

The Effect of Electronic Commerce on Geographic Trade and Price Variance in a Business-to-Business Market

By: Eric Overby and Chris Forman, Georgia Institute of Technology

Abstract: Imbalances in supply and demand often cause the price for the same good to vary across geographic locations. Economic theory suggests that if the price differential is greater than the cost of transporting the good between locations, then buyers will shift demand from high-price locations to low-price locations, while sellers will shift supply from low-price locations to high-price locations. This should make prices more uniform and cause the overall market to adhere more closely to the “law of one price.” However, this assumes that traders have the information necessary to shift their supply/demand in an optimal way. We investigate this using data on over 2 million transactions in the wholesale used vehicle market from 2003 to 2008. This market has traditionally consisted of a set of non-integrated regional markets centered on market facilities located throughout the United States. Supply / demand imbalances and frictions associated with trading across distance created significant geographic price variance for generally equivalent vehicles. During our sample period, the percentage of transactions conducted electronically in this market rose from approximately 0% to approximately 20%. We argue that the electronic channel reduces buyers’ information search costs and show that buyers are more sensitive to price and less sensitive to distance when purchasing via the electronic channel than via the traditional physical channel. This causes buyers to be more likely to shift demand away from a nearby facility where prices are high to a more remote facility where prices are low. We show that these “cross-facility” demand shifts have led to a 25% reduction in geographic price variance during the time frame of our sample. We also show that sellers are reacting to these market shifts by becoming less strategic about vehicle distribution, given that vehicles are increasingly likely to fetch a similar price regardless of where they are sold.

Keywords: electronic commerce, markets, price dispersion, variance, wholesale automotive, auctions, buyer reach, search costs, choice model.

JEL Classifications: C23, C25, D44, D83, L62, R12

Acknowledgements: We thank the NET Institute (www.NETinst.org) for financial support of this research.

1.0 Introduction

Supply and demand forces often cause the price for the same good to vary across geographic locations. Economic theory suggests that if the price differential is greater than the cost of transporting the good between locations, then buyers will shift demand from high-price locations to low-price locations, while sellers will shift supply from low-price locations to high-price locations. These shifts should lead to a market in which price of a good does not differ between any two locations by more than the cost of transport between them. In other words, the market should obey the “law of one price.” However, this expectation is based on the assumption that traders have the information necessary to shift their supply/demand in an optimal way. If this information is costly to acquire, then there is no reason to expect markets to obey the “law of one price.” As such, information search costs have been offered as an explanation for the repeated violation of the “law of one price” in many markets (see Baye et al., 2006 for a review). In his seminal paper on this topic, Stigler (1961) noted that the cost of searching for and acting upon market information will increase with the geographic size and dispersion of the market. This leads to imbalances in supply and demand across locations, creating geographic segmentation within markets.

Because electronic commerce reduces search costs by making information available at the click of a mouse (Brynjolfsson & Smith, 2000), electronic channels have the potential to break down many of the geographic barriers that have traditionally segmented markets. The elimination of these barriers can have dramatic effects for buyer and seller behavior and for overall market efficiency. For example, electronic commerce lowers buyer search costs (Bakos, 1997), which makes it easier for buyers to collect price information from regions outside their local geographies. If buyers discover that prices are lower in a geographic region outside their own, then they can use electronic channels to shift their purchasing to the lower-price region. This shifting of demand should cause prices across regions to become less variable, thereby more closely reflecting the “law of one price.” As prices become more uniform, sellers should become less strategic about how they distribute products across locations, because products will fetch a similar price regardless of where they are sold.

We use the United States wholesale used vehicle market as the context in which to study whether and how buyers use electronic channels to shift their demand geographically, the effect this has on geographic price variance, and how sellers respond to these changes in the market. This market is well-suited to our analysis for several reasons. First, the wholesale used vehicle market has traditionally consisted of a set of non-integrated regional markets centered on market facilities located throughout the U.S. at which transactions were conducted based on physical collocation of buyers, sellers, and vehicles. This geographic segmentation and the associated imbalances in supply and demand caused prices for generally equivalent vehicles to vary across locations by more than the cost of transport. Second, the percentage of trades conducted electronically rose steadily over our sample period, growing from approximately 0% in

2003 to approximately 20% by mid-2008. This evolution permits a longitudinal analysis of whether and how the increase in electronic trading affected buyer behavior, seller behavior, and the variance of prices across market locations. An interesting feature of the empirical context is that the electronic channel is specific to buyers. Sellers do not use the electronic channel, although buyers' use of the channel may prompt behavioral changes by the sellers. This feature permits us to investigate how the electronic channel has affected demand and how supply has adjusted. Third, the wholesale used vehicle market is representative of other geographically segmented markets such as those for agricultural crops, livestock, seafood, fuels, building materials, and heavy machinery (e.g., Aker, 2010; Diekmann et al., 2008; Jensen, 2007). Our findings about whether and how electronic trading has affected trader behavior and market efficiency in the wholesale used vehicle market should help us understand similar phenomena in other markets.

We have detailed transaction data for over 2 million vehicles sold in the market from 2003 to 2008. Among other variables, we observe the transaction price, the make/model of each vehicle, the location of each vehicle (there are over 80 market locations across the U.S. at which vehicles are sold), the location of each buyer based on his zip code, and whether the buyer purchased the vehicle via the traditional physical channel or via the electronic channel. We pose four research questions. First, do buyers use the electronic channel to extend their purchasing reach to more remote locations? We find that they do; buyers are approximately 29% less sensitive to distance when purchasing via the electronic channel than via the physical channel. Second, do buyers use the reach of the electronic channel to shift their demand to other locations to take advantage of lower prices? We find that they do; buyers are approximately 194% more sensitive to price when purchasing via the electronic channel than via the physical channel. Third, are these demand shifts associated with reduced price variance across locations? We find that they are; a one standard deviation increase in the number of vehicles of a given model (e.g., Ford Ranger, Toyota Camry) that buyers purchase from remote locations is associated with a 13% decrease in price variance for vehicles of that model across the U.S. Fourth, how do sellers react to the reduction in geographic price variance in the market? Sellers choose the location at which to sell vehicles, and we find that sellers placed less weight on recent prices at a location when making distribution decisions as electronic trading increased. We conclude that this is because distributing vehicles to locations where prices have recently been high becomes less valuable as buyers become increasingly likely to use the electronic channel to shift demand out of those locations.

The study draws upon and contributes to two main research streams. The first, second, and fourth research questions relate primarily to the research stream on how electronic commerce affects geographic trade (e.g., Blum & Goldfarb, 2006; Hortacsu et al., 2009). The third research question relates primarily

to the research stream on how electronic commerce affects price dispersion¹ (e.g., Chellappa et al., 2011; Clay et al., 2002; Clemons et al., 2002). The second and third research questions create the linkage between these two research streams by asking how the shifting of demand across geographic locations facilitated by electronic commerce affects price variance across those locations.

We contribute to both research streams in several ways. First, prior empirical studies have examined whether reduced buyer search costs in electronic channels lead to lower price dispersion (e.g., Brown & Goolsbee, 2002). Although many of these studies have demonstrated significant changes in price dispersion, they typically have not examined the micro-level buyer behavior that leads to that outcome. We use a discrete choice model using individual buyer transactions across both physical and electronic channels to examine the behavioral mechanism by which reduced buyer search costs lead to lower price dispersion.² Our results show that buyers are more sensitive to price and less sensitive to distance when using the electronic channel compared to the physical channel. As a result, they shift purchases from nearby locations where prices are relatively high to remote locations where prices are relatively low. These cross-location demand shifts -- which have become more common over time -- represent the mechanism through which changes in buyer behavior have led to lower price dispersion across geographic locations. Observing this mechanism is critical for continued empirical research about electronic channels and price dispersion because different assumptions about the mechanism can result in more or less price dispersion when modeled analytically (Baye et al., 2006). We build upon the results of the choice model by using fixed effects panel regression to attempt to quantify the relationship between cross-location demand shifts and geographic price dispersion.

Second, existing empirical research on how electronic channels influence price dispersion has generally ignored the geographic location of products as a factor in influencing their price and therefore their price dispersion. This is because location has been irrelevant for the types of products that have typically been studied (e.g., books, consumer electronics, and tickets), because the cost of shipping the product to the buyer -- which is a component of the overall price paid by the buyer -- is the same regardless of the product's location. However, shipping costs vary significantly with location for products such as automobiles, agricultural commodities, fuels, and metals, the dollar value of trade for which is substantially larger than that for goods more commonly studied.³ For these products, the distance between buyers and products influences the prices that buyers pay, which in turn affects the dispersion of prices

¹ We use the terms "price variance" and "price dispersion" interchangeably.

² For examples of how reductions in seller search costs influence price dispersion in the context of commodity food markets in developing countries, see Aker (2010) and Jensen (2007).

³ For example, U.S. retail sales for automobiles were approximately 30 times larger than that for books and CD's in 2002. Source: U.S. Census, http://www.census.gov/econ/census02/data/us/US000_44.HTM.

for those products at different locations. We examine the location of products and buyers and find that the distance between them influences when buyers use physical vs. electronic channels and weakens the relationship between cross-location demand shifts and price dispersion.

Third, most of the existing research on price dispersion has studied fixed price environments by analyzing prices posted by sellers. By contrast, we study an environment in which prices are determined by auction. This allows us to observe the actual transacted prices in the market, rather than just the posted prices, which is unusual in this stream (but see Ghose and Yao (2011) for a recent exception).

The paper proceeds as follows. Section 2 presents the empirical context and data. Section 3 discusses how the steady increase in electronic trading between 2003 and 2008 might have influenced buyer/seller behavior and geographic price variance. Section 4 presents our empirical models and analysis. Section 5 summarizes the results, limitations, and contributions of the study.

