_S(x, l_S)$ obeys the following Hamilton-Jacobi-Bellman (HJB) equation,

$$\begin{aligned} rV_S(x, l_S) = -\kappa_S(l_S) + \sum_{l_B} \lambda_{SB}(l_S, l_B)\mu_B(l_B)\mathbb{E}\Big[\max\big\{\theta_{SB}[S_{SB}(x, l_S, l_B) + \epsilon], 0\big\}\Big] \\ + \sum_{d, l_D} \lambda_{SD}(d, l_S, l_D)\mu_0(x, l_D)\mathbb{E}\Big[\max\big\{(1 - \theta_{SD}(d))[S_{SD}(x, d, l_S, l_D) + \varepsilon], 0\big\}\Big]. \end{aligned} \tag{9}$$

The seller incurs a flow cost $\kappa_S(l_S)$ to search for buyers. During the search, he meets a location l_B buyer at the rate $\lambda_{SB}(l_S, l_B)\mu_B(l_B)$ and chooses whether to trade. If the joint surplus is positive, the trade occurs and the seller gives up the option value $V_s(x, l_S)$, gets a payment $p_{SB}(\epsilon, x, l_S, l_B)$ determined by equation (4), and leaves the market. By expression (4), his surplus is given by $\theta_{SB}[S_{SB}(x, l_S, l_B) + \epsilon]$. If the joint surplus is negative, the trade will not occur and the seller's surplus is zero. The expectation is taken over ϵ. Similarly, at rate $\lambda_{SD}(d, l_S, l_D)\mu_0(x, l_D)$, he meets a type d vacant dealer at location l_D. They trade if and only if the joint surplus is positive. When they trade, the seller's surplus is $(1 - \theta_{SD}(d))$ fraction of the total surplus. The expectation is taken over ε. Also, since our dataset covers only a single year, we do not model product depreciation while the car is held by any agent.

Buyer's Value Function Next, consider a location l_B buyer's value function. Her value function $V_B(l_B)$ obeys

$$\begin{aligned} rV_B(l_B) = &- \kappa_B(l_B) + \sum_{l_S, x} \lambda_{SB}(l_S, l_B)\mu_s(x, l_S)\mathbb{E}\Big[\max\big\{(1 - \theta_{SB})[S_{SB}(x, l_S, l_B) + \epsilon], 0\big\}\Big] \\ &+ \sum_{x, d, l_D} \lambda_{DB}(d, l_D, l_B)\mu_1(x, d, l_D)\mathbb{E}\Big[\max\big\{(1 - \theta_{DB}(d))[S_{DB}(x, d, l_D, l_B) + \epsilon], 0\big\}\Big]. \end{aligned} \tag{10}$$

The buyer incurs a $\kappa_B(l_B)$ flow cost to search for suppliers. During searching, at rate $\lambda_{SB}(l_S, l_B)\mu_S(x, l_S)$, the buyer meets a location l_S seller with an x-car and draws a random utility ϵ. Trade occurs if and only if two agents' joint surplus is positive. The buyer's surplus is equal to $(1 - \theta_{SB})[S_{SB}(x, l_S, l_B) + \epsilon]$ if she trades; otherwise, her surplus is zero. At rate $\lambda_{DB}(d, l_D, l_B)\mu_1(x, d, l_D)$, she meets a type d dealer with inventory x at location l_D and draws a random utility ϵ. Trade takes place if and only if the joint surplus is positive. If she trades, her surplus is given by $(1 - \theta_{DB}(d))[S_{DB}(x, d, l_D, l_B) + \epsilon]$; otherwise, it is zero. The expectation in the two terms on the right-hand side of the HJB is taken over ϵ.

Dealer's Value Function The value function of a type d dealer at location l_D with an x-car must satisfy

$$rW_1(x,d,l_D) = -c(d,l_D) + \sum_{l_B} \lambda_{DB}(d,l_D,l_B)\mu_B(l_B)\mathbb{E}\Big[\max\{\theta_{DB}(d)[S_{DB}(x,d,l_D,l_B)+\epsilon],0\}\Big]. \tag{11}$$

In words, he incurs an inventory cost $c(d,l_D)$ at each instant, meets a location l_B buyer at rate $\lambda_{DB}(d,l_D,l_B)\mu_B(l_B)$. They trade if and only if the joint surplus is positive. If they trade, the price $p_{DB}(\epsilon,x,d,l_D,l_B)$ is determined by generalized Nash bargaining such that the dealer's surplus is $\theta_{DB}(d)$ proportion of the total gains from trade; otherwise, the dealer's surplus is zero. Similarly, the location l_D type d vacant dealer's value function obeys

$$rW_0(d,l_D) = \sum_{l_S,x} \lambda_{SD}(d,l_S,l_D)\mu_S(x,l_S)\mathbb{E}\Big[\max\{\theta_{SD}(d)[S_{SD}(x,d,l_S,l_D)+\varepsilon],0\}\Big]. \tag{12}$$

At rate $\lambda_{SD}(d,l_S,l_D)\mu_S(x,l_S)$, he meets a location l_S seller with a type x car. They trade if and only if the joint surplus is positive, and they divide the surplus in fixed proportions $\theta_{SD}(d)$.

Steady-State Distribution The final set of equations describes the stationarity conditions of the distributions of buyers, sellers, and occupied dealers. At each location l_B, stationarity requires the inflow $\nu_B(l_B)$ of buyers to equate the outflow, such that

$$\begin{aligned}\nu_B(l_B) =& \mu_B(l_B)\sum_{l_S,x}\lambda_{SB}(l_S,l_B)\mu_S(x,l_S)\alpha_{SB}(x,l_S,l_B)\\ &+ \mu_B(l_B)\sum_{d,l_D,x}\lambda_{DB}(d,l_D,l_B)\mu_1(x,d,l_D)\alpha_{DB}(x,d,l_D,l_B), \quad \forall l_B,\end{aligned} \tag{13}$$

where the first term on the right-hand side of (13) represents the transactions between location-l_B buyers and sellers in any location owning any car, and the second term is the transactions between location-l_B buyers and occupied dealers of any type in any location owning any car.

Similarly, stationarity requires the inflow $\nu_S(x,l_S)$ of sellers with x-cars at location l_S to equate the outflow, such that

$$\begin{aligned}\nu_S(x,l_S) =& \mu_S(x,l_S)\sum_{l_B}\lambda_{SB}(l_S,l_B)\mu_B(l_B)\alpha_{SB}(x,l_S,l_B)\\ &+ \mu_S(x,l_S)\sum_{d,l_D}\lambda_{SD}(d,l_S,l_D)\mu_0(d,l_D)\alpha_{SD}(x,d,l_S,l_D), \quad \forall x,l_S,\end{aligned} \tag{14}$$

where the first term on the right-hand side of equation (14) captures the transactions of location-l_S sellers of x-cars and buyers in any location, and the second term is the transactions between such sellers and vacant dealers of any type in any location.

The inflow and outflow for the mass of each car characteristic x held by each dealer type d at each location l_D must also be equal. That is, for each (x, d, l_D),

$$\begin{aligned}&\sum_{l_B} \lambda_{DB}(d, l_D, l_B)\mu_1(x, d, l_D)\mu_B(l_B)\alpha_{DB}(x, d, l_D, l_B) \\ =&\sum_{l_S} \lambda_{SD}(d, l_S, l_D)\mu_s(x, l_S)\mu_0(d, l_D)\alpha_{SD}(x, d, l_S, l_D),\end{aligned} \tag{15}$$

where $\sum_x \mu_1(x, d, l_D) = m(d, l_D) - \mu_0(d, l_D)$. The left-hand side of equation (15) represents the transaction volume between location l_D type d dealers with type x cars and buyers at all locations, which is the outflow from $\mu_1(x, d, l_D)$. The right-hand side represents the transaction volume between location l_D type d vacant dealers and sellers with x-cars at all locations, which is the inflow into $\mu_1(x, d, l_D)$.

Definition 1. *A steady-state equilibrium consists of*

1. *agents' value functions* $V_S(x, l), V_B(l), W_1(x, d, l), W_0(d, l)$ *and*

2. *agent distributions* $\mu_B(l)$*,* $\mu_S(x, l)$*, and* $\mu_1(x, d, l)$*,*

$\forall x, d, l$*, such that conditions* (9)*,* (10)*,* (11)*,* (12)*,* (13)*,* (14)*, and* (15) *hold.*

In summary, our model allows buyers and sellers incur both monetary and time costs in searching for trading opportunities. Dealers have advantages in both bargaining and search-matching. Due to capacity constraints and search frictions, dealers are more likely to select higher-value goods for trade. Consequently, both intermediated transactions through dealers and direct transactions between sellers and buyers co-exist in equilibrium. The model can generate the dealer price premium based on their advantages in search, bargaining, and product selection, with the premium varying by dealer type, car characteristics, and locations.

3.3 Discussion of Assumptions

Quality Selection The model captures dealers' quality selection through a mechanism driven by search frictions, limited capacity, and the cost of holding inventory. Acquiring a car occupies the dealer's slot until the car is sold. Searching takes time, so maintaining inventory is costly. A dealer chooses to acquire a car only if the car is expected to sell quickly and at a good price. Because high-quality cars are more valued in the retail market, they are more likely to meet this threshold. As a result, dealers disproportionately trade high-quality cars, while lower-quality cars remain held by private sellers. Because quality is unobserved by economists, this selection mechanism helps explain some empirical patterns documented in Section 2.3, for example, dealer cars are sold at higher price, even after controlling for observed car characteristics.

This selection mechanism arises under symmetric information: buyers observe car quality, and negotiation is based on the true value of gains from trade. An alternative approach, following the tradition of Akerlof (1970), attributes dealers' quality selection to their information advantage: buyers cannot observe car quality, but dealers can and act as certifiers by selectively trading high-quality cars (see, e.g., Biglaiser (1993), Biglaiser and Li (2018), and Biglaiser et al. (2020)). This information-based mechanism is also consistent with some (but not all) empirical patterns that we documented, and we believe that it may also motivate dealers to select higher-quality cars in reality.

We choose not to incorporate asymmetric information for several reasons. First, search frictions are central to our analysis, as they allow us to capture dealers' role in facilitating trade and explain variation in car's time on the market across seller types and locations. Once search and capacity constraints are in place, quality selection follows endogenously. Second, while asymmetric information can lead to similar selection behavior, it is empirically difficult to distinguish it from the search-based mechanism. Third, incorporating asymmetric information would require taking a stand on how to model bargaining under incomplete information. It is well known that different assumptions about the bargaining protocol and information sets lead to substantially different predictions (e.g., Ausubel et al. (2002)), and adopting a particular approach involves nontrivial modeling commitments. For example, motivated by industry practice, Larsen (2021) adopts a mechanism design framework, but the implementation of such models typically requires richer data for identification. Finally, our focus is not to explain why dealer selection occurs, but to provide a tractable framework that accommodates its impact while analyzing dealers' advantages in matching and bargaining.

