American Economic Review 2012, 102(5): 2147–2182
http://dx.doi.org/10.1257/aer.102.5.2147

Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline†

By Jean-François Houde*

This paper studies an empirical model of spatial competition applied to gasoline markets. The main feature is to specify commuting paths as the "locations" of consumers in a Hotelling-style model. As a result, spatial differentiation depends in an intuitive way on the structure of the road network and the direction of traffic flows. The model is estimated using panel data on the Quebec City gasoline market and used to evaluate the consequences of a recent vertical merger. Difference-in-difference and counterfactual simulation methods are compared, and the results, to a large extent, validate the assumptions of the demand model. (*JEL* G34, L13, L42, L81, Q41, R41)

The dominance of major retail chains in gasoline markets is a concern for many state governments and antitrust authorities. The presumption that vertically-integrated retailers have significant market power has led to various regulations that protect independent retailers and limit the conduct of major oil companies; including below-cost price regulations (i.e., price floors) and divorcement acts banning vertical integration. The recent wave of mergers among US oil companies has drawn more scrutiny from consumer protection groups and public administrations. For example, the US Senate conducted an investigational hearing to examine the conduct and structure of the industry, and the Federal Trade Commission sponsored a conference on the retrospective evaluation of mergers in gasoline markets, following the publication of a governmental report suggesting that several prominent mergers led to sizable price increases.[1]

Despite this interest from policy makers, little research has been conducted to model demand and supply in gasoline retail markets. Instead, these markets have been at the forefront of the development of a large empirical literature focused on

* Department of Business Economics and Public Policy, The Wharton School, University of Pennsylvania, 1400 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104-6372 (e-mail: houde@wharton.upenn.edu). I am grateful for the advice and support of Chris Ferrall and Susumu Imai. I thank two anonymous referees for their comments and suggestions. I have benefited from the comments and discussions of Jason Allen, Robert Clark, Juan Esteban Carranza, Amit Gandhi, Justine Hastings, Ken Hendricks, John Kennan, Beverly Lapham, Jeremy Lise, Sumon Majumdar, Jack Porter, Shannon Seitz, and seminar participants at Queen's, Concordia, UQAM, HEC-Montréal, Université de Montréal, University of Pennsylvania, Bank of Canada, Simon Fraser University, New York University, New York Fed, University of Wisconsin, the University of Northern-Illinois, ECSM 2006 in Minneapolis, SED 2008 in Prague, the University of Chicago GSB, and Yale University.

† To view additional materials, visit the article page at http://dx.doi.org/10.1257/aer.102.5.2147.

[1] The Senate report and experts' opinions are available here: http://www.gpo.gov/congress/senate/senate12sh107.html. The report of the FTC conference is available here: http://www.ftc.gov/ftc/workshops/oilmergers/index.shtm.

the retrospective analysis of consummated mergers.[2] While this approach offers a transparent evaluation of approved mergers, the results are often difficult to generalize and are sensitive to the sample choice. In contrast, the merger-simulation methodology allows researchers to study counterfactual changes in market structure, but relies on the estimation of a well-specified model of demand and supply; a task that can be difficult to implement in markets with a complex structure of spatial differentiation like gasoline.[3] As a result, the structural approach to merger evaluation has recently been criticized for its lack of robustness and sometimes biased predictions (see, for instance, Peters 2006 and Weinberg and Hosken 2009).[4]

The approach taken in this paper addresses these concerns by combining structural and reduced-form approaches to the evaluation of an existing merger. I first estimate a model of demand for spatially differentiated goods that explicitly takes into account the geography of the market; namely the road network and the direction of traffic flows. Next, I use the structural model to simulate the impact of an existing merger, two years before it occurred. Finally, I compare these merger simulation predictions with difference-in-difference estimates of the merger impact on prices.

The advantages of this comparative approach are twofold. First, the comparison between the natural-experiment and structural merger results can be used to evaluate the assumptions of the structural model, which is estimated without the supply-side restrictions imposed by the observed ownership change. This motivation is related to the idea of validating econometric models using experimental data (see LaLonde 1986). Second, it illustrates the strengths and weaknesses of each approach and enhances our understanding of the potential biases of using one method or the other.[5]

In markets with spatially differentiated products, the market power of retail chains originates from their ability to raise prices by internalizing the cannibalization of sales between their stores. Any attempt to quantify the consequences of counterfactual mergers therefore requires prior knowledge on demand and cost, and in particular on the substitution patterns among stations. Quantifying the degree of spatial differentiation is not a trivial exercise, however, since the geography of actual retail markets is more complex than the line or circle paradigms traditionally used. While the standard Hotelling-style model of spatial competition situates each potential customer at a single point (i.e., single-address), gasoline stations located anywhere along the driving route of a consumer intuitively require similar shopping costs.

The novelty of my approach is to incorporate the mobility of consumers in the product space into a discrete-choice model of demand. This is done by defining

[2] The following papers evaluate the impact of existing gasoline mergers on prices: Hastings (2004); Hastings and Gilbert (2005); Taylor and Hosken (2007); Simpson and Taylor (2008); and Taylor, Kreisle, and Zimmerman (2010). Prominent examples outside of the gasoline literature include: Hortaçsu and Syverson (2007); Dafny (2009); and Dafny, Duggan, and Ramanarayanan (2012).