2.0 Empirical Context

The empirical context of the study is the wholesale used vehicle market. This is a business-to-business market in which buyers and sellers trade used vehicles. The buyers are used car dealers who purchase vehicles in the wholesale market for resale to retail customers. Used car dealers procure approximately 35% of their used vehicle inventory via the wholesale market (Source: NADA Data 2009 (www.nada.org/nadadata), page 10.) The sellers are either other dealers or institutional sellers such as rental car companies and the captive finance arms of automotive manufacturers. The main reason for a dealer to sell in the wholesale market is if he does not wish to (or cannot) sell a vehicle in the retail market. In this case, he will sell the vehicle wholesale to another dealer who will retail the vehicle. Institutional sellers sell in the wholesale market because they often lack retail operations and because the wholesale market is a highly liquid environment for selling multiple vehicles. Approximately 9 million vehicles are exchanged in the market each year (Source: National Auto Auction Association (www.naaa.com))

There are several intermediaries that provide services in the market, including aggregating buyers and sellers, providing storage while vehicles are pending sale (referred to as “marshalling”), and providing transaction assurance. The intermediaries also operate physical market facilities at which transactions are conducted. Market facilities are located throughout the U.S. as well as the world, although our analysis is specific to the U.S. Sellers transport vehicles to market facilities where buyers purchase them via an auction process. Each vehicle that is auctioned is driven – one at a time – into a warehouse-type building into the midst of a group of buyers. A human auctioneer solicits bids from the buyers in an ascending, open outcry format, i.e., a traditional English auction. Once the auctioneer can solicit no additional bids, the seller indicates to the auctioneer whether he will accept the high bid. The vehicle is then driven away,

and the next vehicle is driven into place and the process repeats. It is common for vehicles of the same model j (e.g., Toyota Camry) to be auctioned one after another. It is also common for auctions to be conducted at multiple facilities concurrently. After purchasing a vehicle, the buyer is responsible for transporting the vehicle to his dealership. The cost of transportation is non-trivial and increases with distance.

2.1 The Electronic “Webcast” Channel: The physical process described above remains the predominant method by which vehicles are exchanged in the wholesale market in the United States. However, an increasing percentage of transactions are conducted electronically. The most popular electronic channel is the webcast channel. The webcast channel consists of an Internet browser-based application that streams live audio and video of the physical auctions. Buyers can use the application to place bids on vehicles in competition with the buyers who are physically present at the facility. This has given buyers a choice for how to participate in the bidding for a vehicle: they can either participate physically in the traditional fashion or electronically using the webcast channel.

It is worth highlighting two points about the webcast channel. First, the webcast channel does not affect the price discovery mechanism. The auctioneer solicits bids for each vehicle in an ascending fashion; this price discovery process is not affected by the channel buyers use to place bids. Second, the webcast channel is specific to buyers; sellers do not use it. Sellers present their vehicles in the same fashion -- having them driven through the physical market facility -- regardless of whether buyers are using the physical or the electronic “webcast” channels to place bids.⁴

2.2 Data Description: The data we use to examine our research questions were provided by an intermediary in the wholesale used vehicle market that operates over 80 physical market facilities in the continental U.S., all of which are equipped with the webcast technology. The data consist of all vehicles with between 15,000 and 21,000 miles that were auctioned (both successfully and unsuccessfully) at those facilities between January 2003 and June 2008. The mileage filter reduces heterogeneity in vehicle condition, so that prices for vehicles of the same model across facilities may be more validly compared. The low mileage of vehicles in the sample also increases the likelihood that vehicles are of predictable

⁴ The webcast channel is not the only method by which buyers can purchase vehicles electronically. There are also stand-alone electronic markets that operate in the industry. The key distinction between the two is that the webcast channel augments the physical market, while the electronic markets are separate from it. We limit our analysis to the webcast channel for two reasons. First, the vast majority of electronic transactions are conducted via the webcast channel. In our data, webcast transactions outnumber the other electronic transactions over 8 to 1. Second, the stand-alone electronic markets typically offer a fixed price option to buyers, and there is no human auctioneer to solicit bids. This creates differences in the price discovery mechanism that could confound our results. For simplicity, all references to the electronic channel are specific to the webcast channel.

quality, such that they may be traded effectively via either the physical or electronic channels (Overby & Jap, 2009). The data contain 3,591,443 auctions, 2,340,357 of which resulted in a sale. Of these, 2,059,832 were purchased by buyers using the physical channel (88%) and 280,525 were purchased by buyers using the webcast channel (12%). The percentage of vehicles traded electronically increased from just over 0% to approximately 20% over this time period. The data contain 74,917 unique buyer ID's. According to the 2007 U.S. Economic Census, there are 50,808 automobile dealers in the United States; the number of buyer ID's is higher because many dealers have more than one employee authorized to purchase in the wholesale market. The number of buyer ID's in the data and the likelihood that most used car dealers purchase at least some vehicles with between 15,000 and 21,000 miles increase our confidence that our sample is representative of the buyer population in this market. Table 1 describes the variables in the data. We used the Facility Zip Code and Buyer Zip Code variables to calculate the distance (in miles) between facilities and between buyers and facilities.

Table 1: Variables and descriptive statistics.

Variable	Description	Descriptive Statistics
Facility ID	Denotes the market facility where the vehicle was located.	There are 81 facility ID's.
Facility Zip Code	Zip code of each market facility.	n/a
Buyer ID	Denotes the buyer who purchased a vehicle.	There are 74,917 buyer ID's.
Buyer Zip Code	Zip code of each buyer.	n/a
Sold?	Dummy variable for whether the vehicle was sold (1) or not (0).	Mean: 0.65
Seller ID	Denotes the seller of each vehicle.	There are 28,791 seller IDs.
Webcast?	Dummy variable indicating whether the vehicle was purchased via the webcast (1) or physical (0) channels.	Mean: 0.12
Auction Date	Date each vehicle was auctioned.	From 1/1/2003 to 6/30/2008.
Vehicle Model	Model of each vehicle (e.g., Ford Focus, Nissan Maxima.)	There are 834 vehicle models.
Price	Sales price of each vehicle.	Mean: 14,819; St. Dev.: 6,834
Valuation	Market value estimate for each vehicle on the Auction Date. Calculated by the intermediary based on transactions for similar vehicles in the 30 days prior to the Auction Date.	Mean: 15,005; St. Dev.: 6,768
Normalized Price	Price divided by Valuation.	Mean: 0.99; St. Dev.: 0.09
Condition Grade	Grade (0-5 scale, measured in 0.1 increments) assigned by the intermediary to indicate each vehicle's wear and tear. Higher grades indicate better condition.	Mean: 3.29; St. Dev.: 0.70
Mileage	Odometer reading of each vehicle.	Mean: 18,070; St. Dev.: 1,734
Age	Date the vehicle was auctioned minus January 1 of the vehicle's year. Measured in days. May be negative, e.g., if a 2008 model vehicle was auctioned in December 2007.	Mean: 545; St. Dev.: 714

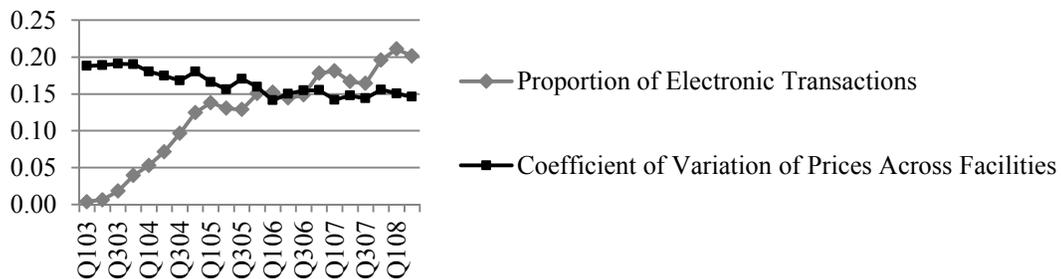
2.2.1: Assigning Buyers to Local Facilities: In order to test some of our hypotheses, it was necessary to assign each buyer to a facility in their geographic area, referred to as the buyer's "local" facility. We did this as follows. First, we calculated the average distance each buyer traveled to make purchases via the traditional physical channel. This statistic was 125 miles (s.d. 185.) We then determined which facilities were within 125 miles of each buyer. If a buyer had exactly one facility within a 125 mile radius,

we assigned him to that facility. If a buyer had more than one facility within this radius, we calculated the facility from which he made the most physical purchases and assigned him to that facility. We assigned buyers who were located farther than 125 miles from the nearest facility to the facility closest to them. This procedure yields the same assignments as a procedure in which each buyer is assigned to the closest facility, with the exception of those buyers who are close to multiple facilities. We used this procedure so that buyers were assigned to nearby facilities that they actually visited, rather than to whichever facility was simply the closest.

3.0 Geographic Price Variance and Buyer/Seller Behavior

3.1 Geographic Price Variance: We first present some summary statistics. We calculated the average price of vehicles of each make/model (e.g., Toyota Camry, Ford Escape) at each facility in each quarter-year. We then calculated the variation in those average prices across facilities (using the coefficient of variation) in each quarter-year. We label this CV_{jt} , where j indexes the make/model and t indexes the quarter-year. For example, assume that the average price for Toyota Camry's in Q1 2003 was \$12,000 at the Charlotte facility, \$14,000 at the Orlando facility, and \$10,000 at the Nashville facility. In this case, $CV_{jt} = (2,000 / 12,000) = 0.16$. We then averaged CV_{jt} across all make/models j in each quarter-year to get a measure of the overall variance of prices in the market. We plotted this statistic for each quarter-year. This is shown in Figure 1, along with the percentage of electronic transactions. Geographic price dispersion declines as electronic transaction activity increases. The main goal of our analysis is to study this relationship.

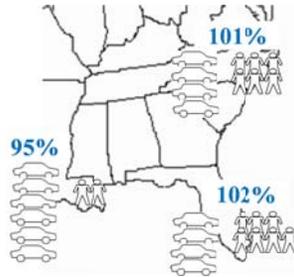
Figure 1: Proportion of electronic transactions and coefficient of variation of prices across facilities per quarter-year.