Dealer Inventory We assume that each dealer holds a single inventory slot. This effectively imposes constant returns to scale in inventory, so that a n-slot dealer is treated as equivalent to n single-slot dealers of the same type. This assumption is adopted for computational tractability. Although it is not difficult to incorporate multiunit inventories into the theoretical framework, the computational burden becomes substantial. With heterogeneous products, the state of an individual dealer becomes the composition of its inventory, and the number of steady-state conditions increases exponentially with capacity size[12] Moreover, the assumption of single-slot

[12]For instance, Li et al. (2024) studies inventory management in a directed search model. While the theoretical extension to incorporate product differentiation is discussed, the quantitative analysis focuses on a homogeneous product setting.

dealers is less limiting than it seems.

Our model allows dealers and private sellers to differ in matching efficiency, a key parameter that captures differences in their trading performance. These differences reflect not only *intrinsic factors* such as location, visibility, and buyer access, but also *scale effects*: Since real-world dealers have more cars, a buyer who visits these dealers can sample more vehicles and is more likely to find a good match. In equilibrium, this leads to different average matching rates for cars held by dealers than those held by sellers. Moreover, we also allow these parameters to differ across dealer groups, accommodating some heterogeneity within the dealer sector.

Crucially, in the next section, we estimate the model using steady-state group-level averages of price, sale speed, and inventory, not individual-level dynamics. Our specification does not prevent us from capturing the average differences between dealers and private sellers (and even between different dealer types) in terms of pricing, product selection, and the effective selling speed of each car they handle. The model attributes these differences to the underlying primitives (e.g., matching efficiency, bargaining power, and search and inventory cost) between sellers and each dealer type.

Admittedly, the estimated differences in matching rates between dealers and sellers do not distinguish dealers' intrinsic advantages from the scale effects. Consequently, when we perform counterfactual exercises in which the average inventory of a particular dealer type differs from our data counterpart, our assumption that the matching efficiency of that dealer type remains the same as our estimated one can cause bias. For example, when we reduce dealers' capacity, their inventory at the new equilibrium tends to be lower, and hence their matching efficiency tends to be lower as well if the ignored scale effect is positive. By keeping their matching efficiency unchanged, we would underestimate the effect of capacity reduction.

4 Estimation

We estimate the model using data from the used car market described in Section 2. Unfortunately, our data lack information on potential buyers, sellers, and activities in the wholesale market. To overcome this problem, we parameterize the model and impose several necessary assumptions. Later, we provide an informal discussion of identification. In particular, we take advantage of the empirical variations in transaction volume, price, time on market, and dealer inventory in vintage cars, dealer types, and locations, most of which have been discussed in Section 2.3. After that, we present our estimation approach that minimizes the distance between model predicted moments and their empirical counterparts. Lastly, we report the estimation results and assess the fit of the model.

4.1 Parametric and Functional Form Assumptions

We choose one week as the unit of time and set the weekly discount rate at $r = 9.615 \times 10^{-4}$ so that the annual interest rate is about 5%. Accordingly, we compute the transaction volume and dealer inventory at the weekly level and then take the averages across weeks as steady-state outcomes. We consider three types of dealers: independent dealers, franchised dealers, and CarMax. We set the dealer capacity (i.e. $m(d, l)$) for each dealer type and county as the maximum number of weekly inventory.

We decompose a car's characteristics into two components: $x = (y, z)$, where y represents the car's observed characteristics (observed by agents and econometricians) and z represents the car's unobserved characteristics or quality (observed by agents but unobserved by econometricians). In practice, we categorize each car into an age group $y \in \{1, 2, ..., 5\}$, where $y = 1$ corresponds to vehicles aged 4 to 5 years, $y = 2$ to 6-7 years, etc. For simplicity, we consider only two types of unobserved quality, i.e. $z \in Z \equiv \{L, H\}$, where L stands for low quality and H stands for high quality. Furthermore, we specify the deterministic component of a buyer's payoff from owing an (y, z)-car to be $u(y, z) = u_z\delta^{y-1}$, where $0 \leq u_L \leq u_H$, and $\delta \in [0, 1]$ is the utility discount rate as cars age. Here, u_H and u_L are the utilities of high-quality and low-quality cars with vintage $y = 1$, respectively.

We assume that the scale parameter of buyers' idiosyncratic utility shocks $\sigma(y, z)$ depends only on car age y, given by $\sigma(y) = \sigma_0 + \sigma_1 \times (y - 1)$, where $\sigma_0 > 0$ measures the dispersion of the utility shock of the youngest group, and $\sigma_1 \lesseqgtr 0$ captures its age trend. As in standard search models (Chapter 6.3.2 of Ljungqvist and Sargent (2018)), our model also predicts that a smaller $\sigma(y)$ implies a lower option value of buyer searching and hence leads to faster trading. In addition, we let the scale parameter of the dealers' pay-off shocks in the wholesale market be a constant, that is, $\sigma_w(x) = \sigma_w$, considering that the data lack information on the wholesale market (seller to dealer) transactions.

For the matching coefficients on the retail market, we only distinguish whether the two trading partners are from the same county or from different counties. That is, for any seller county - buyer county pair (l, l'), we assume

$$\lambda_{SB}(l, l') = \begin{cases} \lambda^0_{SB} & \text{if } l = l' \\ \lambda^1_{SB} & \text{if } l \neq l \end{cases}, \quad \text{and} \quad \lambda_{DB}(d, l, l') = \begin{cases} \lambda^0_{DB}(d) & \text{if } l = l' \\ \lambda^1_{DB}(d) & \text{if } l \neq l' \end{cases}, \forall d,$$

where the superscript $i = 0$ represents within-county trade and $i = 1$ represents cross-county trade. Notice that the matching coefficient still depends on the seller type. On the wholesale side, we assume a common seller-dealer matching efficiency parameter across dealer types and locations, that is, $\lambda_{SD}(d, l, l') = \lambda_{SD}$. We make this assumption because our data lack information on the sources of dealer cars. For the same reason, we assume that dealers of each type have the same type-dependent bargaining weights on both the retail and wholesale markets: $\theta_{DB}(d) = \theta_{SD}(d) = \theta_D(d) \in [0, 1]$ for each dealer type d.

We observe neither the numbers of buyers and sellers arriving at each moment nor the steady-state mass of them. Therefore, we take an indirect approach, rather than estimating the inflows (equal to outflows in steady state) of buyers and sellers for each location and car characteristics, $\nu_B(l)$ and $\nu_S(y, z, l)$. We first estimate the steady-state equilibrium mass of buyers and sellers under the assumption that they are functions of observed local populations. Then, we utilize the first-step estimates and the equilibrium steady-state conditions to recover parameters $\nu_B(l)$ and $\nu_S(y, z, l)$.

Specifically, we specify the steady-state mass of buyers in county l_B to be a function of the total population in that county: $\mu_B(l_B) = \text{Population}(l_B)^{\phi_B}$, where $\text{Population}(l_B)$ is the total population in county l_B in 2017 measured in thousands, and $\phi_B > 0$ is a parameter measuring the elasticity of the buyer mass with respect to population. Similarly, we specify the steady-state mass of sellers owning (y, z) cars in county l_S as a function of the local population, the age and quality of the car, given by

$$\mu_S(y, z, l_S) = \text{Population}(l_S)^{\phi_S} \times \frac{e^{\eta_y \times y}}{\sum_{y' \in Y} e^{\eta_y \times y'}} \times \begin{cases} \dfrac{e^{\eta_z - y}}{1 + e^{\eta_z - y}}, & \text{if } z = H \\ \dfrac{1}{1 + e^{\eta_z - y}}, & \text{if } z = L \end{cases}$$

where $\text{Population}(l_S)$ is the total population in county l_S, $\phi_S > 0$ is a parameter that measures the elasticity of the seller mass with respect to population, and η_y and η_z are parameters governing the age and quality distribution of cars owned by sellers. For example, $\eta_y > 0$ (< 0) implies that more (fewer) sellers own older cars, while $\eta_y = 0$ implies a uniform age distribution of cars owned by sellers. η_z governs the overall level of the fraction of high-quality cars of all car age groups. A larger η_z implies a larger fraction of high-quality cars of all age groups.

In addition, we assume that dealers' inventory cost is dealer-type specific only, i.e., $c(d, l_D) = c(d)$, and assume a common search cost for all buyers at any location and a common search cost for all sellers, i.e., $\kappa_B(l_B) = \kappa_B$ and $\kappa_S(l_S) = \kappa_S$.

The above specifications leave the following 28 parameters to be estimated: (i) parameters associated with buyers' utility from owning cars $(u_L, u_H, \delta, \eta_z)$, (ii) the scale parameters of surplus shocks $(\sigma_0, \sigma_1, \sigma_w)$, (iii) search costs of buyers and sellers (κ_B, κ_S), (iv) dealers' inventory costs $\{c(d)\}_{d \in \mathcal{D}}$, (v) parameters linking the mass of buyers and sellers to population stocks (ϕ_B, ϕ_S, η_y), (vi) bargaining weights $(\theta_{SB}, \{\theta_D(d)\}_{d \in \mathcal{D}})$, and (vii) matching coefficients

$\{\lambda^i_{SB}, \lambda^i_{DB}(d)\}_{d\in\mathcal{D},i=0,1}, \lambda_{SD}$.