[3] The merger-simulation approach has been pioneered by Baker and Bresnahan (1985); Berry and Pakes (1993); Hausman, Leonard, and Zona (1994); Werden and Froeb (1994); and Nevo (2000).

[4] Ashenfelter, Hosken, and Weinberg (2009) summarize the debate and survey the retrospective merger literature and Angrist and Pischke (2010) and Nevo and Whinston (2010) provide conflicting points of view on the usefulness of retrospective merger analysis versus merger simulation.

[5] This paper is not the first to combine structural and reduced-form methods to study mergers. Peters (2006) and Weinberg and Hosken (2009) also use observed mergers to test the assumptions of structural models. See also Hausman and Leonard (2002) for a similar comparative analysis in the context of the introduction of new goods.

consumers' locations as their entire commuting paths. This modeling assumption has two important implications for the analysis of mergers.

The first refers to the estimation of substitution patterns: the elasticity of substitution between two products is a function not only of the distance between their locations, but also their connectivity along the road network and the direction of traffic flows. Competition is therefore not solely localized since consumers can substitute stations far from each other but close to a common commuting path. As a result, even if consumers are not willing to deviate far from their path to shop for gasoline, price differences are unlikely to persist if there is substantial commuting between two regions.[6] The consequences for competition are important: all else being equal retail markets with commuting consumers tend to generate less differentiation, and more intense price competition.[7]

The second implication is related to the definition of the relevant market, a crucial step in the evaluation of proposed mergers. Since the elasticity of substitution between stations mimics the distribution of traffic in cities, price competition spills over locations that potentially include all the neighborhoods of a metropolitan area. A recent merger proposal between Pilot and Flying-J, two of the largest chains of gasoline stations along interstate highways, provides a good example. In this market, trucking companies use sophisticated information systems to track diesel prices and optimize fuel stops accordingly. The distance metric is therefore a function of the delivery route, and the FTC ruled that the relevant market for these two national companies corresponds to the entire country, rather than metropolitan areas or major highway segments.[8]

The econometric analysis is conducted using the Quebec City gasoline market as a case study. This region is particularly well suited since crucial information on the geography of the market is available, including data on the commuting habits of consumers, the road network structure, and accurate measures of stations' locations and sales' volume. The dataset covers an 11 year period, between 1991 and 2001. During this time period, important changes were underway in the structure of the North American gasoline industry associated with the substantial exit of stations and entry of new categories of retailers. This reorganization of the industry provides an important source of variation in the choice set of consumers, which in turn identifies the key parameters of the model. Moreover, midway into the sample the market experienced a merger between two of the largest retail gasoline companies in Canada, Ultramar and Sunoco.[9] Since Ultramar operates as a vertically integrated company in the Quebec market, the merger changed the number of vertically integrated competitors in a neighborhood of each Sunoco station.

Beyond its contribution to the aforementioned literature on merger evaluation, this paper contributes to the literature on the estimation of demand for differentiated

[6] This is consistent with the observation that there is significantly less price dispersion within gasoline markets than in most other retail markets (Lach 2002).

[7] The theoretical results in Claycombe (1991) and Raith (1996) show that, in general, the fact that consumers are able commute between locations reduces prices and limits the capacity of firms to price discriminate relative to the Hotelling model.

[8] More details on the case are available here: http://www.ftc.gov/os/caselist/0910125/index.shtm.

[9] Technically, the two companies did not merge but rather "swapped" part of their network of gasoline stations. The two companies announced in 1996 that Ultramar would operate Sunoco stations in the province of Quebec while Sunoco would operate Ultramar stations in part of Ontario.

products (see Ackerberg et al. 2007 for an extensive review of this literature). The methodology proposed by Berry, Levinsohn, and Pakes (1995) has been extended in several directions to evaluate market responses to policy changes. Closely related to my application, Davis (2006), Manuszak (2001), and Thomadsen (2005) extend this methodology to estimate different variants of the single-address model.[10]

The empirical results are summarized as follows. The first set of results compares the predictions of the multi-address demand specification with a single-address model in which consumers are "located" at home. The model based on commuting behavior is shown to fit the observed distribution of gasoline sales more closely than the standard home-address model. This leads to very different estimates of consumers' transportation cost and markups. Transportation cost is small and statistically insignificant in the single-address model and negative and economically important in the commuting model. Since large-volume gasoline stations tend to be located at the intersection of major commuting paths, the single-address model systematically underestimates the transportation cost of consumers.

This bias has direct implications for market power. The estimated markups under the commuting model are small and similar in magnitude to the observed margins between retail and rack prices, a common proxy for marginal cost. The markups estimated using the alternative demand model are 50 percent higher. These differences originate from the fact that differentiation in the estimated single-address model is mainly related to the presence of idiosyncratic tastes for products, rather than geographic distances (i.e., since the transportation cost is close to zero).

The second set of results is related to analysis of the Ultramar/Sunoco merger. Although the merger led to small changes in the structure of the Quebec City market, the simulation results show that it generated sizable price increases locally, especially among merging parties competing for the same consumers.