A key reason why geographic price variance exists in the market is imbalances in supply and demand across facilities. The supply at a facility is determined by the number of vehicles being auctioned at the facility. The demand at a facility is determined by the number of buyers purchasing at the facility. Figure 2 provides an illustration. The graphic shows three market facilities: Charlotte, Miami, and New Orleans. Supply is illustrated by the vehicle icons; demand is illustrated by the person icons. The blue numbers represent (hypothetical) average vehicle prices as a percentage of market value at each facility on a given

day t . In the graphic, high supply and low demand cause average prices to be low in New Orleans, while low supply and high demand cause average prices to be high in Charlotte and Miami.

Figure 2: Hypothetical illustration of supply / demand imbalances and the corresponding geographic price variance.



Notes: Supply is illustrated by the vehicle icons; demand is illustrated by the person icons. The numbers represent (hypothetical) average vehicle prices as a percentage of market value at each facility on a given day t .

The supply at each facility depends how sellers choose to distribute their vehicles to the different facilities. The demand at each facility depends on how buyers choose the facilities at which to purchase vehicles. It is important to note that seller distribution decisions are made prior to buyer purchase decisions. Sellers distribute vehicles to facilities, and the supply of vehicles available at a facility is typically posted at least a day before the vehicles are auctioned. Buyers can access this information to help choose a facility at which to purchase.

3.2 How Sellers and Buyers Exploit Geographic Price Variance: We consider seller and buyer behavior before and after the electronic channel became available.

3.2.1 Pre-Electronic Channel: We assume that sellers and buyers are aware of the geographic price variance and seek to exploit it. We expect sellers to exploit geographic price variance by distributing vehicles to facilities where prices have recently been high. If sellers do this successfully, then supply and demand will become better balanced and geographic price variance will decrease. For example, the left panel of Figure 3 illustrates that if sellers shift supply from New Orleans to Charlotte, prices in New Orleans should rise and prices in Charlotte should fall, thereby reducing geographic price variance. However, the extent to which sellers do this is limited by two factors: a) their ability to forecast prices at each facility (“seller limitation #1”), and b) the cost of transporting vehicles to each facility (“seller limitation #2”).

Figure 3: Hypothetical illustration of how a shift in supply or demand might affect geographic price variance.



We expect buyers to exploit geographic price variance by choosing to purchase vehicles at facilities where prices are expected to be low. If buyers do this successfully, then the supply and demand will become better balanced and geographic price variance will decrease. For example, the right panel of Figure 3 illustrates that if buyers shift from purchasing in Charlotte to purchasing in New Orleans, prices in Charlotte will fall and prices in New Orleans will rise, thereby reducing geographic price variance. However, the extent to which buyers do this is limited by: a) their ability to forecast prices at each facility (“buyer limitation #1”), b) the costs of transporting vehicles from each facility (“buyer limitation #2”), and c) the costs of their bidding at each facility, which traditionally has required physical attendance at a facility (“buyer limitation #3.”)

3.2.2 Post-Electronic Channel: As discussed above, the electronic “webcast” channel provides live, streaming audio and video of the physical auctions occurring at a facility. Buyers can use the electronic channel to place bids in competition with buyers who are physically at the facility. Thus, the electronic channel allows buyers to purchase vehicles at facilities without having to travel. We reiterate that the electronic channel is specific to buyers; sellers do not use it.

We expect the electronic channel to improve buyers’ ability to exploit geographic price variance, but not sellers’ ability. Specifically, the electronic channel should mitigate buyer limitations #1 and #3. Regarding buyer limitation #1, buyers can use the electronic channel to check prices at each facility by opening multiple browser windows to observe the bidding activity at multiple facilities. This helps buyers forecast prices for yet-to-be-auctioned vehicles at each facility. Regarding buyer limitation #3, buyers who use the electronic channel avoid the costs of physical attendance at a facility, which includes travel cost and opportunity cost. Thus, their cost of bidding is substantially reduced. Conversely, the electronic channel does not improve sellers’ ability to exploit geographic price variance; i.e., it does not mitigate seller limitations #1 or #2. First, it does not help sellers forecast prices, because sellers must distribute vehicles to facilities before the auctions begin. This is in contrast to buyers, who can use the electronic channel to shift their demand across facilities after the auctions have begun. Second, the electronic channel does not affect the cost of vehicle transportation, either for sellers or buyers.

Thus, we posit that buyers use the electronic channel to re-distribute their demand from facilities where prices are high to facilities where prices are low. This should decrease geographic price variance. We posit that sellers react to how buyers are using the electronic channel by placing less weight on recent prices at a facility when making distribution decisions. This is because distributing vehicles to facilities where prices have recently been high becomes less valuable as buyers become increasingly likely to shift demand out of those facilities. Thus, we posit that buyers are using the electronic channel in a way that reduces geographic price variance and that sellers are reacting to buyer behavior in a rational manner.⁵

4.0 Models and Results

In this section, we investigate potential changes to buyer and seller behavior attributable to the electronic channel. We also consider whether these changes might explain the reduction in geographic price variance illustrated in Figure 1.

4.1 Buyer Behavior: As noted in Section 3, the relevant buyer behavior for our analysis is how buyers choose the facilities at which to purchase vehicles. We hypothesize that buyers will make different choices depending on whether they are using the physical or electronic channel (which is itself a choice), with buyers' facility choices less affected by distance but more affected by price when using the electronic channel. We developed a discrete choice model to investigate this. We assume that a buyer is purchasing a vehicle(s) of a particular make/model j (e.g., Ford Taurus) on day t . He can purchase that vehicle at any facility at which vehicles of make/model j are auctioned on day t . Each facility provides differential utility to the buyer based on prices at the facility, the distance to the facility, the supply of vehicles available at the facility, and the condition of the vehicles available at the facility. We measured price ($NormPrice_{jkt}$) as the average normalized price (normalized price = price / valuation; see Table 1)

⁵ It is possible to model this formally via a sequential game in which a seller moves first by choosing a facility at which to auction a vehicle. The two main considerations for the seller are the distance to the facility (due to transportation costs) and historical prices at the facility (to try to exploit geographic price variance by selling where demand is most likely to be high). A buyer moves second and is of one of two types: he either uses the webcast channel or he does not. If the buyer uses the webcast channel, then he has the option to bid at his "local" facility or to shift his bid to a remote facility, which he will do if the price difference between facilities is sufficiently high. Let μ denote the probability that the buyer uses the webcast channel. As μ increases, it becomes more likely that a given buyer will shift his bid away from his local facility if prices there are high. This demand-shifting thwarts (to some extent) the seller's attempt to achieve above average prices; recognizing this, she will place less weight on historical prices (and more on other factors) when making distribution decisions as μ increases.

for vehicles of model j sold at facility k on day t .⁶ $NormPrice_{jkt}$ is unobserved when none of the vehicles of model j auctioned at facility k on day t are sold. This occurs 11.8% of the time. We set $NormPrice_{jkt}$ equal to 1 when unobserved for two reasons. First, a vehicle g 's valuation ($Valuation_g$) is highly correlated with its price ($Price_g$); a simple regression of $Price_g$ on $Valuation_g$ for sold vehicles yields a valuation coefficient of 0.99 and an R^2 of 0.96. Thus, using an imputed value = 1 (meaning that $Price_g = Valuation_g$) is a reasonable strategy. Second, we assume that buyers impute prices in this same manner, i.e., when price information is not available, they assume that prices, if observed, would be equal to valuations.⁷ We measured distance ($Distance_{ik}$) as the number of miles between buyer i 's zip code and facility k 's zip code. We measured supply ($Supply_{jkt}$) as the number of vehicles of model j auctioned at facility k on day t . We measured average condition ($AvgCondition_{jkt}$) as the average condition grade assigned to vehicles of model j auctioned at facility k on day t . The condition grade of vehicles ($ConditionGrade_g$) is not recorded for approximately one-third of the vehicles. We imputed $ConditionGrade_g$ for those vehicles based on: a) the average condition grade for vehicles of the same model j sold at the same facility k over the prior 21 days, b) the vehicle's mileage, and c) the vehicle's age. Specifically, we regressed $ConditionGrade_g$ (when observed) on these three variables and used the resulting coefficients to impute $ConditionGrade_g$ when not observed. We also included alternative-specific constants in the choice model to capture the latent utility of each facility. Table 2 provides descriptive statistics for these and other variables described later in this section.

Table 2: Descriptive statistics for variables in the buyer choice model.

Variable	Description	Mean (St. Dev.)
$NormPrice_{jkt}$	Average normalized price (price / valuation) for vehicles of model j sold at facility k on day t .	0.99 (0.13)
$Distance_{ik}$	Average distance between buyer i and each facility k in buyer i 's choice set.	451.99 (348.77)
$Supply_{jkt}$	Number of vehicles of model j auctioned at facility k on day t .	6.32 (11.59)
Avg $Condition_{jkt}$	Average condition grade of vehicles of model j auctioned at facility k on day t .	3.23 (0.59)
Electronic_ $Propensity_{it}$	Proportion of buyer i 's 10 (or fewer if necessary) purchases prior to day t made via the webcast channel.	0.08 (0.22)

⁶ We also included in the model the average normalized price for vehicles of model j sold at facility k for the 21 days prior to day t ($RecentNormPrice_{jkt}$) to account for the possibility that buyers consider recent historical prices at a facility when making decisions. This variable was not significant at the 10% level.

⁷ Price (and other variables) is commonly imputed in choice models, because the price of non-chosen alternatives in a choice set is often unobserved. This often results in choice models in which price is imputed $[(n-1)/n]*100$ percent of the time, where n is the number of alternatives (e.g., see Bucklin et al., 2008, p. 480; Chiou, 2009, p. 292). Our level of imputation is much less because we usually observe average prices of non-chosen alternatives.

The complete buyer choice set is quite large, as is the number of parameters given the inclusion of the alternative-specific constants for each facility. This size makes model estimation unstable and convergence difficult, particularly for choice model formulations other than the conditional logit. To achieve an estimable model, we took a geographic subset of the sample by analyzing only the purchases made by buyers local to the facilities in the western U.S., which consists of the facilities in Arizona, California, Colorado, Nevada, New Mexico, Oregon, Utah, and Washington. We also limited the choice set to facilities within this region. We estimated the model using this sub-sample for two reasons. First, it is large enough geographically to allow us to examine whether buyers' sensitivity to distance is affected by the electronic channel. Second, 93.7% of purchases by buyers local to facilities within this region are from facilities within this region, which allows us to consider it a microcosm of the entire market. The filtered choice data set consists of 313,252 choices and 18 facilities.