Let Θ summarize all parameters to be estimated. Our estimation follows the classical minimum distance method, such that our estimator $\hat{\Theta}$ is chosen to minimize the following objective function:

$$(m(\Theta) - \bar{m})'\Omega(m(\Theta) - \bar{m}),$$

where $m(\Theta)$ is the vector of moments computed from the equilibrium outcomes of the model given a parameter vector Θ, $\bar{m}$ is the vector of corresponding sample moments, and Ω is a weighting matrix. Specifically, for each parameter vector Θ, we compute the following seven groups of moments: (i) the average transaction price of cars with vintage y between private sellers in county l_S and buyers in county l_B, which equals $P_{SB}(y, z, l_S, l_B)$ given by equation (6) integrated over the distribution of unobservable car quality z,[13] (ii) the minimum price of cars with vintage y between private sellers in county l_S and buyers in county l_B, which equals $V_S(y, L, l_S)$ given by equation (4) because the buyer values low-quality car ($z = L$) less, (iii) the average transaction price of cars with vintage y between type d dealers in county l_D and buyers in county l_B, which equals $P_{DB}(y, z, d, l_S, l_B)$ given by equation (7) integrated over the distribution of unobservable car quality z, (iv) the minimum transaction price of cars with vintage y between type d dealers in county l_D and buyers in county l_B, which equals $W_1(y, L, d, l_D) - W_0(d, l_D)$, (v) the transaction volumes of cars with vintage y purchased by buyers in county l_B from private sellers in county l_S and type d dealers in county l_D, given by

$$\lambda_{SB}(l_S, l_B)\mu_B(l_B)\mathbb{E}_z\Big[\mu_S(y, z, l_S)\alpha_{SB}(y, z, l_S, l_B)\Big],$$
$$\lambda_{DB}(d, l_D, l_B)\mu_B(l_B)\mathbb{E}_z\Big[\mu_1(y, z, d, l_D)\alpha_{DB}(y, z, d, l_D, l_B)\Big],$$

(vi) the average time on the market of cars with vintage y sold by dealers of type d in county l_D, given by

$$\mathbb{E}_z\Big[\frac{1}{\sum_{l_B}\lambda_{DB}(d, l_D, l_B)\mu_B(l_B)\alpha_{DB}(y, z, d, l_D, l_B)}\Big],$$

and (vii) the mass of occupied dealer slots $\mu_1(d, l)$ for each (d, l), calculated using the equilibrium steady-state condition (15).

We have a total of 959 individual moments to match. We set the weighting matrix Ω to be a diagonal one such that the expected transaction prices are weighted by the transaction

[13]To economize on the number of statistics to which we fit the model, we only target the average price of vintage-y cars sold by sellers located in each county l_S, the average price of vintage-y cars purchased by buyers located in each county l_B, and the average price of vintage-y cars traded within each county l. Accordingly, we do the same aggregation when we compute the sample moments of the seller-buyer transaction volumes and price, as well as dealer-buyer transaction volumes and price.

volumes, and all the other moments are uniformly weighted within each moment group.[14] An alternative weighting matrix is the one efficiently chosen by bootstrapping the data and inverting the resulting variance-covariance matrix of the empirical moments. This alternative is less preferred for two reasons. First, the empirical moments are computed from two datasets, but the two datasets are hardly independent. Second, even if we treat the two datasets as independent, some statistics, such as a subset of dealer transaction prices and volumes, are much more precisely estimated than others, such as inventory distribution and time on the market. The resulting weighting matrix is therefore highly unbalanced. Therefore, we end up choosing a weighting matrix that helps to achieve a more balanced fit to the targeted moments.

4.2 Identification

Identification shares key similarities with the structural labor search literature (Eckstein and Van den Berg, 2007).[15] Our model is highly nonlinear, with all parameters jointly affecting all equilibrium outcomes. Nonetheless, identification can be intuitively understood by noting that certain moments crucially depend on specific parameters, but not others.

Bargaining Parameters We begin with the case in which car quality z is observable to econometricians. By equation (4), the minimum price set by a seller for an x-car at location l must be equal to the seller's continuation value, $V_S(x, l)$. For each car type x, equation (5) implies a linear relationship between the average retail price and the minimum retail price, with slope $1 - \theta_{SB}$. Therefore, θ_{SB} can be identified from the joint variation of the average and minimum prices across seller locations.

When car quality is unobservable, the identification argument proceeds in three steps. For each (y, l), the retail price distribution is a mixture of two truncated components, each conditional on quality $z \in \{L, H\}$, with the lower support bound given by $V_S(y, z, l)$. The first step is to identify $V_S(y, L, l)$, the continuation value of the seller for a low quality car. Because a low-quality car generates lower utility, we have $V_S(y, L, l) < V_S(y, H, l)$, and the minimum price of direct transactions for an age-y car at location l reveals $V_S(y, L, l)$. The second step is to identify $V_S(y, H, l)$, the minimum price of high-quality cars. Following Flabbi and Moro

[14]More specifically, the moments are divided into five groups: mean transaction prices, minimum transaction prices, transaction volumes, time on the market, and the distribution of occupied dealer slots. The weights across groups are uniform. Within the group of mean transaction prices, the weights of individual moments are proportional to the corresponding transaction volumes and sum up to one; within each of the four other groups, the weights of individual moments are uniform and sum up to one.

[15]See the canonical work of Flinn and Heckman (1982), and the following papers that included unobserved heterogeneity (see, e.g., Eckstein and Wolpin (1995) and Flabbi and Moro (2012)) and location (see, e.g., Todd and Zhang (2022)).

(2012), this can be achieved by locating the discontinuity (or jump) in the density of the price distribution. Any price between $V_S(y, L, l)$ and $V_S(y, H, l)$ must be generated by low-quality cars, while prices above $V_S(y, H, l)$ come from both type of quality. The third step is to compute the conditional mean price of low-quality cars within the interval $(V_S(y, L, l), V_S(y, H, l))$ and repeat the logic from the case of observable quality. The gap between this conditional mean and the minimum price, combined with the average prices, identifies θ_{SB}. An analogous procedure applies to the identification of the dealers' bargaining parameters $\theta_{DB}(d)$.

Payoff Parameters Three model primitives can contribute to the age patterns of endogenous outcomes: mechanical depreciation of utility in car age (δ), age-dependent distribution of unobserved quality (η_z, u_L/u_H), and the age trend of the scale of surplus shocks (σ_1). The identification of these parameters relies on the empirical age patterns of average time on the market and retail prices and their variations across transaction types and locations.

Intuitively, as cars age, their average value declines due to two main forces: mechanical depreciation and a reduction in average unobserved quality. These factors jointly contribute to the downward trend in average retail prices with age. Importantly, the age pattern of average retail prices reflects the combined effects of these two forces, while the age pattern of the price premium, that is, the price gap between dealer and non-dealer sales, captures dealers' selective trading of higher quality vehicles at different ages. Together, these two age patterns allow us to separately identify the mechanical depreciation rate (δ) and the parameters that govern the distribution of unobserved quality (η_z, u_L/u_H). This has been numerically confirmed by Figures 3a and Figure 3b that plot the marginal contribution of each relevant parameter to the equilibrium prices of cars sold by private sellers and to the equilibrium prices of cars sold by franchised dealers, respectively. These two figures also show that a larger δ (slower depreciation) tends to flatten the age pattern of both types of transactions, while η_z and u_L/u_H primarily influence the age profile of the price premium. Specifically, an increase in η_z (more high-quality cars in each age group) moderately increases the price premia of older cars; an increase in u_L/u_H reduces the price premia of all age groups, more so for younger cars.

However, a decline in car value with age, implied by both depreciation and quality decay, would predict that older cars remain on the market longer because of lower valuations of cars. Therefore, the observed decline in time on the market with age (see Table 1) must be driven by a countervailing force, which pins down the negative sign of σ_1. Intuitively, $\sigma_1 < 0$ implies that the dispersion of buyer taste narrows with car age so that sellers of older cars are less likely to encounter buyers with very high idiosyncratic valuations. This reduces the option value of waiting and leads to faster sales for older vehicles. Figure 3c plots the marginal contribution of each relevant parameter to the equilibrium time on market of cars listed by franchised dealers

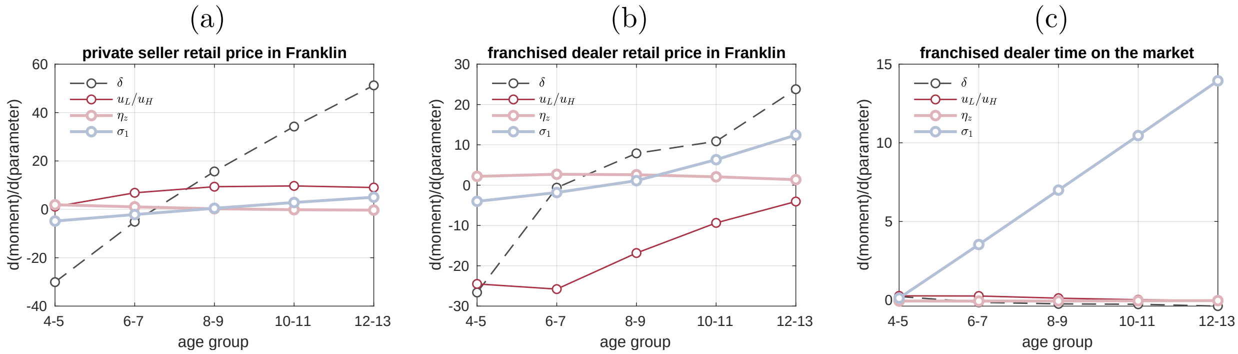

Figure 3: Marginal Impacts of δ, u_L/u_H, η_z, σ_1 on Key Moments

Note. Each panel plots the value of the corresponding elements in the numerically approximated Jacobian matrix of moments with respect to each relevant parameter ($G := \partial m(\Theta)/\partial\Theta'$). Every $\circ$ represents a particular entry of G, that is, the sensitivity of a model implied moment to a particular parameter. Each line style represents a particular parameter's impacts on the corresponding moments for each age group: δ (dashed black), u_L/u_H (thin dark red), η_z (thick light red), σ_0 (thin dark blue), and σ_1 (thick light blue). Panel (a) plots how the selected parameters affect model-implied time on the market of cars listed by franchised dealers ($d = 2$) for each age group. Panel (b) shows how each of selected parameters affects the private sellers' retail prices in Franklin county ($l = 3$), and panel (c) displays their impacts on franchised dealers' retail prices in Franklin county.

for each car age group. It shows that the decline of the time on the market in car age is primarily driven by $\sigma_1 < 0$. An increase in σ_1 (less negative) prolongs the time on market of older cars more than that of younger cars. The impacts of depreciation and quality decay are in the opposite direction, albeit small.

Finally, η_y is identified by the empirical car age distribution in the sample, the cross-sectional average retail price identifies the level of u_H, and the retail price dispersion identifies σ_0.

Matching and Population Parameters For each seller type, the difference in matching coefficients, λ^i_{SB} and $\lambda^i_{DB}(d)$, between within-county ($i = 0$) and cross-county ($i = 1$) can be identified by how the volume of within-county transactions differs from that of cross-county transactions. The overall level of time on market can help identify the levels of matching efficiency.

The equilibrium conditions for a stationary inventory distribution across dealer types and car characteristics identify the wholesale market parameters λ_{SD} and σ_D. Across buyer locations, equilibrium buyer stock varies with local population, and so does the volume of dealer-to-buyer transactions. Given the previously identified parameters and the observability of dealer inventories, the joint variation in population and dealer-buyer transaction volume across locations identifies ϕ_B. Consequently, by the same logic, across seller locations, the joint variation in population and seller-buyer transaction volume identifies ϕ_S.