To evaluate the validity of these results, the merger simulation predictions are compared with observed price changes estimated using a difference-in-difference approach similar to Hastings (2004). The results suggest that the merger led to statistically significant price increases in line with other estimates obtained in US markets. Importantly, the estimates are also shown to align well with the merger simulation predictions. On average, prices were predicted to increase by 0.38 cents per liter (cpl) in the neighborhood of Sunoco stations acquired by Ultramar. In comparison, the observed impact of the merger is estimated to be between 0.15 and 0.45 cpl depending on the specifications. These changes translate into 4–11 percent increases in average retail margins.

The rest of the paper is organized as follows. Section I introduces the data and describes some of the key institutional features of the market. I then present an analysis of gasoline demand starting with a description of the model specifications (Section II), the estimation methodology (Section III), and finally a discussion of the empirical results (Section IV). Section V is devoted to the analysis of the vertical

[10] McManus (2007) estimates a similar model in the market for coffee shops and focuses on the extent of non-linear pricing, and Smith (2004) studies demand for grocery products using micro-data from a UK household survey. Pinkse, Slade, and Brett (2002) developed a different approach based on a direct approximation of the equilibrium pricing rules and apply their techniques to study the distribution of wholesale gasoline prices. Manuszak and Moul (2009) estimate consumers' transportation cost for gasoline using observed tax differences at the boundaries of Illinois.

merger. I conclude the paper and discuss extensions in Section VI. Further computational and data construction details are placed in the online Appendix.

I. Overview of the Data and Market

In the following subsections, I describe the two main sources of data used in the empirical analysis. The first is a detailed survey of gasoline stations, and the second is a transportation survey used to compute the distribution of commuting paths within the market. I also present a set of descriptive statistics documenting the correlation between the distribution of gasoline sales and the distribution of consumers' locations.

A. *Gasoline Data and Market Trends*

The gasoline station data were collected by Kent Marketing, the leading survey company for the Canadian gasoline market. The panel spans 46 bimonthly periods (every two months) between 1991 and 2001 for every station in the Quebec City market.[11] The survey offers accurate measures of market shares and station characteristics since each site is physically visited at the end of the survey period and volume sold in liters per day is measured by reading the pump's meters.[12]

The observed characteristics of stations include the type of convenience store (small, medium, or large), a car-repair shop indicator, the number of service islands and pumps, opening hours, brand name, type of service, and an indicator for the availability of a car wash. The sample includes 14,263 observations for 429 different gasoline station sites. On average, each product is observed for 42 periods, and 75 percent of stations are observed for more than 15 periods.

Table 1 summarizes the key variables. The summary statistics highlight the large amount of heterogeneity across stations. On average, stations sell 5,000 liters of gasoline per day, with a standard-deviation above 3,000. This heterogeneity is also reflected in the capacity of stations; measured by the number of pumps and service islands. All three measures of size increased significantly over the sample period.

The price variable represents the end-of-period unleaded price, posted on the day of the survey, normalized by the relevant period consumer price index for the Quebec City region. Since most stations in the city are surveyed on the same day, this represents a reliable picture of the cross-sectional distribution of prices.[13] This variable, however, ignores the important time-series volatility of prices. Since the objective of the model is to predict the distribution of market shares, the key assumption maintained in the empirical analysis is that the distribution of prices is stable throughout the period.

[11] For the first four years I observe only the fourth period.

[12] Note that about 8 percent of the sample stations refused to participate in the survey for some or all periods. For those observations, the station characteristics (including prices) are accurately measured, but the volume sold is not available. Since the estimation procedure cannot accommodate missing values in the quantity variable, I imputed the missing values using linear regression methods. The explanatory variables include the average neighborhood market shares, predicted traffic, population density, a polynomial function of the geographic coordinates of the locations, prices, characteristics, and lagged sales (for stations that were previously participating in the survey).

[13] Over the whole period, 64.5 percent of stations were surveyed on the same day and 97 percent were surveyed within a two day window around the median date. The survey company improved its performance over time: 100 percent of stations were surveyed on the same day in 2001.

Table 1—Summary Statistics on Station Characteristics, Volumes and Prices

	Fall 1991		Fall 2001		Full sample	
	Avg.	SD	Avg.	SD	Avg.	SD
Volume (liter/day)	3,949	2,415	6,271	3,812	4,934	3,289
Price (cpl)	65.13	1.24	62.07	0.22	63.83	1.78
Absolute price deviation (cpl)	0.90	0.85	0.13	0.17	0.57	0.73
Number of pumps	7.77	5.40	11.37	8.12	9.29	6.92
Number of islands	2.07	1.26	2.46	1.43	2.24	1.35
Large convenience store	0.23	0.42	0.34	0.47	0.27	0.45
Full service	0.54	0.50	0.28	0.45	0.43	0.50
Open 24 hours	0.34	0.47	0.42	0.49	0.37	0.48
Car wash	0.18	0.39	0.18	0.38	0.18	0.38
Repair shop	0.25	0.43	0.14	0.35	0.20	0.40
Major brands	0.68	0.47	0.66	0.47	0.67	0.47

The second and third rows summarize the level and dispersion of prices in the data. Price dispersion is measured as the absolute deviation of prices expressed relative to the period average. On average, stations post a price that is slightly more than half a cent different from the mean. However, this statistic decreased substantially over time as most stations posted the modal price during the last two years of the data.