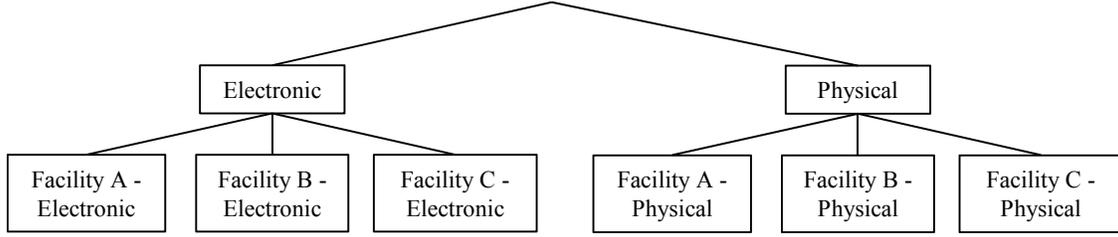
As a first step in our analysis, we divided the sample into observations in which the buyer purchased via the physical channel and those in which the buyer purchased via the electronic channel. We then estimated conditional logit models for the two sub-samples. This provides estimates of how each variable influences buyers' location choices, conditional on the buyer having already chosen a channel. Results appear in Table 3.

Table 3: Results of buyer facility choice model conditional on having chosen a channel.

	Conditional on Physical Channel	Conditional on Electronic Channel
	Coefficient	Coefficient
β_1 : NormPrice	-0.415 (0.039) ***	-0.848 (0.153) ***
β_2 : Distance	-0.007 (0.000) ***	-0.005 (0.000) ***
β_3 : Supply	0.067 (0.001) ***	0.143 (0.002) ***
β_4 : AvgCondition	0.120 (0.009) ***	0.216 (0.029) ***
$\beta_{5,k}$: Facility constants	included	included
n (number of choices)	207,484	20,236
Log likelihood	-82407.82	-8258.58
Robust standard errors in parentheses.	*** p < 0.01	

As a next step, we included the channel choice in the specification by using a nested logit model, which essentially allows us to estimate the two sub-samples shown in Table 3 simultaneously. The nests represent the buyer's option to purchase from a facility using either the physical or electronic channels. This allows us to model the buyer's choice of facility *and* channel, although our focus remains on how buyers choose a facility conditional on having chosen a channel. The nesting structure is illustrated in Figure 4.

Figure 4: Illustration of the nesting structure for the nested logit model.



In the nested logit specification, each facility appears twice in a buyer's choice set: once in the physical channel nest and once in the electronic channel nest. For example, assume that vehicles of make/model j are being auctioned in Charlotte and New Orleans on day t . Because buyer i can purchase using either the physical or electronic channels, he has four alternatives for purchasing this vehicle: physically in Charlotte, physically in New Orleans, electronically in Charlotte, and electronically in New Orleans. This means that a change in a variable such as $NormPrice_{jkt}$ affects both the physical and electronic alternatives for facility k in a choice set. A feature of the nested logit is that it increases the probability that buyers substitute within nests (i.e., channels in our case) rather than across nests (Train, 2009). In our case, the nested logit increases the probability that a change in a variable such as an increase in $NormPrice_{jkt}$ will cause a buyer to buy from a different facility using his preferred channel (i.e., to substitute within nests) rather than to buy from the same facility using the other channel (i.e., to substitute across nests.)

We included in the buyer's utility function a dummy variable for alternatives in the electronic nest ($Electronic_c$) and interacted that with the price, distance, supply, and condition variables to capture how their explanatory power differ across channels. To capture dynamics in buyers' use of the electronic channel vis-à-vis the physical channel, we constructed $Electronic_Propensity_{it}$, which is the proportion of buyer i 's 10 (or fewer if necessary) purchases prior to day t made via the electronic channel. It measures the strength of buyer i 's preference for the electronic channel. We interacted $Electronic_Propensity_{it}$ with $Electronic_c$. This allows the latent utility of the electronic channel to change as a buyer i uses it more (or less). To summarize, we modeled the utility of each alternative (which consists of a channel c / facility k combination) for buyer i purchasing a vehicle(s) of model j on day t as:

$$\begin{aligned}
 U_{ijckt} = & \beta_1 * NormPrice_{jkt} + \beta_2 * Distance_{ik} + \beta_3 * Supply_{jkt} + \beta_4 * AvgCondition_{jkt} + \\
 & \beta_5 * Electronic_c * NormPrice_{jkt} + \beta_6 * Electronic_c * Distance_{ik} + \beta_7 * Electronic_c * Supply_{jkt} + \\
 & \beta_8 * Electronic_c * AvgCondition_{jkt} + \sum_{k=2}^K \beta_{9,k} * Facility(k) + \beta_{10} * Electronic_c + \sum_{k=2}^K \beta_{11,k} * \\
 & Electronic_c * Facility(k) + \beta_{12} * Electronic_c * Electronic_Propensity_{it} + \varepsilon_{ijckt}
 \end{aligned} \tag{1}$$

$\sum_{k=2}^K Facility(k)$ are dummy variables for each facility. Note that the combination of $\beta_{9,k}$, β_{10} , and $\beta_{11,k}$ form the alternative-specific constants for this model. Results appear in Table 4.

Table 4: Results of buyer channel/facility choice model using a nested logit specification.

	Coefficient (Std. Error)
β_1 : NormPrice	-0.388 (0.038) ***
β_2 : Distance	-0.007 (0.000) ***
β_3 : Supply	0.068 (0.001) ***
β_4 : AvgCondition	0.105 (0.008) ***
β_5 : Electronic*NormPrice	-0.752 (0.144) ***
β_6 : Electronic*Distance _i	0.002 (0.000) ***
β_7 : Electronic _c *Supply	0.064 (0.002) ***
β_8 : Electronic _c *AvgCondition	0.239 (0.028) ***
β_{10} : Electronic	-13.872 (0.490) ***
β_{12} : Electronic*Electronic_Propensity	19.262 (0.696) ***
$\beta_{9,k}$: Facility(k)	included
$\beta_{11,k}$: Electronic*Facility(k)	included
Inclusive value (physical nest)	0.424 (0.009)
Inclusive value (electronic nest)	0.273 (0.010)
n	313,252
Log likelihood	-149784.1
*** p < 0.01	

The coefficient for *Electronic_c* ($\beta_{10} = -13.872$) is negative and significant. This indicates that purchasing via the electronic channel generates substantial disutility. However, the positive and significant coefficient for *Electronic_c*Electronic_Propensity_{i,t}* ($\beta_{12} = 19.262$) indicates that the utility of the electronic channel increases with a buyer's propensity to use it. The coefficient for *NormPrice_{j,k,t}* ($\beta_1 = -0.388$) is negative and significant, as is the coefficient for *Electronic_c*NormPrice_{j,k,t}* ($\beta_5 = -0.752$). The combined coefficient ($\beta_1 + \beta_5 = -1.140$) represents the disutility of price when purchasing via the electronic channel, which is 194% greater ($[(\beta_1 + \beta_5)/\beta_1] - 1$, formatted as a percentage) than the disutility of price when purchasing via the physical channel. The coefficient for *Distance_{i,k}* ($\beta_2 = -0.007$) is negative and significant, while the coefficient for *Electronic_c*Distance_{i,k}* ($\beta_6 = 0.002$) is positive and significant. This indicates that the disutility of distance is 29% less when purchasing via the electronic channel. The combined coefficient ($\beta_2 + \beta_6 = -0.005$) is negative and significant. This indicates that buyers still prefer nearby facilities when using the electronic channel, but less so than when using the physical channel.

We summarize these results as follows. A buyer is more likely to be in the electronic “nest” (i.e., to choose the electronic channel) as his experience with the electronic channel grows. When in the electronic “nest”, a buyer is less sensitive to distance when choosing a facility, because he doesn't have to travel to the facility if he is using the electronic channel. When in the electronic “nest”, a buyer is more sensitive to price when choosing a facility, because the electronic channel provides him with better information about average prices across facilities. Put together, these results indicate that buyers are more likely to shift their

demand from nearby facilities where prices are high to more distant facilities where prices are low as their use of the electronic channel increases.

4.2 Seller Behavior: As noted in Section 3, the relevant seller behavior for our analysis is how sellers choose the facilities at which to sell vehicles. Although sellers do not use the electronic channel, we posit that buyers' use of the channel will influence sellers' distribution choices. We examine this via a discrete choice model.

Each auctioned vehicle represents a choice made by the seller to sell that vehicle at a given facility rather than alternative facilities. Each facility provides differential utility to the seller based on factors such as recent prices at the facility, the recent supply of similar vehicles, and the historical propensity of the seller to sell vehicles at the facility. Specifically, we model the utility of facility k for seller s selling a vehicle of make/model j at time t as:

$$U_{jkst} = \beta_1 * \text{RecentNormPrice}_{jkt} + \beta_2 * \text{PctBuyerElectronic}_{kt} + \beta_3 * \text{DistributionPropensity}_{jkst} + \beta_4 * \text{RecentSupply}_{jkt} + \beta_5 * \text{RecentNormPrice}_{jkt} * \text{PctBuyerElectronic}_{kt} + \beta_6 * \text{DistributionPropensity}_{jkst} * \text{PctBuyerElectronic}_{kt} + \sum_{k=2}^K \beta_{7,k} * \text{Facility}(k) + \varepsilon_{jkst} \quad (2)$$

$\text{RecentNormPrice}_{jkt}$ is the average normalized price (normalized price = price / valuation; see Table 1) for vehicles of model j sold at facility k in the 3 weeks prior to week t . This accounts for the (dis)utility of recent (low) high prices at a facility. $\text{PctBuyerElectronic}_{kt}$ is the percentage of purchases by buyers assigned to facility k that were made in the electronic channel in week t . This influences the utility of facility k for a seller because it represents whether buyers local to facility k are a "captive" buying group or are likely to shift their demand out of facility k . $\text{DistributionPropensity}_{jkst}$ is the number of vehicles of model j that seller s sold at facility k divided by the number of vehicles of model j that seller s sold at all facilities, both in the 52 weeks (or fewer for observations in year 2003) prior to week t . This captures unobserved factors that influence seller distribution decisions, including habit. $\text{RecentSupply}_{jkt}$ is the number of vehicles of model j sold by any seller at facility k in the 3 weeks prior to day t . $\sum_{k=2}^K \text{Facility}(k)$ are dummy variables for each facility and serve as alternative-specific constants in the model. Table 5 provides descriptive statistics.