Search and Inventory Cost Parameters Lastly, having above parameters been identified, one can compute the equilibrium trade probability $\alpha_{SB}(\cdot), \alpha_{DB}(\cdot)$ and $\alpha_{SD}(\cdot)$ and the equilibrium value for each agent type $V_S(\cdot), V_B(\cdot), W_1(\cdot)$ and $W_0(\cdot)$. Then, the system of HJB equations identifies the remaining parameters, including buyers' search cost κ_B, sellers' search cost κ_S, and the inventory cost of each dealer type $c(d)$.

4.3 Estimation Results

Table 4 reports the estimates of the utility parameters, the surplus shock scale parameters, the search costs of buyers and sellers, the inventory costs of dealers, and the parameters that link the steady-state mass of buyers and sellers to the local population. Table 5 reports the estimates of the parameters related to the bargaining power and the matching technology.

Utility and Surplus Shock Scale Parameters The base utility of a 4–5-year-old car is \$33,608 when quality is high and drops to 70.9% of this value when quality is low, indicating that buyers value low-quality cars substantially less than high-quality ones. The age discount rate $\delta = 0.954$ suggests that buyers' utility mechanically depreciates with car age. Figure 4a plots the implied utilities of the two types of quality $u(y, z)$ for each vintage.

The baseline scale parameter of buyers' surplus shocks for 4-5 years old cars is $\sigma_0 = 1.042$, suggesting that buyers' idiosyncratic valuations of cars are noticeably heterogeneous. The negative sign of $\sigma_1 = -0.09$ indicates that buyers' idiosyncratic utility shocks are less dispersed for older cars than for younger cars. This is not surprising, considering that as designs become outdated and features lose relevance over time, buyers' idiosyncratic valuations of these features tend to converge. For example, a 2005 car may have different types of CD players, but such features are almost irrelevant to buyers in 2017. As a result, as the car ages, the variation in how buyers value this car becomes smaller. Lastly, the scale parameter of dealers' surplus shocks in the wholesale market $\sigma_w = 0.011$, which is small in magnitude. It suggests that dealers' valuations of wholesale cars are fairly homogeneous, after controlling for car vintage.

Search Costs and Inventory Costs Buyers' weekly search cost is \$35, and private sellers' weekly search cost is \$47. Multiplying them by their estimated search time gives their expected total search costs, which are \$218.4 for buyers and \$404.2 for sellers. An occupied slot costs the dealer \$106, \$196, or \$511 per week for independent dealers, franchised dealers, and CarMax, respectively. Similarly, multiplying by the average time on the market, we obtain the dealers' average total cost of selling a car, which ranges from \$700 to \$1000. These estimates may appear to be high. However, notice that the inventory cost in our model includes all costs involved in selling a car, including mortgage payments and insurance payments of holding the car, rent of

Table 4: Estimates of Agent-Specific Parameters

Utility	u_H ($\times\$1,000$)	δ	u_L/u_H	
	33.608 (0.0007)	0.954 (0.0004)	0.709 (0.0004)	
Surplus shock scale	σ_0	σ_1	σ_w	
	1.042 (0.0214)	-0.090 (0.0080)	0.011 (0.0003)	
Search cost κ ($\times\$1,000$)	buyer	seller		
	0.035 (0.0004)	0.047 (0.0003)		
Inventory cost $c(d)$ ($\times\$1,000$)	independent dealer	franchised dealer	CarMax	
	0.106 (0.0003)	0.196 (0.0004)	0.511 (0.0004)	
Buyer and seller mass	ϕ_B	ϕ_S	η_y	η_z
	0.986 (0.0019)	0.950 (0.0005)	0.225 (0.0004)	0.445 (0.0638)

Note: Standard errors (in parentheses) are calculated using 150 bootstrapped samples from the data.

space, utilities, inspections and cleaning, labor costs of test drives, and other transaction costs such as title transfer, etc.

Buyer and Seller Mass Parameters Estimates of ϕ_B and ϕ_S suggest that as the local population increases by 1%, the mass of buyers and the mass of private sellers in steady state increase by 0.986% and 0.950%, respectively. The positive sign of η_y indicates that the mass of private sellers on the market increases with car age. The estimate of η_z is difficult to interpret on its own. Based on the estimate, we plot the share of high-quality cars held by private sellers for each car vintage in Figure 4b. This share starts around 40% for the youngest group and decreases to less than 2% for the oldest group.

Bargaining Parameters Our estimates suggest that dealers have higher bargaining weights than private sellers, indicating that they have a greater ability to extract trade surplus. Compared to private sellers who extract 64% of the trade surplus, independent dealers can extract a slightly higher share, franchised dealers can extract 76%, and CarMax can extract almost all. This is in line with the observation that CarMax and most franchised dealers actually use a non-haggling pricing strategy.

Matching Efficiency Parameters The estimates of matching coefficients indicate that dealers are more efficient at matching with buyers compared to private sellers. This advantage is

(a) $u(y,z)$ (b) fraction of $z=H$

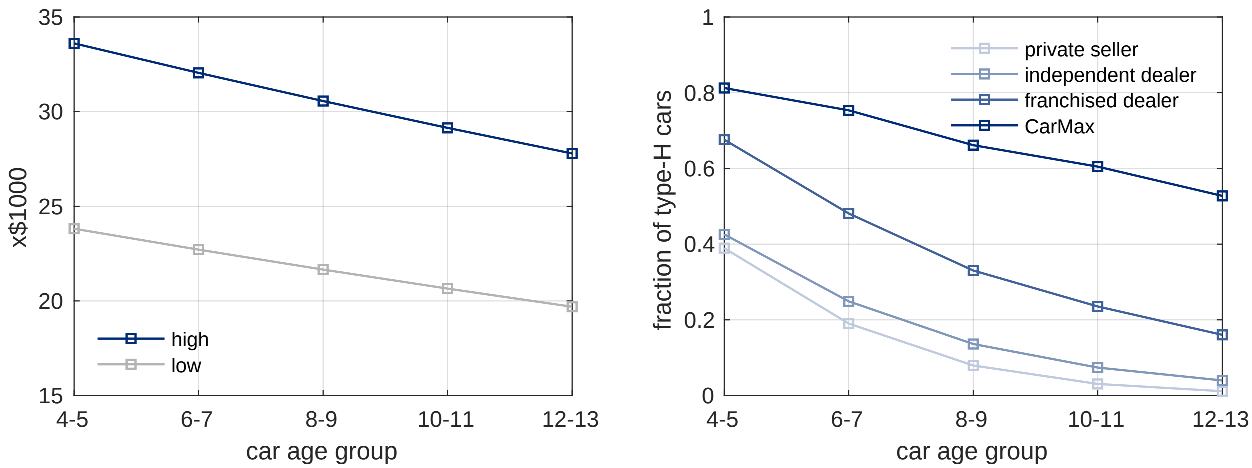

Figure 4: Estimated Utility Function and Quality Distribution

most pronounced for CarMax, followed by franchised dealers. Dealers' advantage in matching efficiency reflects their prominence in the retail market.[16]

Furthermore, the estimates reveal two spatial features on matching efficiency in the used car market. First, for both direct and intermediate trade, within-county matching coefficients are larger than the across-county ones, which is expected, as it is easier for trading partners to meet if they live in the same county than if they live in different counties. This explains why used-car transactions are largely localized. Meanwhile, it suggests that the used car market is spatially segmented, in which geographic distance acts as a source of friction and determines who meets whom and, ultimately, who trades with whom. Segmentation remains a central feature of decentralized used-car markets.

The second and more interesting feature is that the extra efficiency of independent and franchised dealers in matching with within-county buyers relative to across-county buyers is smaller than that of private sellers. This is probably because, compared to individual sellers, dealers typically operate formal businesses with established reputations, advertising networks, and more online presence, making them more visible and attractive to buyers from a broader geographic range. As a result, their trading partners are geographically less constrained. This result also suggests one role for used car dealers: improving matching efficiency across locations and hence alleviating spatial frictions. By enhancing market integration, dealers can potentially intensify competition, reduce disparities in pricing, product variety, and transaction quality,

[16]It may also reflect the economies of scale: in reality, a buyer who visits a large dealer can see multiple cars at once, increasing the probability of each inventory being sold.

Table 5: Estimates of Bargaining and Matching Parameters

	Bargaining	Matching: Retail (0.001)		Matching: Wholesale (0.001)
		within-county	across-county	
private seller - buyer (SB) or dealer (SD)	θ_{SB}	λ^0_{SB}	λ^1_{SB}	λ_{SD}
	0.639	0.159	0.122	5.615
	(0.0004)	(0.0007)	(0.0005)	(0.0004)
dealer - buyer (DB)	$\theta_D(d)$	$\lambda^0_{DB}(d)$	$\lambda^1_{DB}(d)$	
independent dealer	0.656	0.291	0.267	
	(0.0004)	(0.0030)	(0.0011)	
franchised dealer	0.760	0.455	0.446	
	(0.0150)	(0.0087)	(0.0075)	
CarMax	0.989	1.068	0.531	
	(0.0196)	(0.0195)	(0.0513)	

Note: Standard errors (in parentheses) are calculated using 150 bootstrapped samples from the data.

creating welfare equality across locations.

One implication of dealers' better matching technology is that their trading speed is faster than private sellers, enabling them to realize their gain from trade faster. In particular, it accelerates the trading speeds even more for high-value cars, thereby enabling even faster realization of the associated gains from trading these cars. As a result, dealers are more likely to hold younger cars of higher quality. Figure 4b plots the share of high-quality cars held by each dealer type for each car vintage at the steady state equilibrium predicted by our model. Clearly, dealers hold a larger proportion of high-quality cars than private sellers for each car age group, and such a quality selection is more pronounced for larger dealers. For example, among the youngest cars held by the largest dealer, more than 80% are of high quality, which is more than doubled compared to private sellers. Finally, because dealers tend to purchase higher-quality cars, their selection behavior leaves a disproportionate share of lower-quality vehicles in the pool held by individual sellers. This residual composition mirrors the classic pattern of adverse selection in which market intermediaries absorb better products and leave behind those of inferior quality.