The remaining summary statistics describe the amenities and services offered. Over the period studied, the average station changed significantly its characteristics. In particular, the proportion of self-service stations increased by ten percentage points, and the proportion of stations with repair shops decreased by the same margin. Similar trends are observed in terms of the presence of large convenience stores, car wash, and opening hours. The proportion of stations selling branded gasoline (i.e., major) remained constant over time. In Quebec City, five major gasoline brands are represented: Shell, Esso, Petro-Canada, Irving, and Ultramar.

These changes reflect the important reorganization of the North American gasoline retail industry observed during the 1990s associated with massive exit and entry of new categories of stations. These changes were mainly caused by technological innovations common to most retailing sectors which increased the efficient size of stations (e.g., automatization of the service, better inventory control systems), as well as changes in the value of certain amenities that occurred throughout the 1980s and 1990s (e.g., mainly decreased in the usage of small repair garages). New environmental regulations of underground storage tanks also explain some of these changes as older stations were forced to replace their equipment. The main consequence of these changes is that newer stations are significantly larger and offer more automatized services than exiting stations.

Figure 1, panel A illustrates the exit process of stations and the associated increase in the market share of major brands. Over the 11 years of the panel, the number of stations in Quebec City dropped by 30 percent, from 382 in 1991 to 283 in December 2001. Major chains are mainly responsible for these changes, which explains their larger market share at the end of the period. In fact, the majors have steadily increased their market share, without expanding their retail network.

Figure 1, panel B also presents two important facts about this market: (i) the small cross-sectional dispersion in prices and (ii) the low level of markups. Over most of the period, the inter-quartile range of prices (i.e., difference between the 75th and

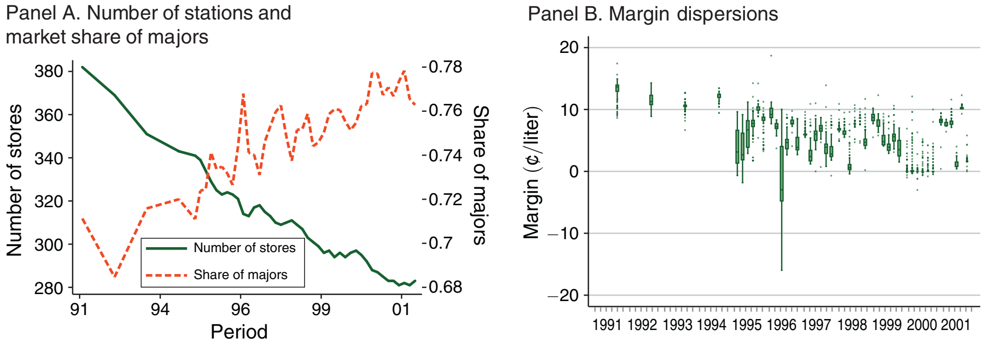

FIGURE 1. TWO KEY MARKET TRENDS: NUMBER OF STATIONS AND MARGIN DISPERSION

the 25th quantile) was smaller than 1 cent per liter, or about 1.5 percent of the average price. The markups, calculated as the ratio of rack price over the average retail price, oscillated between 8 percent and 15 percent for most of the period.[14] This figure also illustrates the presence of two price war episodes during the sample periods. The first episode occurred in the summer of 1996 and was associated with two phenomena: excess capacity associated with the slow reorganization of the market and the introduction by Ultramar, the leading chain, of a low-price guarantee marketing campaign.[15] In 2000, average markups again dropped to zero, after several stations chose to set their price equal to a regulated price floor.[16]

B. *Empirical Distribution of Commuters*

The geography of the market, here defined as the Quebec City Census Metropolitan Area (CMA), is described by a grid of L location areas where people reside and/or work and a road network described by a set of street intersections (or nodes) and road segments (or arcs).[17] The construction of the distribution of commuters across the road network involves two elements: (i) the route used to commute between these points and (ii) the empirical distribution of people over origin and destination locations.

Route Choice.—I consider two types of consumers: local and outside commuters. The location of local commuters corresponds to the centroid of their area of residence and the location of their main occupation (i.e., work or study), denoted by (s, d).[18] Outside commuters, on the other hand, are assumed to travel along the main highways of the city, and therefore each outside commuter origin/destination

[14] The rack price is the price of wholesale gasoline posted at the terminal ramp. It is typically used as a proxy for the marginal cost of gasoline stations.

[15] A similar, but less severe, price war episode also occurred in 1995 before the introduction of Ultramar's low-price guarantee.

[16] The *law on petroleum products* was created in December of 1996 and administrated by the Régie de l'énergie du Québec. The price floor is the sum of the minimum rack price published every week by refiners, taxes, and an estimate of the transportation cost from the terminal to each station. See Carranza, Clark, and Houde (2009) for a more detailed discussion of the regulation and its impact on market structure.

[17] Each location is defined as a dissemination area, which corresponds to the smallest census agglomeration for which data is publicly available. There are over 1,200 Dissemination Areas in Quebec City.

[18] Conceptually, it is feasible to expand this definition to include multiple destinations (e.g., shopping or leisure trips). However it would significantly increase the computation burden of the estimation. Moreover, as long as these

locations correspond to the beginning and end of a particular highway segment (I consider 11 such segments in the empirical analysis).