Another variable that affects the seller's utility is the distance between a vehicle's location prior to entering the market and each facility k . We cannot include this variable in the model because the vehicle's location prior to entering the market is unobserved. However, we believe that some of this effect is captured in $\text{DistributionPropensity}_{jkst}$. This is because one of the reasons that a seller is likely to have a high (low) propensity to sell vehicles at a facility is because the facility is close (far) to the vehicle.

Table 5: Descriptive statistics for variables in the seller choice model.

Variable	Description	Mean (St. Dev.)
RecentNormPrice _{jkt}	Average normalized price (price / valuation) for vehicles of model j sold at facility k in the 3 weeks prior to day t .	0.98 (0.09)
PctBuyerElectronic _{kt}	Percentage of purchases by buyers assigned to facility k that were made in the electronic channel in week t .	0.10 (0.12)
DistributionPropensity _{jkst}	Number of vehicles of model j that seller s sold at facility k divided by the number of vehicles of model j that seller s sold at all facilities, both in the 52 weeks (or fewer for observations in year 2003) prior to week t .	0.03 (0.15)
RecentSupply _{jkt}	Number of vehicles of model j sold by any seller at facility k in the 3 weeks prior to day t .	5.93 (16.75)

We estimated the model using a conditional logit specification. Results appear in Table 6.

Table 6: Results of seller distribution choice model.

	Coef.	Std. Err	z-stat
β_1 : RecentNormPrice _{jkt}	0.114	0.036	3.14 ***
β_2 : PctBuyerElectronic _{kt}	-0.495	0.331	-1.50
β_3 : RecentNormPrice _{jkt} *PctBuyerElectronic _{kt}	-1.155	0.333	-3.47 ***
β_4 : RecentSupply _{jkt}	0.002	0.000	24.44 ***
β_5 : DistributionPropensity _{jkst}	8.718	0.028	316.39 ***
β_6 : DistributionPropensity _{jkst} *PctBuyerElectronic _{kt}	6.152	0.250	24.56 ***
$B_{7,j}$: Facility(k)	included		
n (number of choices)	509,497		
Log likelihood	-606,901		
*** p < 0.01			

The coefficient for *RecentNormPrice_c* ($\beta_1 = 0.110$) is positive and significant, and the coefficient for *RecentNormPrice_{jkt}*PctBuyerElectronic_{kt}* ($\beta_3 = -1.093$) is negative and significant. This shows that recent prices provide positive utility to sellers when making distribution decisions but that this utility becomes weaker with increased use of the electronic channel by the buyers local to that facility. This suggests that the demand shifting enabled by the electronic channel is causing sellers to become less strategic about vehicle distribution, which is as expected based on the logic outlined in Section 3. Further evidence of this is provided by the coefficients for *DistributionPropensity_{jkst}* and *DistributionPropensity_{jkst}*PctBuyerElectronic_{kt}* (β_5 and β_6), which are both positive and significant. This indicates that sellers tend to sell vehicles where they have sold them in the past (which is likely due to geographic proximity, at least in part), and they have become even more likely to do so as electronic trading has increased.

In addition to the separate estimation of supply and demand just presented, we also investigated simultaneous estimation. See Appendix B for details.

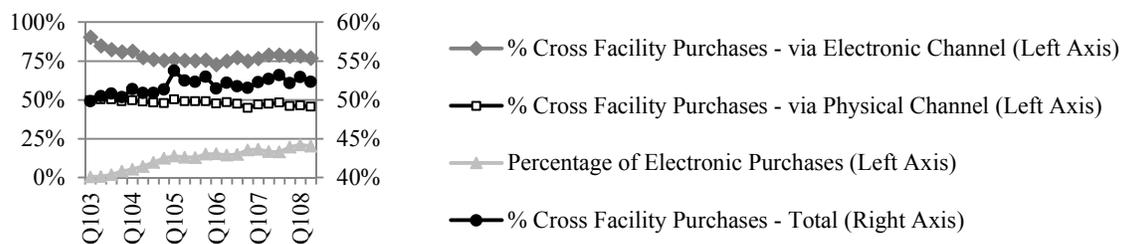
4.3 The Effect of Behavioral Changes on Geographic Price Variance: The results in Sections 4.1 indicate that buyers exploit geographic price variance by purchasing from facilities where prices are low

(and thereby eschewing purchasing from facilities where prices are high.) This demand-shifting is facilitated by the electronic channel and has become more pronounced as use of the electronic channel has grown. The results in Sections 4.2 indicate that sellers exploit geographic price variance by allocating supply to facilities where prices have been high. Contrary to the demand side, however, this behavior has become *less* pronounced as use of the electronic channel has grown. Essentially, sellers are becoming less strategic about vehicle distribution as electronic channel use has grown, because any pricing inefficiencies that a seller might otherwise exploit are increasingly likely to be eliminated once the auctions start by buyers using the electronic channel.

The change in buyer behavior should lower geographic price variance by better matching demand to the available supply across facilities. On the other hand, the change in seller behavior should *not* lower geographic price variance. If anything, it might *increase* geographic price variance, because sellers are less likely to shift supply to match expected demand. Thus, any relationship we find between changes in buyer behavior and geographic price variance should be in the correct direction but may be conservative.

To examine (and attempt to quantify) how the change in buyer behavior due to the electronic channel has lowered geographic price variance, we focused on the specific mechanism through which this should occur: what we refer to a “cross-facility purchase.” We define this as a purchase in which a buyer local to facility A purchased from facility B. Figure 5 shows that most purchases made via the electronic channel are “cross-facility” and that the total number of cross-facility purchases increased over time. This increase is attributable to increased use of the electronic channel over time.

Figure 5: Number of cross-facility purchases over time.



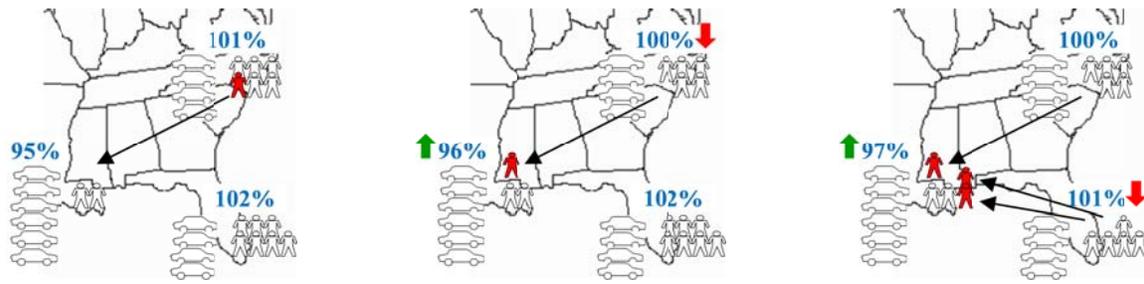
Notes: The large cross-facility statistics for the first few quarters for electronic transactions are an artifact of the low number of electronic transactions for these quarters.

Figure 6 illustrates how cross-facility purchases should lower geographic price variance by providing a hypothetical example of supply (represented by the vehicle icons) and demand (represented by the people icons) in Charlotte, New Orleans, and Miami. Note that prices (as a percentage of market value) are higher in Charlotte than in New Orleans. Because buyers using the electronic channel are more sensitive to price and less sensitive to distance (as shown in the buyer choice model), an electronic buyer in Charlotte (shown in red) is likely to shift his demand to New Orleans. Because prices are determined via an ascending auction, this “cross-facility purchase” will cause the expected prices in Charlotte to fall

and the expected prices at New Orleans to rise. This will reduce the geographic price variance between these facilities. As electronic channel use grows, more buyers will engage in these “cross-facility purchases” to exploit geographic price variance. The right panel of Figure 6 provides an example of buyers in Miami using the electronic channel to shift their demand to New Orleans. This reduces the price variance not only between Miami and New Orleans, but also between Charlotte and New Orleans (because prices in New Orleans creep up towards those in Charlotte). See Appendix A for a simple proof of this.

Although prices should become less variable, the cost of vehicle transport will prevent price variance from dropping to 0. Note that the “law of one price” allows prices to vary by the cost of transport between locations.

Figure 6: Illustration of how cross-facility purchase should lower geographic price variance.



Notes: The shaded buyers use the electronic channel. Their cross-facility purchases shift demand between locations, thereby lowering geographic price variance.