4.4 Model Fit

To assess how well our estimated model fits the data, we compare the moments predicted by the estimated model with the data counterparts. Figure 5 presents the key endogenous outcomes predicted by the model and observed in the data, including average prices and trading volumes

(a) average retail prices ($\times$\$1000) by car age

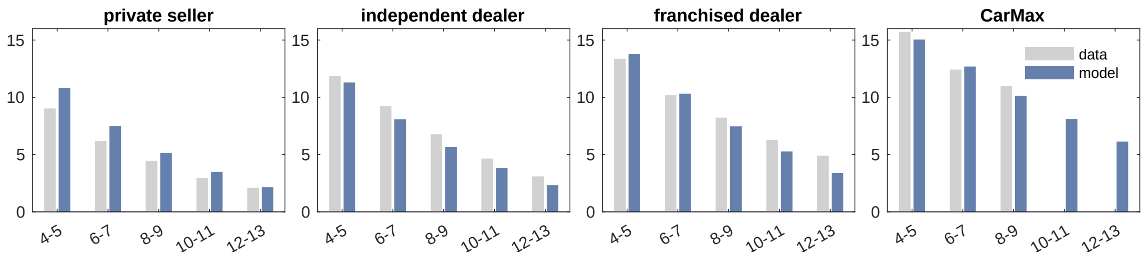

(b) retail transaction volumes by car age

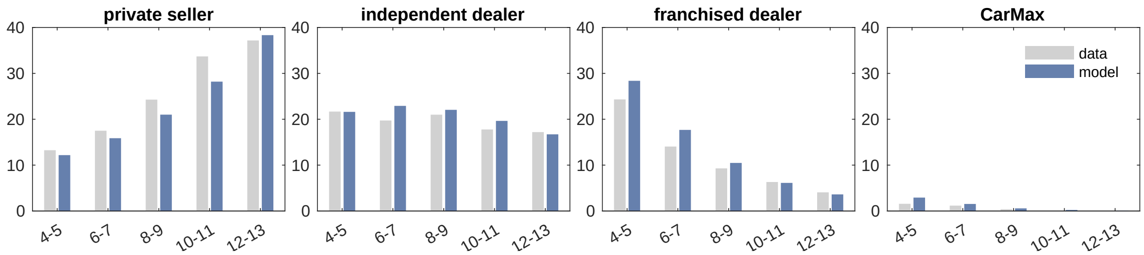

(c) time on the market by location (left) and mean-min price difference (right)

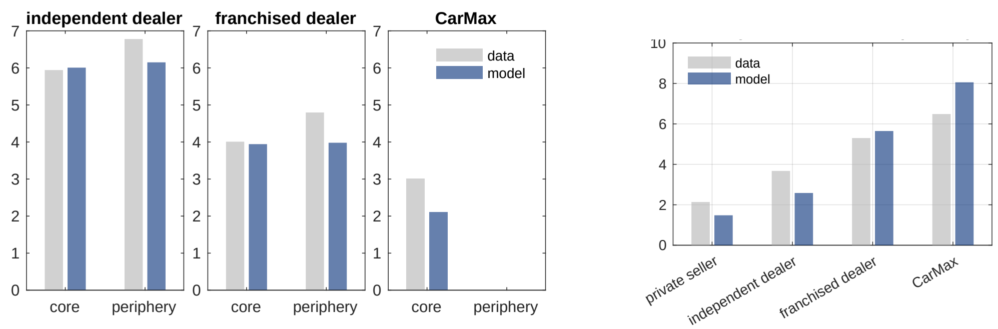

Figure 5: Model Fit

by age, the average and minimum price gaps, and time on the market.

Overall, the model fits the data well. In particular, as shown in Figure 5a, the model is able to reproduce (i) the monotonic declining of transaction price in car age for each seller type and (ii) the sizable and age-varying price premia for each dealer type. Notice that the model predicts higher seller prices than the data, which is consistent with the conjecture that direct transaction prices are under-reported. In addition, our model is able to match the mean-min price gap, trading volumes, and time on market of dealer cars, as shown in Figure 5b and Figure 5c. The model also matches the differences in transaction price and volume across locations. See the Appendix for details.

5 Counterfactual Analyses

We now conduct two sets of counterfactual analyses to address the following questions: (i) how important used-car dealers are and what drives their impact; and (ii) how spatial search frictions affect prices, trading delays, and welfare. What is the social value of dealers and why? (ii) How do spatial search frictions affect price, trading speed, and welfare?

5.1 Quantifying and Disentangling the Effects of Dealers

The first set of counterfactual simulations quantifies the impact of intermediaries in the used car market. We first reduce the capacity of used car dealers to quantify their overall impact. Since our estimates indicate that used car dealers have greater bargaining power and superior matching efficiency than private sellers, we further conduct counterfactual simulations to examine the contribution of each advantage. Table 6 displays the results.

5.1.1 Overall Impact of Dealers

The economics of reducing dealer capacity is as follows. As the dealer capacity decreases, trading shifts from through dealers (more efficient channel) to direct trading (less efficient channel), increasing the search frictions and slowing down the average trading speed. In addition, along with the reduction of dealers' presence, more high-quality cars are present in the private market. As a result, dealers' advantage in selecting high-quality cars diminishes and the gap in the trading speed between high-quality and low-quality cars narrows.

Reducing dealer capacity has different welfare impacts on different groups of agents. Buyers need to search longer. This also weakens their bargaining position and leads to higher retail

Table 6: Dealer Role Counterfactual Results

Outcome	Baseline	(1) Dealer Capacity	(2) Bargaining Power	(3) Dealer Matching
Buyer Value (V_B, \$)	18,574	17,142	16,531	16,463
Seller Value (V_S, \$)	4,195	5,748	7,118	6,655
Empty Dealer Value (W_0, \$)	6,874	9,580	2,351	4,113
Occupied Dealer Value (W_1, \$)	15,051	18,823	15,982	15,031
Avg. Dealer Price (P_{DB}, \$)	8,224	9,643	10,516	10,535
Avg. Seller Price (P_{SB}, \$)	4,667	6,217	7,605	7,141
Time on Market (Buyer, weeks)	6.242	6.429	7.241	6.585
Time to Sell (Dealer, weeks)	5.379	5.268	5.618	5.365
Dealer Share of High Quality	0.304	0.310	0.325	0.305
Seller Share of High Quality	0.094	0.104	0.153	0.122

Note: "Baseline" refers to the predictions of the model. "Dealer Capacity" refers to decreasing total dealer economy-wide capacity by 10%. "Bargaining Power" refers to reducing all dealers' bargaining power to the private sellers' value ($\theta_{SB} = 0.639$). "Dealer Matching" refers to reducing dealers' matching advantage by 10%.

prices. Hence, buyers tend to be worse off. In contrast, sellers and remaining dealers are better off because they face less competition.

The quantitative effect of reducing dealer capacity is remarkable. Column (1) in Table 6 shows that a 10% reduction of each dealer type's capacity at each location leads to \$1,432 welfare loss for an average buyer and \$1,554 welfare gain for an average seller. Meanwhile, it causes a value increase of \$3,943 for an average remaining dealer slot. Multiplying these changes per agent by the respective population sizes, the total value accruing to dealers increases by \$2.29 million, substantially exceeding the total welfare loss incurred by buyers and sellers combined.

In this case, the buyer's welfare loss and the seller's gain are roughly symmetric in magnitude, indicating that the welfare effects for buyers and sellers are primarily redistribution rather than net loss or gain. The major source of net surplus is the increased profitability of the remaining dealers: reducing capacity softens competition in the retail market, allowing surviving dealers to extract higher margins. These results suggest that the current level of dealer capacity may be excessive, both from the standpoint of dealers collectively and from the perspective of overall market efficiency.

It is important to note that dealers' matching coefficients remain the same as their estimated values. If dealers' matching technology exhibited increasing returns to scale, then a reduction in steady-state inventory level would diminish the scale effect, effectively leading to lower matching efficiency ($\lambda^i_{DB}(d), \forall i, d$). As we discuss later, this effect further discourages dealers from participating and amplifies the impact of reducing their capacity. Consequently, our quantification results tend to understate the welfare impact of dealers on buyers and sellers while exaggerating the welfare impact on the remaining dealer slots. This bias also exists when we diminish dealers' advantage in bargaining and matching efficiency one by one in the next counterfactual analysis.

To assess the robustness of these findings, we extend the analysis to a wider range of capacity reductions. Our simulation results confirm that the key qualitative impact of reducing the presence of dealers, including redistributing welfare between buyers and sellers, increasing dealer profits, and leading to net gains when modestly reducing the capacity, holds consistently across the spectrum. Figure A.4 presents the simulation results when dealers' capacity is reduced by other proportions.

5.1.2 Decomposition of Intermediaries' Advantages

Having examined the overall impact of intermediaries, we now separately quantify the impact associated with each of their two key advantages over private sellers: greater bargaining power and superior matching efficiency. In this analysis, we can also evaluate whether the social value of superior matching is worth the price premium associated with dealers' market power.

Effect of Dealers' Stronger Bargaining Power We first examine the impact of dealers' stronger bargaining power by decreasing their advantage. The welfare impact of reducing dealers' bargaining power can be non-monotonic, driven by the interaction of several countervailing effects. First, lowering dealers' bargaining power reduces the dealers' retail price, which in turn lowers private sellers' retail price due to competition. This direct effect benefits buyers, but harms both dealers and sellers. Meanwhile, with lower bargaining power, dealers have weaker incentives to place wholesale orders and maintain inventories. As fewer dealer slots are occupied, buyers' search value diminishes, weakening their bargaining position in negotiations with sellers. This indirect effect results in higher retail prices in direct seller-buyer transactions. The spillover effect further weakens buyers' bargaining positions in negotiations with dealers, ultimately raising dealer retail prices as well. Additionally, reduced intermediation slows the transaction speed of high-quality cars, despite dealers becoming more selective about car quality.

To quantify the comparative statics discussed above, we eliminate dealers' advantage in bargaining by setting the bargaining weights of all dealer types equal to that of sellers. The simulation results are presented in column (2) of Table 6. We find that when dealers have the same bargaining ability as private sellers, the equilibrium number of occupied dealer slots decreases by 21%, the average buyer value drops by \$2,043, the average seller value increases by \$2,923, and the average dealer value decreases by \$2,581.

In the Appendix we present additional counterfactual results when dealers' bargaining power is lowered gradually towards that of private sellers. As the gap in the bargaining power between dealers and private sellers narrows, the retail prices of both channels slightly go down, the average buyer value slightly increases while the average seller value slightly decreases, and the average dealer value substantially declines. Again, the welfare impacts on buyers and sellers are primarily reallocating.

Effect of Dealers' Superior Matching Technology To quantify the impact of dealers' better matching technology, we worsen dealers' matching technology toward that of private sellers, both for within- and across-county trade.