Local commuters are assumed to choose the route that minimizes the travel time between their home and main occupation locations. This assumption corresponds to the deterministic route choice model used to predict traffic patterns in the transportation literature (see Oppenheim 1995). It generates a single path for each type of consumer and a commuting time and distance, denoted by $r(s,d)$, $t(s,d)$, and $m(s,d)$. This path abstracts from any congestion or unobserved preferences considerations.[19] It is calculated using a version of the Dijkstra's Shortest Path Tree (SPT) algorithm (more details are provided in the online Appendix).

Information on the Quebec City CMA street network was obtained from the CanMap RouteLogistics database DMTI-Spatial (2004), the leading road data provider in Canada. The street network data is extremely fine. It includes more than 30,000 street segments and the average travel time per segment is less than 30 seconds.

Distribution of Origin/Destination Locations.—The distribution of consumers in period t across origin/destination locations is given by T^t_{sd}. I decompose this number into four components: the number of workers, full-time students, unemployed, and the number of outside commuters. For workers and students, the probability of commuting between (s,d) are organized into $L \times L$ matrices. For outside commuters, the pairs of origin/destination correspond to beginning and end points of each highway segment.

I approximate the number of outside commuters by the average number of occupied hotel rooms in the city in period t. This is due to the fact that the transportation survey is related only to local commuters, and I do not have access to any traffic count data on the highway network.

Commuting probabilities are computed from three surveys conducted in 1991, 1996, and 2001 by the Quebec government in the Quebec City CMA. The results of the survey are available in the form of aggregate Origin/Destination (OD) tables, providing the predicted number of commuters between every pair of Traffic Area Zones (TAZs).[20] I use the results of two OD matrices for every transportation mode, representing work and study trips over a full day window. Each matrix is then rescaled to calculate commuting probabilities.[21]

The empirical distributions of local commuters living in each origin location s are computed by combining data from the three most recent censuses (i.e., 1991, 1996, and 2001) and the monthly Canadian Labour Force survey (available at the CMA level) to predict the month-to-month variations in the relevant population measures.

additional trips are close to home or workplace locations, adding them in the model would not add any significant explanatory power.

[19] The no-congestion assumption is realistic in the Quebec City area, since the population is spread over a large territory and the road network is well developed. It has been used also by Thériault et al. (1999) to study the distribution of commuting trips in the Quebec City metropolitan area using similar data.

[20] The aggregate OD matrices are freely available on the government website. The survey reports available on the same website provide further details on the conduct of the survey and the method used to aggregate individual responses (MTQ 2002).

[21] The OD probabilities were further disaggregated into a finer grid in order to predict commuting patterns more accurately. In order to predict traffic in between the two survey dates, the OD probabilities are interpolated assuming a constant growth rate. The online Appendix describes in detail the method used to compute the probability for each location pairs.

In addition, the 2001 survey provides aggregate data on work trips by mode of transportation. I use this additional piece of information to construct the proportion of trips between (s, d) that use a car, which corresponds to the aggregate probability of buying gasoline in the model. I use this information in the estimation to identify some of the preference parameters.

C. *Preliminary Evidence on the Importance of Commuting*

Before turning to the model, it is interesting to compare the empirical population distribution and work commutes with the distribution of gasoline sales. The idea is to correlate the observed distribution of gasoline sales with the predicted demand if consumers are restricted to buy along their commuting path or close to home. To do so, I compute the predicted market share of each store under two assumptions: (i) consumers select at random a store located within a certain distance b of their commuting path, or (ii) consumers buy at random from the set of stores located within a distance b of their home.

Figure 2, panel A illustrates the evolution of the period-by-period correlations for a distance band equal to ½ a minute. The striking feature is that the average correlation generated from the simple commuting buffer model is systematically larger than the one generated from the home buffer model. In fact, the correlation is close to zero or negative using distance from home, while the correlation coefficients from the commuting model all lie around 0.3. The correlation is also stable over time, which suggests that the structure of the road network and the distribution of commuting paths did not change too much between 1991 and 2001.

Figure 2, panel B shows that by increasing the distance band to 5 minutes the average correlation decreases under 0.2 using the distance from commuting path, but increases towards 0.1 for the home distance measure (i.e., low transportation cost). This suggests that the single-address model requires a larger distance band to rationalize the distribution of gasoline sales (i.e., low or negative transportation cost).

In sum, these correlations reveal that the distribution of work commuters is closely related to the distribution of gasoline sales in the market, while the distribution of the nearby population is not. These differences motivate the modeling assumptions regarding the location of consumers in the product space. In particular, they confirm that the locations of gasoline stations reflect the distribution of home to work commuting paths, rather than the population distribution.

II. Demand Model

I model demand for gasoline as a discrete choice problem over $J + 1$ options.[22] In particular, a consumer has the option of buying gasoline from one of the J stores or using an alternative mode of transportation (i.e., option 0). The indirect utility of buying from store j for consumer i is given by

$$(1) \qquad u_{ij} = \begin{cases} X_j\beta + g_i(p_j) + \lambda_1 D\big(r(s_i, d_i), l_j\big) + \xi_j + \epsilon_{ij} & \text{if } j \neq 0, \\ -\lambda_0 C(s_i, d_i) + \epsilon_{i0} & \text{otherwise,} \end{cases}$$

[22] Time subscripts are omitted in this section.