We used the following fixed effects panel regression model to examine the relationship between cross-facility purchases and geographic price variance.

$$CVPrice_{jkl,t} = \alpha + \beta_1 CrossFacilityPurchases_{jkl,t} + \beta_2 CrossFacilityPurchases_{jkl,t} * Distance_{kl} + \beta_3 SameFacilityPurchases_{jkl,t} + \beta_4 SameFacilityPurchases_{jkl,t} * Distance_{kl} + \sum_{t=2-2008} \beta_t time_t + c_{jkl} + \epsilon_{jkl,t} \quad (3)$$

The dependent variable $CVPrice_{jkl,t}$ is the coefficient of variation of mean prices between facilities k and l for vehicles of model j in quarter-year t . We calculated $CVPrice_{jkl,t}$ as follows. First, we calculated the mean price of vehicles of model j at each facility in each quarter-year. Second, we calculated the coefficient of variation of the mean prices for each facility pair at which vehicles of model j were traded in each quarter-year. For example, if the mean price of a Honda Accord was \$10,000 at facility k and \$11,000 at facility l at time t , then $CVPrice_{jkl,t} = 707/10,500 = 0.067$. The advantage of using the coefficient of variation as the variance measure is that it accounts for differences in value across vehicle models and across time (Baye et al., 2006). To measure the amount of cross-facility purchasing ($CrossFacilityPurchases_{jkl,t}$), we: a) counted the number of vehicles of model j purchased at facility l by

buyers local to facility k , b) counted the number of vehicles of model j purchased at facility k by buyers local to facility l , and c) summed the two. We constructed $SameFacilityPurchases_{jkl}$ analogously, except that $SameFacilityPurchases_{jkl}$ measures purchases made by buyers local to facility k (facility l) at facility k (facility l). The sum of $CrossFacilityPurchases_{jkl}$ and $SameFacilityPurchases_{jkl}$ is the total number of purchases of vehicles of model j made by the buyers local to either of the facilities in the pair at either of the facilities in the pair. Thus, including $SameFacilityPurchases_{jkl}$ allows us to control for this overall volume. $Distance_{kl}$ is the distance in miles between facilities k and l , which we scaled by dividing by 1,000. We interacted $Distance_{kl}$ with $CrossFacilityPurchases_{jkl}$ and $SameFacilityPurchases_{jkl}$. Because $Distance_{kl}$ does not vary over time, we could not include it as a main effect. $\sum_{t=2}^{2-2008} time_t$ represent dummy variables for each quarter, and c_{jkl} represents a fixed effect. Descriptive statistics for the variables in specification 2 appear in Table 7.

Table 7: Descriptive statistics for variables in the panel regression model.

Variable	Description	Mean (St. Dev.)	
		A	B
CVPrice _{jkl} _t	Coefficient of variation of mean prices between facilities k and l for vehicles of model j in time period t .	0.16 (0.25)	0.11 (0.18)
CrossFacilityPurchases _{jkl} _t	Number of purchases at facility l by buyers local to facility k (and vice versa) for vehicles of model j in time period t .	0.23 (1.62) ^a	0.84 (3.00) ^a
SameFacilityPurchases _{jkl} _t	Number of purchases at facility k by buyers local to facility k (and same for facility l) for vehicles of model j in time period t .	15.05 (30.22)	25.48 (42.18)
Distance _{kl}	Average distance in miles between facilities k and l , divided by 1000.	1.058 (0.635)	0.671 (0.503)
Column A: All facility pairs; Column B: Only facility pairs for which there was cross-facility purchasing. ^a $CrossFacilityTrans_{jkl}$ measures the purchases by buyers local to facility k at one other facility l (and vice versa.) When purchases by buyers local to facility k at all other facilities are calculated, the volume of cross-facility purchasing is roughly equal to that of same-facility purchasing, as shown in Table 9.			

Results of specification 2 appear in column A of Table 8. There were 621,141 facility pairs in the panel. There was no cross-facility purchasing for 532,360 of these facility pairs. Column B shows the results after excluding these facility pairs; this allowed us to develop estimates for facility pairs with active cross-facility purchasing.

Table 8 shows that cross-facility purchases have a negative relationship with price variance, as expected. Of interest is that both same-facility and cross-facility purchases negatively influence price variance, but the effect of cross-facility purchases is stronger; β_1 and β_3 are statistically different ($p < 0.01$.) We attribute this to the following. Any transaction increases trading volume, which reduces price variance because it increases the amount of price information in the market. However, a cross-facility purchase has a particularly strong effect because buyers purchase across facilities partly to exploit price discrepancies, which accentuates the reduction in geographic price variance. The coefficient for the interaction between $CrossFacilityTrans_{jkl}$ and $Distance_{kl}$ (β_2) is positive and significant, showing that the

effect of cross-facility transactions weakens with distance. We believe that this is because although the electronic channel reduces some of the frictions associated with distant trade, buyers must still ship vehicles back to their locations. Thus, especially long-distance trades are less likely to be motivated by price considerations and should therefore have a smaller effect on geographic price variance.

Table 8: Results of facility-pair price variance regressions.

	A: All Facility Pairs	B: Facility Pairs With Cross-Facility Trades
	Coefficient	Coefficient
β_1 : CrossFacilityTrans _{jkl,t}	-0.0016 (0.0001) ***	-0.0017 (0.0001) ***
β_2 : CrossFacilityTrans _{jkl,t} * Distance _{kl}	0.0013 (0.0002) ***	0.0013 (0.0002) ***
β_3 : SameFacilityTrans _{jkl,t}	-0.0003 (0.0000) ***	-0.0003 (0.0000) ***
β_4 : SameFacilityTrans _{jkl,t} * Distance _{kl}	-0.0000 (0.0000) ***	-0.0000 (0.0000)
Intercept	0.1659 (0.0005) ***	0.1237 (0.0007) ***
Time dummies	included	included
Vehicle model / Facility pair fixed effects	included	included
R ² , including fixed effects	0.53	0.43
n	3,986,978	1,100,421
Robust standard errors in parentheses.	*** p < 0.01.	

The average distance between buyer and facility for electronic transactions is 251 miles. Setting this as the distance and using the estimates from column B of Table 8, an additional cross-facility purchase by the buyers local to a facility pair decreases the coefficient of variation of an average vehicle model by 0.0014. The mean of $CVPrice_{jkl,t}$ for facility pairs with cross-facility purchasing is 0.11. This indicates that each additional cross-facility purchase for a facility pair separated by 251 miles is associated with a 1.3% decrease in the coefficient of variation of prices between the facility pair.

There is a risk of reverse causality in the panel regression because low price variance between facilities might lead to few cross-facility transactions. However, if this were the direction of the effect, then there would be a positive relationship between $CVPrice_{jkl,t}$ and $CrossFacilityPurchasesTrans_{jkl,t}$, which would make it more difficult to recover the negative relationship shown in Table 8. In addition, the results of the buyer choice model provide a clear mechanism through which causality flows from cross-facility purchasing to price variance. Nevertheless, for robustness against this potential endogeneity concern, we instrumented $CrossFacilityPurchases_{jkl,t}$ with $CrossFacilityPurchases_{jkl,t-1}$, i.e., the first lag. The first lag is a useful instrument when the endogeneity concern stems from reverse causality because it is not contemporaneous to the dependent variable. In the first stage regression, $CrossFacilityPurchases_{jkl,t-1}$ is positively correlated with $CrossFacilityPurchases_{jkl,t}$ ($\beta = 1.07$, std. error = 0.04). After instrumentation, the β_1 and β_2 coefficients in column B of Table 8 become -0.0060 (s.e. = 0.0005) and 0.0061 (s. e. = 0.0007), respectively, which suggests that the results shown in Table 8 are in the correct direction but potentially conservative in magnitude.

As an additional test, we created a panel containing the total number of cross-facility purchases ($CrossFacilityPurchases_{jt}$), the total number of same-facility purchases ($SameFacilityPurchases_{jt}$), and the coefficient of variation of price ($CVPrice_{jt}$) for vehicles of model j in each time period t across all facilities in the U.S. (instead of between discrete facility pairs as above.) Table 9 shows descriptive statistics for these variables.

$$CVPrice_{jt} = \alpha + \beta_1 CrossFacilityPurchases_{jt} + \beta_2 SameFacilityPurchases_{jt} + \sum_{t=2}^{2008} \beta_t time_t + c_j + \varepsilon_{jt}. \quad (4)$$

Table 9: Descriptive statistics for variables in specification 3.

Variable	Description	Mean (St. Dev.)
$CVPrice_{jt}$	Coefficient of variation of price for vehicles of model j in time period t .	0.23 (0.23)
$CrossFacilityPurchases_{jt}$	Number of cross-facility purchases for vehicles of model j in time period t .	130.34 (307.12)
$SameFacilityPurchases_{jt}$	Number of same-facility purchases for vehicles of model j in time period t .	125.42 (338.05)

We estimated specification 4 using time dummies and fixed effects for vehicle models. Table 10 shows the results. β_1 is negative and significant. A one standard deviation increase in $CrossFacilityTrans_{jt}$ is associated with a 13% decrease in $CVPrice_{jt}$. For reasons similar to those discussed above, we instrumented $CrossFacilityTrans_{jt}$ with the first lag. In the first stage regression, $CrossFacilityTrans_{j,t-1}$ is positively correlated with $CrossFacilityTrans_{jt}$ ($\beta = 0.08$, std. error = 0.03). After instrumentation, β_1 becomes -0.0004 (std. error = 0.0001). This suggests that β_1 is in the correct direction but conservative.

Table 10: Results of country-wide price variance and price mean regressions.

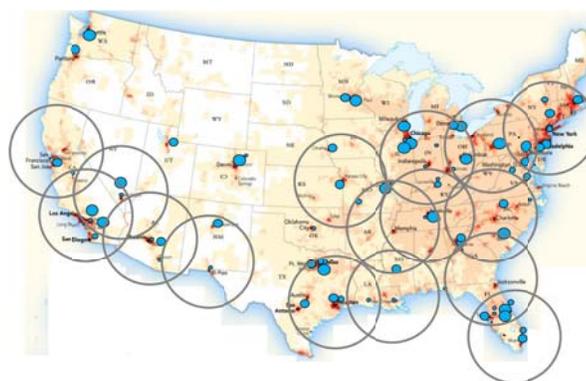
	Coefficient
β_1 : $CrossFacilityTrans_{jt}$	-0.0001 (0.0000) ***
β_2 : $SameFacilityTrans_{jt}$	0.0000 (0.0000)
Intercept	0.2333 (0.0056) ***
Time dummies	included
Vehicle model fixed effects	included
R^2 , including fixed effects	0.66
n	8,017
Robust standard errors in parentheses.	*** $p < 0.01$.

The results shown in Table 10 show that prices across the U.S. are becoming more uniform or what we call “flatter”, while the results shown in Table 8 show that the flattening effect of electronic trading weakens with distance. This presents a paradox: how can prices across the entire country be flattening when the flattening effect operates primarily on a regional basis?