As dealers become less efficient in matching, their bargaining position weakens, discouraging them from holding inventory, given that it is costly to do so. With less participation of intermediaries, trading shifts towards the less efficient direct channel. So buyers need to stay on market longer before trading occurs. In addition, as dealers become less motivated to hold inventory, they become more selective: they hold inventory only if it is highly profitable to do so. When dealers' matching efficiency drops substantially, they only trade high-quality cars.

Our simulation results, reported in column 3 of Table 6, show that dealers' advantage in matching plays a significant role in shaping market outcomes. For instance, when dealers' matching efficiency moves close to private sellers' by 10%, the average buyer value drops by \$2,111 while the average seller value increases by \$2,460. Meanwhile, the steady-state number of dealer slots occupied drops by 13%, and the average dealer value decreases by \$1,786. To gain a complete picture, we consider a range of reductions in dealer matching efficiency in the Appendix.

5.2 Quantifying the Impact of Spatial Search Frictions

Our estimation results reveal a pronounced spatial disparity in matching technology: Trading parties are significantly more likely to match within the same location than across locations. In addition, dealers can match with distant buyers more efficiently than private sellers. In the second set of counterfactual simulations, we aim to examine the importance of such spatial frictions in determining the market outcomes and shaping the roles of intermediaries.

In this set of counterfactual exercises, we alleviate geographic frictions by increasing the efficiency of across-location matching. In some sense, the replicates the advantages that may be enjoyed by more recent online search platforms and mobile applications, like Carvana and Vroom. Our simulation results suggest how market outcomes – that is, trade volumes, intermediary roles, and price dispersion – are attributable to spatial frictions, and how they might evolve in a more integrated digital marketplace.

Specifically, we increase the across-location matching efficiency on the retail side towards its within-location counterpart. Intuitively, improving the efficiency of cross-location matching affects market outcomes through two channels. The more direct channel is that lowering trade barriers implies more competition on the retail side, which drives down the retail price and hence benefits buyers. However, a second, more subtle effect arises when it becomes easier for private sellers to match with buyers, and more trade takes place in the direct market (less efficient), which countervails the competition effect and puts upward pressure on prices.

To illustrate, we narrow the difference between within- and across-location search frictions by 20% and report the simulation results in Table 7.[17] To better understand the economic mechanisms behind the two effects discussed above, we decompose this exercise into two steps. In the first step, we only narrow the spatial heterogeneity when buyers match with dealers, and report the simulation outcomes in column (4). Comparing column (4) with the baseline column that reports the outcomes in the benchmark case, we can quantify the importance of only improving dealers' matching efficiency across locations. In the second step, we also narrow the spatial heterogeneity when buyers match with private sellers, and report the simulation results in column (5). By comparing column (5) with the baseline column, we can measure the total effect of the direct competition channel and the indirect trade-shifting channel.

When we only increase the across-location matching coefficients for dealer-buyer trade, dealers and buyers at different locations are more likely to meet, intensifying competition among dealers. As a result, buyers gain a stronger bargaining position in negotiations with both dealers and sellers, leading to a decrease in retail prices. Our simulation results indicate that on average the dealer price is \$353 lower and the average private price is \$331 lower on the retail side. On average, buyers are better off by \$303 while sellers are worse off by \$332. Notice that occupied dealers are worse off by \$375 due to intensified competition in the retail market, whereas unoccupied dealers are better off by \$653 as a result of the weaker bargaining position of sellers in the wholesale market. These results align with standard search theory intuition: When dealers are more likely to match with buyers, competition among dealers increases, driving down retail prices, lowering dealers' profit, and benefiting buyers.

Furthermore, the improved bargaining position of buyers extends to their negotiations with sellers, further reducing the retail prices and value of sellers. In response, dealers have stronger

[17] We simulate the model for various grid points of w ranging from 0 to 1, and the impacts are monotonic.

Table 7: Effect of Narrowing Spatial Discrepancy in Matching Efficiency

Outcome	Baseline	(4) Dealer-only	(5) Dealer and Seller
Buyer Value (V_B, \$)	18,574	18,878	17,240
Seller Value (V_S, \$)	4,195	3,865	5,316
Empty Dealer Value (W_0, \$)	6,874	7,527	5,167
Occupied Dealer Value (W_1, \$)	15,051	14,676	15,310
Avg. Dealer Price (P_{DB}, \$)	8,224	7,871	10,056
Avg. Seller Price (P_{SB}, \$)	4,667	4,335	5,779
Time on Market (Buyer, weeks)	6.242	6.195	6.378
Time to Sell (Dealer, weeks)	5.379	5.403	5.356
Dealer Share of High Quality	0.304	0.303	0.332
Seller Share of High Quality	0.094	0.092	0.074

Note: "Baseline" refers to the predictions of the model. "Dealer-only" refers to reducing dealer spatial friction discrepancy by 20%: $\lambda_{DB}^{1,\text{counterfactual}}(d) = 0.2\lambda_{DB}^{0}(d) + 0.8\lambda_{DB}^{1}(d)$. "Dealer and Seller" refers to also reducing sellers' spatial friction discrepancy by 20%: $\lambda_{SB}^{1,\text{counterfactual}} = 0.2\lambda_{SB}^{0} + 0.8\lambda_{SB}^{1}$.

incentives to order more inventory in the wholesale market for two reasons: (i) they can sell cars faster, reducing the role of inventory costs, and (ii) sellers are willing to accept lower wholesale prices due to their worse outside option in the retail market. Consequently, buyers and sellers spend less time on searching, while the number of occupied dealer slots increases.

When we further increase the across-location matching coefficient for seller-buyer trade, sellers are also more likely to meet buyers from different locations. As a result, sellers become less reliant on dealers and are more likely to sell directly to buyers. This shift strengthens sellers' bargaining position, increasing their retail prices. Furthermore, this also weakens the role of dealers as intermediaries, discouraging them from holding inventory unless it is highly profitable. As dealers trade less frequently, the average transaction speed declines, particularly for high-quality cars. Our simulation results suggest that the indirect trade-shifting effect dominates the direct competition effect. As reported in Table 7, improving the matching efficiency of all types of trade across locations actually increases the retail price by \$1,832 (22%) for intermediated trade and \$1,113 (24%) for direct trade. As a result, buyers lose \$1,334 and sellers gain \$1,121 on average. Meanwhile, occupied dealers slightly gain \$259 and unoccupied dealers lose \$1,707.

This analysis demonstrates that reducing search friction does not necessarily lead to lower price due to intensified competition, which the standard search literature has long recognized. When multiple trading channels of different efficiency levels coexist, such as the more efficient intermediated channel and the less efficient direct channel in our context, the reduction of search frictions may also shift the trading from the efficient channel to the less efficient channel and lead to higher prices.

This analysis also highlights that dealers' role as intermediaries crucially depends on the spatial heterogeneity of their advantage in matching over private sellers. According to our estimates, dealers' visibility is less restricted by geography, allowing them to trade with distant buyers more efficiently than private sellers. If it becomes easier for private sellers to trade with distant buyers, they tend to bypass dealers, which could weaken dealers' advantage in facilitating trade. This could result in misallocation by shifting trade away from efficient intermediaries, eventually harming buyers and benefiting sellers.

6 Conclusion

This paper develops and estimates a spatial search and bargaining model to understand the role of dealers in decentralized used car markets. Using detailed transaction data from an Ohio metropolitan area, we document substantial heterogeneity in pricing, transaction speed, and inventory composition across dealer types and locations. We then use our model to quantify how dealers' advantages in matching efficiency, bargaining power, and spatial positioning shape equilibrium outcomes.

Our results show that dealers trade vehicles with higher values at a price premium, enabled by their superior matching technology and stronger bargaining positions relative to private sellers. These advantages allow dealers to reduce the time on the market and capture surplus, though this comes at the cost of crowding out direct transactions. Counterfactual analyses highlight several key mechanisms. Limiting dealer capacity causes a shift in trade volume from efficient intermediated channels to less efficient direct trade, increasing frictions and distorting allocation. Reducing dealer bargaining power or matching efficiency leads to longer search times, higher prices, and reduced welfare, particularly for buyers. We also show that spatial frictions and dealer location are critical: lowering geographic trade barriers by equalizing within- and across-location matching efficiency reduces spatial frictions, weakens dealer selection, and shifts market power, ultimately distorting the allocation of vehicles and reducing overall welfare.

These results demonstrate the complex but important role of dealers in facilitating trade and shaping welfare. Their advantages generate both benefits and distortions, and their market impact is critically dependent on spatial frictions, inventory constraints, and their interaction with decentralized trade. Our framework provides a foundation for evaluating policies that affect intermediary behavior, such as dealership entry, online platforms, or inventory regulation.

References

Akerlof, G. A. (1970). The market for" lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 488–500.

Allen, J., R. Clark, and J.-F. Houde (2014). Price dispersion in mortgage markets. *The Journal of Industrial Economics 62*(3), 377–416.

Allen, J., R. Clark, J.-F. Houde, S. Li, and A. V. Trubnikova (2023). The role of intermediaries in selection markets: Evidence form mortgage lending. Technical report, National Bureau of Economic Research.

Ausubel, L. M., P. Cramton, and R. J. Deneckere (2002). Bargaining with incomplete information. *Handbook of game theory with economic applications 3*, 1897–1945.

Biglaiser, G. (1993). Middlemen as experts. *The RAND journal of Economics*, 212–223.

Biglaiser, G. and F. Li (2018). Middlemen: the good, the bad, and the ugly. *The RAND Journal of Economics 49*(1), 3–22.

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

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2020). Geography, transportation, and endogenous trade costs. *Econometrica 88*(2), 657–691.

Brancaccio, G., M. Kalouptsidi, T. Papageorgiou, and N. Rosaia (2023). Search frictions and efficiency in decentralized transport markets. *The Quarterly Journal of Economics 138*(4), 2451–2503.

Brancaccio, G. and K. Kang (2022). Search frictions and product design in the municipal bond market. Technical report, National Bureau of Economic Research.

Brancaccio, G., D. Li, and N. Schürhoff (2017). Learning by trading: The case of the us market for municipal bonds. Technical report, Princeton University.

Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies 89*(2), 556–591.

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, 955–969.

Demsetz, H. (1968). The cost of transacting. *The quarterly journal of economics 82*(1), 33–53.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Eckstein, Z. and G. J. Van den Berg (2007). Empirical labor search: A survey. *Journal of econometrics 136*(2), 531–564.

Eckstein, Z. and K. I. Wolpin (1995). Duration to first job and the return to schooling: Estimates from a search-matching model. *The Review of Economic Studies 62*(2), 263–286.