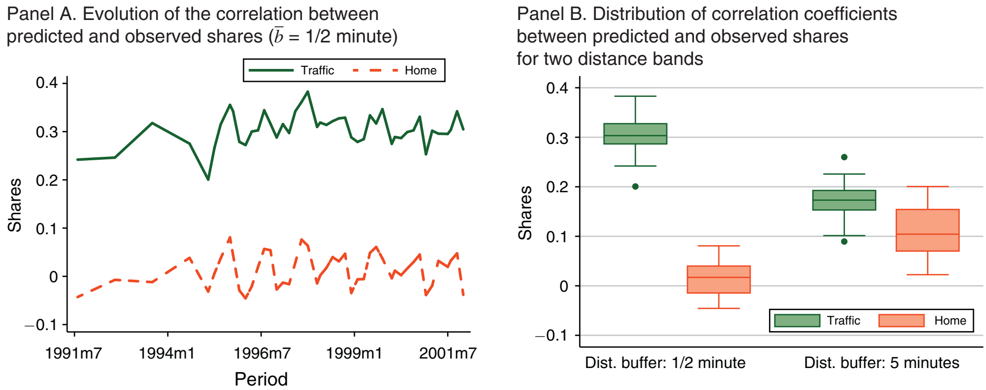

Figure 2. Correlations between Observed Predicted Market Shares from the Home and Commuting Buffer Models

where X_j is a vector of observed station characteristics, p_j is the posted price, $D(r(s,d), l_j)$ is the distance between path $r(s_i, d_i)$ and the location of station j (i.e., l_j), ξ_j is an index of unobserved (to the econometrician) station attributes, $C(s,d)$ is an indicator variable equal to one if consumer i's workplace and home are located in different traffic zones, and ϵ_{ij} is an i.i.d. random utility shock distributed according to a type-1 extreme value distribution.[23] The inclusion of $C(s,d)$ in the indirect utility function is used to proxy for the fact that long-commuters are more likely to commute with a car from home to work, since the region does not have a well-developed system of public transportation.

The set of time varying station characteristics includes the number of gas pumps, the number of service islands, the type of service, the type of convenience store (if any), dummy variables indicating whether the station offers car-repair and/or car wash services, an indicator for brand, and a set of time-dummy variables capturing unobserved period-specific variables (e.g., weather, price of public transportation). Since these variables mainly characterize the type of amenities offered by station j, the unobserved attribute ξ_j measures the set of characteristics of the location which are valued positively by all consumers.

Throughout the paper I assume that consumers have inelastic demand for gasoline. Consumers demand heterogenous quantities of gasoline, denoted by $\overline{q}(r(s,d)) = c_0 + c_1 m(s,d)$, where $m(s,d)$ measures the length of path $r(s,d)$ in kilometers. In this representation, individual consumption is split between heterogeneous work commuting needs and a common fixed quantity c_0, representing leisure and shopping trips. In the empirical analysis, the value of c_1 will be fixed to the average gasoline consumption of a car in a city (i.e., 0.10 liters/km) and c_0 will be estimated. This function determines the size of the market:

$$M = \sum_s \sum_d \overline{q}(s,d)\, T_{s,d}, \tag{2}$$

[23] The Quebec Department of Transportation defines 64 traffic zones for Quebec City.

where as before $T_{s,d}$ is the measure of commuters between s and d.

Although gasoline consumption is assumed to be inelastic, the purchasing elasticity is allowed to vary across income groups. In particular, I use the following linear form for the disutility of prices:

$$g_i(p_j) = p_j(\overline{\alpha} + \alpha y_i),$$

where y_i is the log hourly wage of consumer i.[24]

In the main specification of the model, $D(r(s,d),l)$ measures the cost of deviating from the optimal path $r(s,d)$ in order to purchase gasoline from a store located at point l. There exist many possible distance metrics since $r(s,d)$ corresponds to a vector of street intersections characterizing the path from s to d. In order to keep the model tractable I constructed a measure of the extra time (in minutes) that consumers must incur when stopping at location l on their way to location d. Formally, $D(r(s,d),l)$ is defined by

$$D(r(s,d),l) = t(s,l) + t(l,d) - t(s,d),$$

where $t(s,l)$ and $t(l,d)$ measure the optimal driving time to commute between (s,l) and (l,d) respectively.[25] Notice that the more traditional single-address or home location model is nested in this specification. In particular, the distance metric of this model is given by

$$D^h(s,l) = 2 \times t(s,l).$$

As noted by Petrin (2002) and others, the inclusion of the utility shock, ϵ_{ij}, in the model can generate unrealistic substitution patterns across products due to the embedded independence of irrelevant alternatives assumption. In particular, without other sources of heterogeneity between consumers, the cross-elasticities of substitution depend solely on the relative valuation of products. In the current model, the location of consumers with respect to stations is an important source of heterogeneity between consumers. In particular, consumers are willing to substitute toward products that are "close" to each other, given the distance metric $D(r(s,d),l))$.

[24] The results of a more general discrete-continuous model are available upon request. In that specification, the price disutility function takes an exponential form (i.e., $g_i(p_j) = \sigma(\alpha y_i + \overline{\alpha} p_j + c_1 m(s_i, d_i))\exp(-\alpha p_j))$, and conditional demand is derived using Roy's identity as in Smith (2004) (see also Hanemann 1984 and Dubin and McFadden 1984). The key demand results are unchanged under this alternative model, but the supply results suggest that firms behave as strategic substitutes after the simulated merger (i.e., competing prices go down after a merger). This counterintuitive result is likely due to the fact that the conditional demand function is inelastic for most consumers. The functional form is probably not an accurate description of the market. A more flexible discrete-continuous demand model is unlikely identified using only aggregate data.