We offer two possible explanations. First, regional price flattening eliminates extremely high and low prices from the market, which will lower price variance for the country as a whole. Second, it is possible

that the flat regions overlap, which causes prices across the entire country to flatten via transitivity. Figure 7 illustrates this. The map shows the population density of the United States. The small circles represent market facilities, with the size of the circle reflective of transaction volume. The large circles are centered on major population centers and have a radius of approximately 251 miles, which is the average distance between buyer and facility for electronic purchases in our data. Our results suggest that buyers in each population center are flattening prices within (and sometimes beyond) their circle. Thus, each large circle represents a flat region. Because Boston and Baltimore are in the same region, prices between them are flattening; the same holds for Baltimore and Detroit. By transitivity, this means that prices in Boston and Detroit are flattening, as are, by extension, prices in Boston and San Antonio.

Figure 7: Illustration of how individual “flat” regions overlap to yield a “flat” country. (Map downloaded from <http://www.maps.com/map.aspx?pid=17073>.)



5. Conclusion

We analyzed over 2 million wholesale transactions from 2003 to 2008 for vehicles with between 15,000 and 21,000 miles. Use of the electronic channel to purchase vehicles grew from approximately 0% to approximately 20% of transactions over this time. Our analysis shows that the electronic channel has made buyers more sensitive to price but less sensitive to distance. This means that many buyers who would have otherwise used the physical channel to purchase a vehicle of model j from a nearby facility A for price p are using the electronic channel to purchase a vehicle of model j at a remote facility B for price p^* , with $p^* < p$. The shifting of demand across facilities facilitated by the electronic channel should cause prices at facility A to fall and prices at facility B to rise, thereby making prices more uniform or “flatter.” Our results provide evidence that this is happening, thereby causing the market to more closely reflect the “law of one price.” Due to the cost of vehicle transportation, the price flattening effects attenuate with distance, although the effects are evident for the U.S. as a whole. We find that sellers are responding rationally to these market shifts by becoming less strategic about vehicle distribution, given that buyers are increasingly using the electronic channel to eliminate any pricing inefficiency that the seller might otherwise try to exploit.

6.1: Limitations: Although we limited heterogeneity among vehicles by restricting our sample to vehicles with between 15,000 and 21,000 miles, no two vehicles of the same model in our sample are identical. For example, a low mileage Chevrolet Malibu located at facility A differs from a low mileage Chevrolet Malibu at facility B on more dimensions than just location. Even if the two vehicles were identical when first manufactured (e.g., same color, same option packages), they are no longer identical after having been driven 15,000 miles. We address this by including multiple controls for vehicle quality, including normalizing price by valuation and including vehicle condition data. In our analysis of the change in geographic price variance, we cannot determine how much of the change we observe is due to increased cross-facility purchasing and how much is due to heterogeneity within vehicle models. However, we have no reason to believe that the heterogeneity within vehicle models changed significantly over the time span covered by our data, but we know that cross-facility purchasing increased during this time. Thus, we believe that the reduction in price variance we observe is attributable to demand shifts facilitated by the electronic channel, and that this holds despite noise due to vehicle heterogeneity.

6.2: Intended Contributions: Our study links two research streams: 1) how electronic commerce affects price dispersion, and b) how electronic commerce affects geographic trade. The price dispersion stream has not considered the geographic location of products, despite the fact that it plays an important role in the trade of many products such as automobiles, food products, and raw materials. The geographic trade stream has not considered whether or how electronic commerce leads to lower price dispersion across locations. In joining these two streams, we make the following contributions. First, we document the behavioral mechanism by which electronic trading affects price dispersion, which is critical for continued empirical research in this stream (Baye et al. 2006). Specifically, our results show that buyers are more sensitive to price and less sensitive to distance when using the electronic channel than when using the physical channel. This increases the likelihood that buyers will eschew purchasing at a nearby facility to purchase at a remote facility where prices are lower. These cross-facility purchases are the mechanism that has led to lower price variance between facilities. Second, we show that the location of products and buyers plays an important role in buyer behavior and price dispersion. The distance between products and buyers influences how buyers choose the facility and channel from which to purchase and moderates the relationship between cross-facility trading and price dispersion.

Appendix A: Simple Proof of Why Cross-Facility Purchasing Should Reduce Geographic Price Variance

What follows is a simple proof of why shifting a bidder from a high-priced location to a low-priced location lowers the price variance between locations.

Setup: Let there be two open-outcry ascending auctions, each for a vehicle of model j (e.g., Honda Accord), where one auction occurs at facility A and one auction occurs at facility B. Denote the number of bidders at facilities A and B as n_a and n_b , respectively. Each bidder draws a valuation v_i for vehicles of model j from $F(\cdot)$, where $v_i \in [\underline{v}, \bar{v}]$. Denote the prices of the auctions at facilities A and B as pr_a and pr_b , respectively. Let $|pr_a - pr_b| = \delta$. δ measures the variance between pr_a and pr_b because the absolute value of the difference of two numbers is equal to $\sqrt{2}$ times the standard deviation of the two numbers.

Proposition: If the variance of prices between the two facilities is nonzero (i.e., if $\delta > 0$), then shifting a bidder from the facility with the higher price to the facility with the lower price will reduce the variance of prices (or keep the variance the same.) More formally, if $pr_a - pr_b = \delta > 0$, then shifting a bidder from facility A (such that n_a becomes $n_a - 1$) to facility B (such that n_b becomes $n_b + 1$) will yield $|pr_a' - pr_b'| = \delta' \leq \delta$, where pr_a' and pr_b' are the auction prices after shifting the bidders. The converse is true if $pr_b - pr_a = \delta > 0$.

Proof: The proof is based on properties of order statistics. Using a well-known result from auction theory about open outcry auctions, the expected price is the valuation of the second-highest bidder. In an auction with n bidders, this means that the expected price is the $(n-1)^{\text{st}}$ order statistic from the distribution of those bidders' valuations. To simplify the math, we assume that bidder valuations are drawn from a uniform distribution bounded by $[\underline{v}, \bar{v}]$, although uniformity is not required. From the properties of a uniform distribution bounded by $[\underline{v}, \bar{v}]$, the expected value of the $(n-1)^{\text{st}}$ order statistic is $\underline{v} + \frac{n-1}{n+1}(\bar{v} - \underline{v})$

(Arnold et al., 1992). Thus, $pr_a = \underline{v} + \frac{n_a - 1}{n_a + 1}(\bar{v} - \underline{v})$ and $pr_b = \underline{v} + \frac{n_b - 1}{n_b + 1}(\bar{v} - \underline{v})$, respectively. The proof is

accomplished by showing that $\delta \geq \delta'$, i.e., that

$$\left| \underline{v} + \frac{n_a - 1}{n_a + 1}(\bar{v} - \underline{v}) - \left(\underline{v} + \frac{n_b - 1}{n_b + 1}(\bar{v} - \underline{v}) \right) \right| \geq \left| \underline{v} + \frac{(n_a - 1) - 1}{(n_a - 1) + 1}(\bar{v} - \underline{v}) - \left(\underline{v} + \frac{(n_b + 1) - 1}{(n_b + 1) + 1}(\bar{v} - \underline{v}) \right) \right| \quad (1)$$

To show the proof, we manipulate this expression assuming $pr_a > pr_b$. The proof is analogous for $pr_b > pr_a$. For $pr_a > pr_b$, we can remove the absolute value operator from the left-hand side of (1) because it is always positive. We then examine two cases: a) $pr_a' \geq pr_b'$, and b) $pr_a' < pr_b'$. For $pr_a' \geq pr_b'$, we can remove the absolute value operator from the right-hand side of expression 1. Collecting the n_a and n_b

terms on each side of the above expression and performing some algebraic manipulation yields:

$$\frac{2}{(n_b + 2)(n_b + 1)} \geq \frac{-2}{n_a(n_a + 1)},$$

which always holds. For $pr_a' < pr_b'$, we can remove the absolute value operator from the right-hand side of expression 1 after multiplying by -1. Note that when $pr_a > pr_b$, the only way for $pr_a' < pr_b'$ to hold is if $n_a = n_b + 1$. Substituting $n_a = n_b + 1$, collecting the n_a and n_b terms on each side of the above expression, and conducting some algebra yields:

$$\frac{n_b}{n_b + 2} + \frac{n_b - 1}{n_b + 1} \geq \frac{n_b}{n_b + 2} + \frac{n_b - 1}{n_b + 1}.$$

Both sides of the expression are equal, so the expression holds.
 QED.

Appendix B: Simultaneous Estimation of Supply and Demand in a Choice Setting in Which Prices are Determined by Auction

This appendix discusses the possibility of simultaneous estimation of supply and demand in our context. As background, consider that the most common simultaneous models of supply and demand in a non-linear discrete choice framework are based on the work of industrial economists such as Berry, Levinsohn, and Pakes (Berry et al., 1995; Berry, 1994). These models assume that sellers set prices based on an equilibrium pricing rule specified by the econometrician (see Train, 2009: chapter 13 for a thorough discussion.) For example, the econometrician might assume that the seller sets prices based on marginal cost or that the seller chooses a profit-maximizing price in an oligopoly setting, etc. These models do not apply in our context because sellers do not set prices in our context; prices are determined by auction. Also, it is not clear how a new method that is appropriate for an auction context might be constructed. Thus, standard methods for incorporating the supply side in a discrete choice framework are not applicable to our context. However, if we develop linear expressions for supply and demand, then we can estimate them simultaneously using established simultaneous equations methods. However, as we discuss below, we concluded that this approach is inferior to the separate estimation reported in the paper. Accordingly, we did not implement this method, although we describe it below.

To develop linear expressions for supply and demand, we began with a structural model of individual choice based on utility maximization and then aggregated over individuals to yield linear models akin to “gravity models” (e.g., Blum & Goldfarb, 2006; Hortacsu et al., 2009). We begin with the construction of the linear demand model.

Linear Demand Model

Let s = seller, b = buyer, m = vehicle model, i, j = facility, c = channel, and t = week. Each buyer b has a “local” facility i determined based on his zip code.