Flabbi, L. and A. Moro (2012). The effect of job flexibility on female labor market outcomes: Estimates from a search and bargaining model. *Journal of Econometrics 168*(1), 81–95.

Flinn, C. and J. Heckman (1982). New methods for analyzing structural models of labor force dynamics. *Journal of econometrics 18*(1), 115–168.

Frechette, G. R., A. Lizzeri, and T. Salz (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review 109*(8), 2954–2992.

Gavazza, A. (2016). An empirical equilibrium model of a decentralized asset market. *Econometrica 84*(5), 1755–1798.

Gavazza, A. and A. Lizzeri (2021). Frictions in product markets. In *Handbook of Industrial Organization*, Volume 4, pp. 433–484. Elsevier.

Gavazza, A., A. Lizzeri, and N. Roketskiy (2014). A quantitative analysis of the used-car market. *American Economic Review 104*(11), 3668–3700.

Gillingham, K., F. Iskhakov, A. Munk-Nielsen, J. Rust, and B. Schjerning (2022). Equilibrium trade in automobiles. *Journal of Political Economy 130*(10), 2534–2593.

Hornstein, A., P. Krusell, and G. L. Violante (2011). Frictional wage dispersion in search models: A quantitative assessment. *American Economic Review 101*(7), 2873–2898.

Judge, K. (2022). *Direct: The Rise of the Middleman Economy and the Revolution Underway.* HarperCollins.

Lagos, R. (2000). An alternative approach to search frictions. *Journal of Political Economy 108*(5), 851–873.

Lagos, R. (2003). An analysis of the market for taxicab rides in new york city. *International Economic Review 44*(2), 423–434.

Larsen, B. and A. L. Zhang (2021). Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry. *Available at SSRN 3990290*.

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

Li, F., C. Murry, C. Tian, and Y. Zhou (2024). Inventory management in decentralized markets. *International Economic Review 65*(1), 431–470.

Ljungqvist, L. and T. J. Sargent (2018). *Recursive macroeconomic theory*. MIT press.

McFadden, D. (1977). Modelling the choice of residential location.

Murry, C. and H. S. Schneider (2016). The economics of retail markets for new and used cars. In *Handbook on the Economics of Retailing and Distribution*, pp. 343–367. Edward Elgar Publishing.

Rubinstein, A. and A. Wolinsky (1987). Middlemen. *The Quarterly Journal of Economics 102*(3), 581–593.

Salz, T. (2022). Intermediation and competition in search markets: An empirical case study. *Journal of Political Economy 130*(2), 310–345.

Todd, P. E. and W. Zhang (2022). Distributional effects of local minimum wages: A spatial job search approach. Technical report, National Bureau of Economic Research.

A Appendix: More Figures and Tables

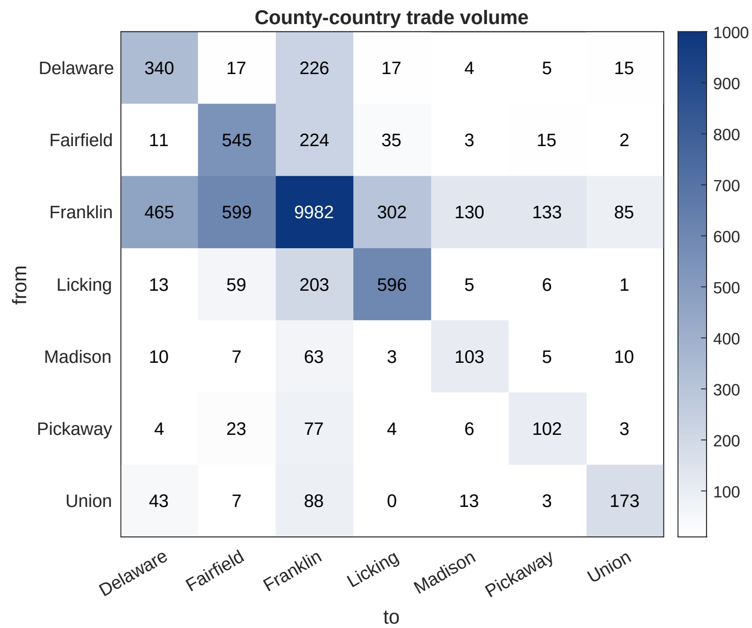

Figure A.1: Transaction Volume by County

Note: This matrix displays the total inflows and outflows of used cars among the seven counties in the Columbus area in 2017. The diagonal entries correspond to transactions within each county, whereas off-diagonal entries correspond to transactions across counties. In Franklin County, about 85% of within-county transactions take place, and it is simultaneously the largest exporter and importer to other counties. In other counties, the transaction volumes are much lower, and most transactions are either within-county or with Franklin. In fact, most locations are either large net importers (inflows exceeding outflows) or large net exporters (outflows exceeding inflows). Franklin is the largest net exporter, and Union's outflow is slightly higher than its inflow. The other five counties are net importers, with outflows falling short of inflows.

Table A.1: Hedonic Price Regression Results

	(I)	(II)	(III)
independent dealer	2.838	2.945	2.943
	(0.061)	(0.060)	(0.061)
Franchised dealer	4.354	4.186	4.171
	(0.072)	(0.071)	(0.072)
CarMax	6.325	6.286	6.276
	(0.205)	(0.200)	(0.200)
Car age group dummies			
6-7 years	-2.207	-2.272	-2.266
	(0.079)	(0.078)	(0.078)
8-9 years	-4.017	-4.174	-4.171
	(0.084)	(0.082)	(0.082)
10-11 years	-5.522	-5.749	-5.742
	(0.089)	(0.088)	(0.088)
12-13 years	-6.433	-6.745	-6.745
	(0.093)	(0.092)	(0.092)
log(Mileage)	-1.586	-1.603	-1.599
	(0.052)	(0.051)	(0.051)
Constant	26.150	26.480	26.440
	(0.579)	(0.565)	(0.566)
Transaction monthly FE	Y	Y	Y
Car model FE	N	Y	Y
Seller county $\times$ buyer county FE	N	N	Y
R^2	0.683	0.699	0.701

Note: An observation is a single transaction from the sample described in the text. The dependent variable is the transaction price measured in $1,000. Similar to Biglaiser et al. (2020), our strategy is to compare prices of four observably equivalent cars (same car model, same odometer mileage, same vintage group, same transaction month, and same seller county - buyer county), with the first one being sold directly by the owner, the second one being sold by a independent dealer, the third one being sold by a franchised dealer, and the last one being sold by a CarMax store, and we examine how much different the prices of the second, third, and the last cars relative to the first car. All specifications include transaction monthly fixed effects. Specification (II) includes monthly and car model fixed effects. Specification (III) includes monthly, car model, and seller county - buyer county fixed effects.