[25] In previous versions of the paper I measured $D(r,l)$ by the shortest Euclidian distance between any point in $r(s,d)$ and l_j. This measure provides very similar empirical results, but I found that the one based on commuting time deviations yields a better fit.

Given the distributional assumption on ϵ_{ij}, the conditional probability of buying from store j, for a consumer commuting along route $r(s_i, d_i)$ takes the familiar multinomial logit form with heterogeneous coefficients

$$P_j(r(s_i,d_i),y_i\,|\,\boldsymbol{\delta},\mathbf{p}) = \frac{\exp(\delta_j + \alpha p_j + \mu_{ij})}{1 + \sum_k \exp(\delta_j + \overline{\alpha} p_k + \mu_{ik})}, \tag{3}$$

where $\delta_j = X_j\beta + \xi_j$ is the mean value of store j, and $\mu_{ij} = \alpha p_j y_i + \lambda_0 C(s_i, d_i) + \lambda_1 D(r(s_i, d_i), l_j)$ is the heterogeneous valuation term.

The predicted demand at the station level is obtained by aggregating individual choice probabilities over every OD pair and income:

$$Q_j(\mathbf{p}) = \sum_s \sum_d \int \overline{q}(s,d)\, P_j(r(s,d),y\,|\,\boldsymbol{\delta},\mathbf{p})\, dF(y\,|\,s)\, T_{s,d}, \tag{4}$$

where the implicit dependence of $Q_j(\mathbf{p})$ and δ is omitted. In practice, because the number of OD pairs is very large and income is continuous, I approximate the demand function using simulation techniques (see details in the online Appendix).

Notice that the previous demand equation implicitly assumes that the income distribution is independent of the work destinations of consumers. While this is restrictive, it is justified by a lack of data on the joint distribution of income and origin/destination pairs. The distribution of income is available at the Dissemination Area level from the 2001 census, which allows me to condition the income distribution on the home location of consumers.

III. Estimation Methodology

The set of preference parameters estimated is given by $\boldsymbol{\theta} = \{c_0, \lambda_0, \lambda_1, \alpha, \overline{\alpha}, \beta\}$, and the main dataset used is an unbalanced panel of observed sales and product characteristics. To estimate the model, I adapt the nonlinear GMM estimator developed by Berry (1994) and Berry, Levinsohn, and Pakes (1995). In this section, I first describe the GMM problem, and then discuss the identification of the parameters.

The moment conditions and excluded instrumental variables serve two purposes: (i) correcting for a simultaneity problem between the structural error term ξ_{jt} and prices, and (ii) identifying the nonlinear preference parameters $(c_0, \lambda_0, \lambda_1, \alpha)$. I use two sets of moment conditions to identify the model. Following Nevo (2001), the first set of moment conditions combine Instrumental Variables (IV) with fixed-effects at the station location level. If $\tilde{w}_{jt} = w_{jt} - \frac{1}{n_j}\sum_t w_{jt}$ denotes the within $-j$ transformation of a variable w_{jt}, this first set of empirical moment conditions are given by

$$\overline{g}_n^1(\boldsymbol{\theta}) = \frac{1}{n}\sum_{j,t} g_{jt}^1(\theta) = \frac{1}{n}\sum_{j,t} \tilde{\xi}_{j,t}(\boldsymbol{\theta})\widetilde{\mathbf{W}}_{jt}^1, \tag{5}$$

where n is the number of observations, and $\widetilde{\mathbf{W}}_{jt}^1$ is a vector of predetermined variables including characteristics of stations in $\mathbf{X}_{jt}$, L instrumental variables $\mathbf{Z}_{jt}$, and period dummy variables. Notice that because of the fixed-effects, the moment conditions are expressed over the transitory component of the unobserved quality of a station (i.e., $\tilde{\xi}_{jt} = \xi_{jt} - \bar{\xi}_j$). The fixed component refers to characteristics of the location associated with the road network and the organization of the city. These include, for instance, how easy it is to enter the parking lot of a store, and on which side of the street the station is located. The transitory component is associated with temporary changes made to the quality of the location, such as employee turnover or temporary road repair.

The second set of moment conditions matches the proportion of people who use their car to go to work (or study) in the model with the empirical frequencies from the Origin/Destination survey conducted in Quebec City in 2001. In particular, letting $U_{sd}(\boldsymbol{\theta})$ and $\hat{U}_{sd}$ be the predicted and observed proportion of car users for the pair of origin and destination (s, d), I compute the following moment conditions:[26]

$$(6) \qquad \bar{g}^2_{n_2}(\boldsymbol{\theta}) = \frac{1}{n_2} \sum_{s,d \in \text{Workers}} (U_{sd}(\boldsymbol{\theta}) - \hat{U}_{sd})\, \mathbf{W}^2_{sd},$$

where $\mathbf{W}^2_{sd}$ includes a constant, a dummy variable for long commutes (i.e., $C(s, d)$), and the simulated income for a consumer of type (s, d). In order to calculate this moment condition in practice, the simulated consumers are aggregated at their respective traffic zone area. The number of observations n_2 is therefore the number of aggregate traffic-zone OD pairs included in the survey, which is smaller than the number of simulated consumers. Note that this specification of the micro-moments is similar to an indirect inference approach, which forces the model to replicate the observed linear reduced-form relationship between car-usage, log-income and long-commuting. Using a simulated sample of 3,000 consumers from the 2001 OD survey, this relationship is given by

$$(7) \qquad \hat{U}_i = \underset{(0.033)}{0.247} + \underset{(0.0141)}{0.0495} \times y_i + \underset{(0.0163)}{0.315} \times C_i, \qquad R^2 = 0.47.$$