Let N_{jm}^B = the total number of buyers of vehicles of model m who are local to facility j . The probability that a buyer b purchases a vehicle of model m from facility j at time t is:

$$P_{jmt}^B = \frac{N_{jm}^B e^{V_{jmt}^B}}{\sum_{j \in J} N_{jm}^B e^{V_{jmt}^B}} \quad (\text{B1})$$

(B1) is a standard logit probability, weighted by N_{jm}^B . Weighting is necessary because the probability of purchasing at facility j increases with the number of buyers who are local to facility j .

Let N_{mt}^v be the number of vehicles of model m that buyers purchased at time t at any facility. Let Q_{jmt}^B be the number of vehicles of model m that buyers purchased at facility j at time t .

$$Q_{jmt}^B = N_{mt}^v P_{jmt}^B = N_{mt}^v \frac{N_{jm}^B e^{V_{jmt}^B}}{\sum_{j \in J} N_{jm}^B e^{V_{jmt}^B}} \quad (\text{B2})$$

Taking the log of (B2) yields:

$$\ln Q_{jmt}^B = \ln N_{mt}^v + \ln N_{jm}^B + V_{jmt}^B - I_{jt}^B \quad (\text{B3})$$

where $I_{jt}^B = \ln \sum_{j \in J} N_{jm}^B e^{V_{jmt}^B}$ and represents the population-weighted expected utility that a buyer receives from the choice situation. Absorb I_{jt}^B into a constant term, replace $V_{jmt}^B = \sum_{h=1}^H \alpha_h X_{h,jmt}^B$, and add coefficients, an intercept, and a disturbance to yield the demand equation:

$$\ln Q_{jmt}^B = \alpha_{0,jm} + \sum_{h=1}^H \alpha_h X_{h,jmt}^B + \delta_1 (\ln N_{mt}^v) + \delta_2 (\ln N_{jm}^B) + \varepsilon_{jmt}^B \quad (\text{B4})$$

Linear Supply Model

Define the ‘‘grounding city’’ as the city in which a vehicle is located prior to its entering the market. We do not observe the grounding city, but we assume that the probability that a vehicle is in a given grounding city is proportional to the size of the grounding city. E.g., a vehicle is more likely to be grounded in Dallas, TX than in Darlington, SC. Let Pop_j = the population of the city in which facility j is located. If we assume that sellers prefer to sell vehicles at facilities near the vehicles’ grounding city, all else equal, then we can write the probability that a seller s chooses facility j to sell a vehicle of model m at time t as:

$$P_{jmt}^S = \frac{Pop_j e^{V_{jmt}^S}}{\sum_{j \in J} Pop_j e^{V_{jmt}^S}} \quad (\text{B5})$$

Let N_{mt}^v = the total number of vehicles of model m that sellers sell at time t at any facility. Let Q_{jmt}^S = the number of vehicles of model m that sellers sell *at facility j* at time t .

$$Q_{jmt}^S = N_{mt}^v P_{jmt}^S = N_{mt}^v \frac{Pop_j e^{V_{jmt}^S}}{\sum_{j \in J} Pop_j e^{V_{jmt}^S}} \quad (\text{B6})$$

Taking the log of (B6) yields:

$$\ln(Q_{jmt}^S) = \ln(N_{mt}^v) + \ln(Pop_j) + V_{jmt}^S - I_{jt}^S \quad (\text{B7})$$

where $I_{jt}^S = \ln \sum_{j \in J} Pop_j e^{V_{jmt}^S}$. Absorb I_{jt}^S into a constant term, replace $V_{jmt}^S = \sum_{k=1}^K \beta_k X_{k,jmt}^S$ and add coefficients, an intercept, and a disturbance to yield the supply equation:

$$\ln Q_{jmt}^S = \beta_{0,jm} + \sum_{k=1}^K \beta_k X_{k,jmt}^S + \gamma_1 (\ln N_{mt}^v) + \gamma_2 (\ln Pop_j) + \varepsilon_{jmt}^S \quad (\text{B8})$$

Q_{jmt}^B (Q_{jmt}^S) is the number of vehicles of model m purchased (sold) at facility j at time t . Variables in $X_{h,jmt}^B$ include: a) average price of vehicles of model m at facility j , b) number of vehicles auctioned (not necessarily sold) at facility j , c) the average propensity for buyers local to facility j to use the electronic channel, and d) interactions. $X_{k,jmt}^S$ is similar. The interaction terms in these models allow us to test how buyers' and sellers' sensitivities to price changes as use of the electronic channel grows. The equilibrium condition is $\ln Q_{jmt}^B = \ln Q_{jmt}^S$ (i.e., demand = supply.)

Some Considerations with Simultaneous Estimation

The distance to each facility for buyers and sellers is not considered in either of the linear supply or demand equations. This is unavoidable for sellers, because we don't know where sellers store vehicles prior to their entering the market. We observe distance for buyers, but cannot include it in the linear demand equation because the linear demand equation measures Q_{jmt}^B . Measuring Q_{jmt}^B allows us to use the "supply=demand" equilibrium condition $\ln Q_{jmt}^B = \ln Q_{jmt}^S$. Q_{jmt}^B is the total number of vehicles of model m purchased at facility j at time t by buyers located throughout the country. This means that a distance measure would have to be a weighted average of the distance to facility j from all buyers' locations or some other measure of "remoteness." This would provide a limited and indirect means of examining how the electronic channel affects buyers' sensitivity to distance. Abstracting away from distance is a significant limitation because distance is what generates the friction that causes geographic price variance in the first place.

Another consideration is that sellers have the option to sell or retain each auctioned vehicle (i.e., they can accept or reject the high bid). Thus, there are two seller behaviors at work in our context: a) where they choose to distribute vehicles, and b) whether they choose to sell the vehicles once they have distributed them. Section 4.2 directly models the first behavior. The linear supply equation above models some combination of both behaviors in that it accounts for how many vehicles sellers choose to sell at a facility k , with that number not to exceed the number of vehicles they had distributed there. We do not believe it possible to separate the two behaviors in this approach. Although both behaviors are relevant, the first behavior is more pertinent to the research question. This is because it is reasonable that sellers might systematically distribute vehicles in a way that affects geographic price variance (see the logic in Section 3.) It is less reasonable to expect that sellers systematically accept and reject high bids in a way that affects geographic price variance.

Given these issues, we determined that separate estimation was the most appropriate approach for our context. Also, any concern about a simultaneity bias in our estimation is mitigated by our use of daily data and the sequential nature of buyer/seller behavior: sellers choose where to distribute vehicles before

buyers choose where to purchase vehicles. The sequential nature of the decisions allows us to include the relevant supply variables directly in the demand model and vice versa, and the daily data ensure that the sequence of behavior is not confounded by aggregating across multiple days.

References

- Aker, J. C. 2010. Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger. *American Economic Journal: Applied Economics* **2**(3) 46-59.
- Arnold, B. C., N. Balakrishnan, H. N. Nagaraja. 1992. *A First Course on Order Statistics*. New York: John Wiley & Sons.
- Bakos, J. Y. 1997. Reducing Buyer Search Costs: Implications For Electronic Marketplaces. *Management Science* **43**(12) 1676-1692.
- Baye, M. R., J. Morgan, P. Scholten. 2006. Information, Search, and Price Dispersion. In T. Hendershott (Ed.), *Handbook on Economics and Information Systems*. Amsterdam: Elsevier.
- Berry, S., J. Levinsohn, A. Pakes. 1995. Automobile Prices in Market Equilibrium. *Econometrica* **63**(4) 841-890.
- Berry, S. T. 1994. Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics* **25**(2) 242-262.
- Blum, B., A. Goldfarb. 2006. Does the internet defy the law of gravity? . *Journal of International Economics* **70**(2) 384-405.
- Brown, J. R., A. Goolsbee. 2002. Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry. *Journal of Political Economy* **110**(3) 481-507.
- Brynjolfsson, E., M. D. Smith. 2000. Frictionless Commerce? A Comparison of Internet and Conventional Retailers. *Management Science* **46**(4) 563-585.
- Bucklin, R. E., S. Siddarth, J. Silva-Risso. 2008. Distribution Intensity and New Car Choice. *Journal of Marketing Research* **45**(3) 473-486.
- Chellappa, R. K., R. G. Sin, S. Siddarth. 2011. Price-formats as a Source of Price Dispersion: A Study of Online and Offline Prices in the Domestic U.S. Airline Markets. *Information Systems Research* **22**(1) 83-98.
- Chiou, L. 2009. Empirical Analysis of Competition between Wal-Mart and Other Retail Channels. *Journal of Economics & Management Strategy* **18**(2) 285-322.
- Clay, K., R. Krishnan, E. Wolff, D. Fernandes. 2002. Retail Strategies on the Web: Price and Non-Price Competition in the Online Book Industry. *Journal of Industrial Economics* **50**(3) 351-367.
- Clemons, E. K., I. H. Hann, L. M. Hitt. 2002. Price Dispersion and Differentiation in Online Travel: An Empirical Investigation. *Management Science* **48**(4) 534-549.
- Diekmann, F., B. Roe, M. Batte. 2008. Tractors on eBay: Differences between Internet and In-Person Auctions. *American Journal of Agricultural Economics* **90**(2) 306-320.
- Ghose, A., Y. Yao. 2011. Using Transaction Prices to Re-Examine Price Dispersion in Electronic Markets. *Information Systems Research* **22**(2) 269-288.

- Hortacsu, A., F. A. Martinez-Jerez, J. Douglas. 2009. The Geography of Trade in Online Transactions: Evidence from eBay and MercadoLibre. *American Economic Journal: Microeconomics* **1**(1) 53-74.
- Jensen, R. 2007. The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *Quarterly Journal of Economics* **122**(3) 879-924.
- Overby, E., S. Jap. 2009. Electronic and Physical Market Channels: A Multiyear Investigation in a Market for Products of Uncertain Quality. *Management Science* **55**(6) 940-957.
- Stigler, G. 1961. The Economics of Information. *Journal of Political Economy* **69**(3) 213-225.
- Train, K. E. 2009. *Discrete Choice Methods with Simulation* (Second Edition ed.). Cambridge: Cambridge University Press.