(a) average retail prices (×$1000) by location

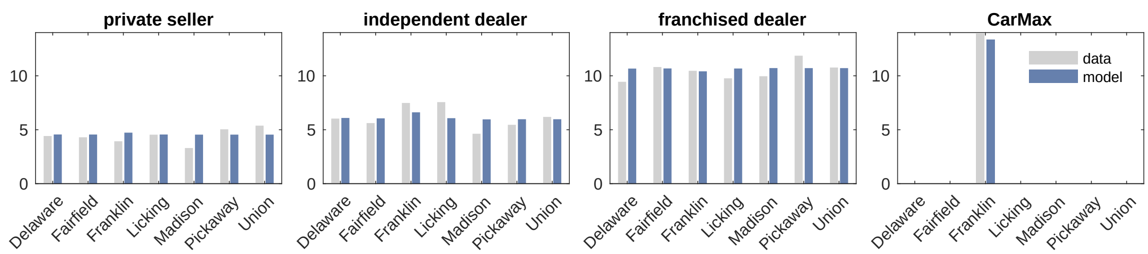

(b) retail transaction volumes by location

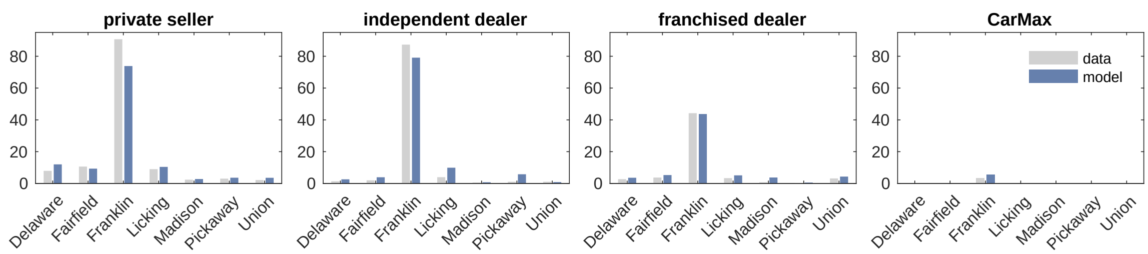

Figure A.2: Model Fit: Transaction Price and Volume by Location

(a) time on the market by car age

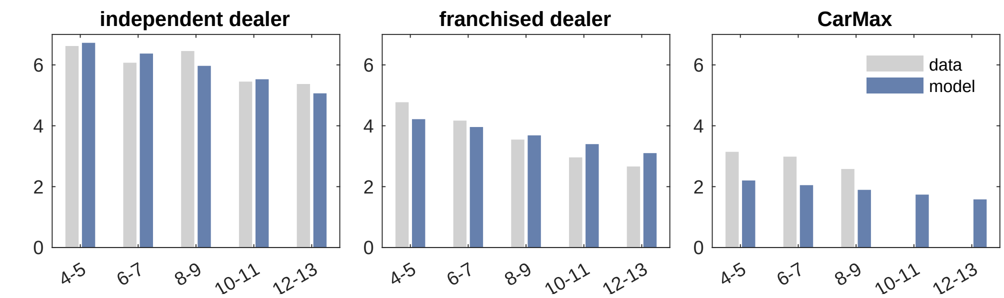

(b) dealers' inventory by car age

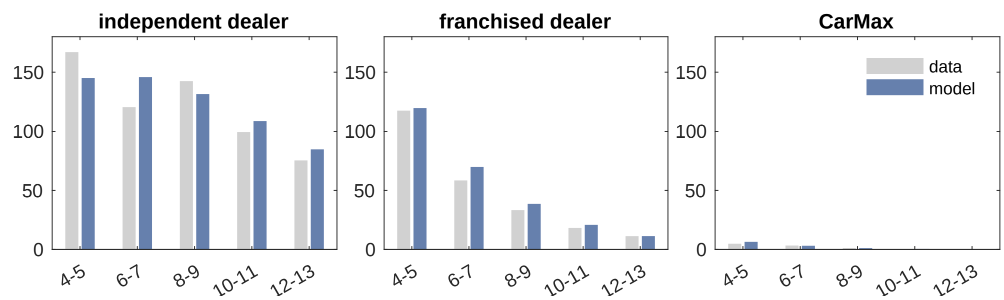

(c) dealers' inventory by location

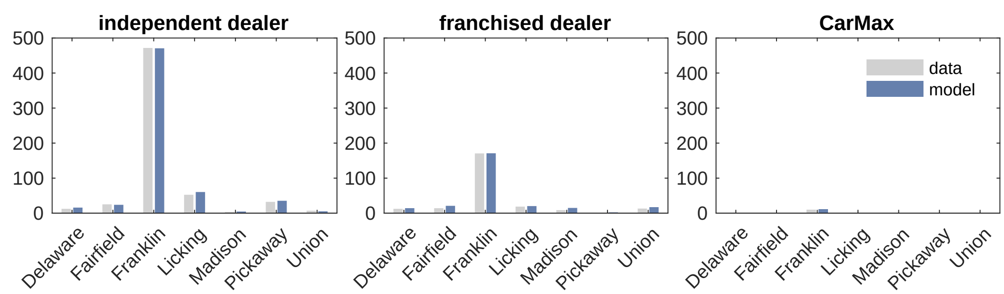

Figure A.3: Model Fit: Time on Market, and Dealer Inventory

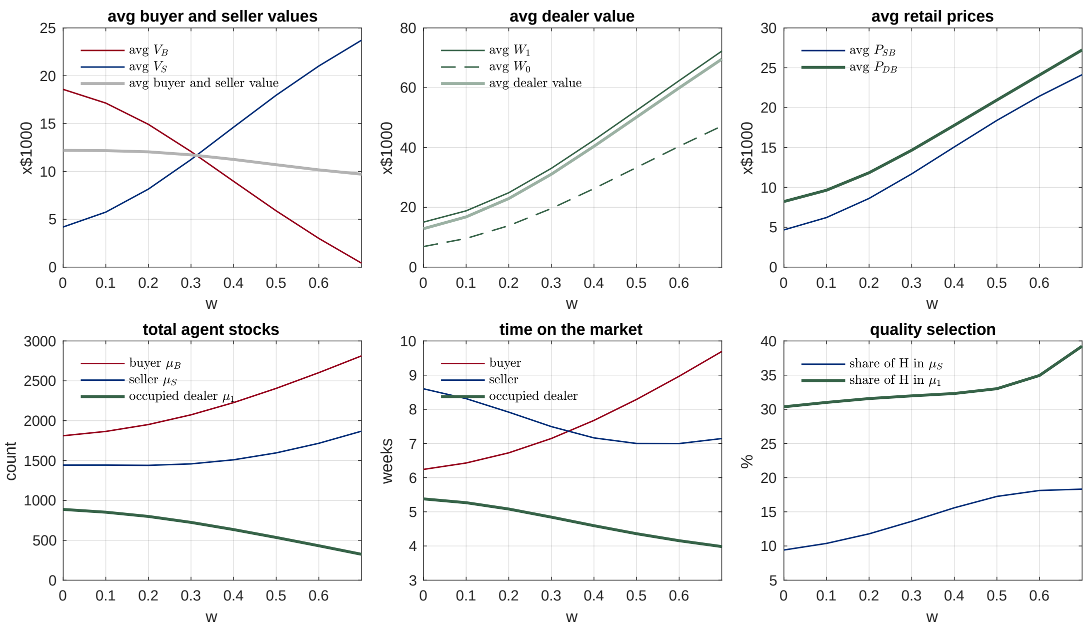

Figure A.4: Impact of Reducing Dealer Capacity

Notes: The x-axis shows the fraction w of dealer capacity reduced, such that the counterfactual dealer capacity $m^{\text{counterfactual}}(d, l; w) = (1-w)m(d, l), \forall d, l$ where $m(d, l)$ is the mass of type-d dealers at location l and $w \in [0, 1]$ represents the fraction of capacity removed. The case $w = 0$ corresponds to the baseline with observed dealer capacity, while $w = 1$ represents a market with no dealers. The average buyer and seller value is the average of V_B and V_S weighted by μ_B and μ_S. The average dealer value is the average of W_1 and W_0 weighted by μ_1 and μ_0. The average retail price is weighted by the corresponding transaction volumes in equilibrium. Total agent stocks are the cross-sectionally aggregated μ_B, μ_S, and μ_1 in equilibrium, respectively. The average buyer's time on the market is the expected time needed to buy a car from any source, averaged over μ_B; the average seller's time on the market is the expected time to sell a car either to a buyer or to a dealer, averaged over μ_S; and the occupied dealer's time on the market is the expected time to sell a car to a buyer, averaged over μ_1. Quality selection plots the respective shares of high-quality cars held by all sellers μ_S in the market and by all occupied dealers μ_1.

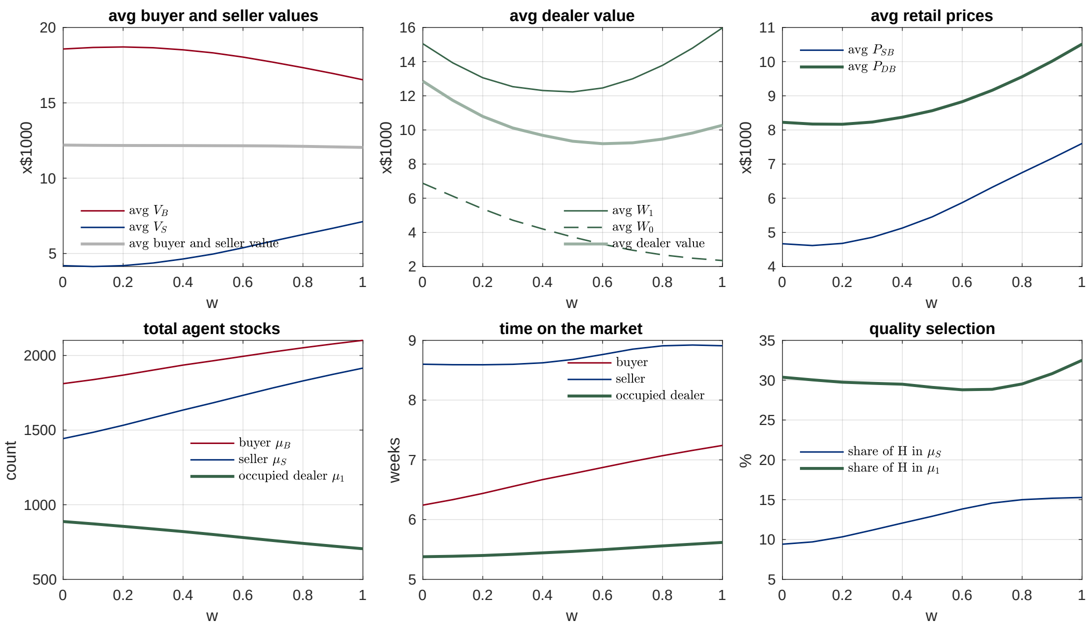

Figure A.5: Effect of Reducing Dealers' Bargaining Power

Note: The x-axis represents the weight $w \in [0, 1$ of dealers' bargaining weights in the counterfactual scenario, such that $\theta_{DB}^{\text{counterfactual}}(d) = (1-w)\theta_{DB}(d) + w\lambda_{SB}$. The polar case $w = 0$ represents the equilibrium benchmark; whereas the other polar case $w = 1$ represents the case where all dealer types' bargaining power are equal to the one of sellers. We interpret $\theta_{DB}(d) - \theta_{SB}$ as type-d dealer's *bargaining power advantage* comparing to sellers. As w increases from 0 to 1, type d dealer becomes less advantageous. See the note of the previous figure for the constructions of outcome variables of interest.

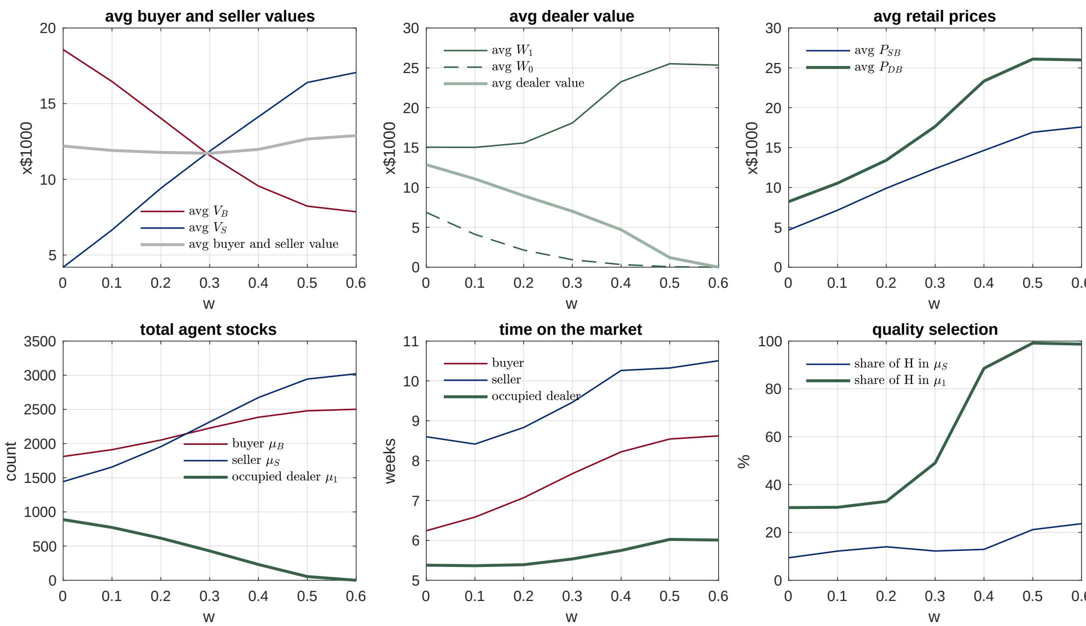

Figure A.6: Effect of Reducing Dealers' Matching Efficiency

Note: The x-axis represents the weight w of dealers' matching efficiency in the counterfactual scenario, such that $\lambda_{DB}^{i,\text{counterfactual}}(d;w) = w\lambda_{SB}^{i} + (1-w)\lambda_{DB}^{i}(d)$, for $i = 0$ (within-county trade) and 1 (across-county trade). As before, we interpret $\lambda_{DB}^{i}(d) - \lambda_{SB}^{i}$ as type-d dealer's matching advantage, and as w increases, type-d dealer becomes less advantageous. See the note of the previous figure for the constructions of outcome variables of interest.