This indicates that consumers living far from their workplace are significantly more likely to use their car (by 31 percentage points), while a 1 percent increase in income is associated with a 0.049 percentage point increase in the probability of using a car for work commutes. Intuitively, the intercept of this equation identifies the parameter c_0 determining the market size (see equation (2)), while the other two terms identify the marginal utility of commuting and income (i.e., λ_0 and α).

The parameters are estimated by minimizing a quadratic function of the empirical moment conditions. The weighting matrix used in the estimation allows for an arbitrary level of heteroskedasticity in the errors associated with the second moment conditions, and controls for potential serial and spatial correlation in the transitory demand shocks $\tilde{\xi}_{jt}$ (see Section 3 in the online Appendix for more details). Since the number of nonlinear parameters is relatively small, I use the nested fixed-point algorithm proposed by Berry, Levinsohn, and Pakes (1995): given a guess that $\boldsymbol{\theta}$, the vector of structural errors ξ_{jt} solve a system of nonlinear equations equating

[26] The predicted usage rate corresponds to one minus the probability of the outside option. The observed usage rate is calculated from the OD survey described in Section IB.

predicted and observed demand at each store.[27] More details on the procedure are included in the online Appendix.

A. *Identification*

This section discusses the assumptions necessary to identify the key parameters of the model. I focus on the identification of two parameters: the price coefficient α and the transportation cost parameters λ_1. Heuristically, the other nonlinear parameters are identified by imposing moment conditions coming from the transportation survey. The common valuation parameters β are identified by assuming that product characteristics (other than prices) evolve independently of the demand unobservables conditional of location fixed effects (i.e., $\tilde{\xi}_{jt}$).

The identification of α and λ_1 is linked to the more general problem of identifying own and cross elasticities using market-level data. In a linear-demand setting, one must find valid instruments for each product's own and competing prices, to solve a standard simultaneity problem. Here, a simultaneity problem between p_{jt} and ξ_{jt} arises for two reasons. First, gasoline prices are known to adjust frequently, on a weekly or even daily basis (see, for instance, Noel 2007 and Lewis 2008). This introduces measurement error in prices if the adjustment process affects the position of a station in the price distribution. Secondly, since firms and consumers observe the quality index ξ_{jt} when making their decisions, prices will adjust in the short-run to changes in the unobserved product quality.

Since prices enter additively in the expression for the moment conditions, relevant and valid instruments correspond to the set of variables that are correlated with individual prices even after controlling for location and period fixed-effects, and are plausibly independent of unobserved demand shocks $\tilde{\xi}_{jt}$. It is more complicated to discuss identification of the transportation cost parameter λ_1, since it enters nonlinearly in the calculation of the moment conditions.

A common strategy, that exploits the cross-sectional variation of the data, is to assume that the unobserved location attributes are independent of neighboring station characteristics (see, for instance, Pinkse, Slade, and Brett 2002; Davis 2006; and Thomadsen 2005). If entry and location choices are independent of ξ_{jt} but correlated with the observed distribution of consumers, these moment conditions can identify λ_1 by restricting the correlation between market shares and traffic to operate only through the shopping model rather than the unobserved attributes.[28] Of course, this argument is invalid when stores endogenously make their location choices based on traffic patterns and unobserved product attributes. This will typically bias the parameter upwards (i.e., transportation cost close to zero or positive).

[27] As pointed out by Metaxoglou and Knittel (2008) and Dube, Fox, and Su (2011), the numerical solution of this problem is sensitive to the choice of convergence criteria and starting values. During the estimation I set the convergence criteria for the inner-loop equal to 10^{-12}, and use a tighter criteria for the calculation of the standard-errors (i.e., 10^{-14}). The parameter estimates are also robust to starting values. Notice also that the computation cost is important because of the large number of products and time periods (i.e., 46 time periods and over 300 products). To alleviate this burden, I use a derivative-free quasi-newton algorithm that inverts the demand system faster than the contraction-mapping algorithm suggested by Berry, Levinsohn, and Pakes (1995). The model is estimated using Ox (Doornik 2007).

[28] Under this assumption stores tend to cluster around high-traffic areas, rather than high quality location (unobserved). The number or average size of neighboring firms must therefore be orthogonal to the error term.

The identification strategy followed in this paper uses a similar argument, but relies more heavily on the panel dimension of the data. As discussed in Section IA, the reorganization of the industry over the 1990s created sharp changes in the structure of local markets. In particular, the number of stores decreased by nearly 30 percent, and the average size of stores increased significantly. Importantly these changes were driven mainly by technological and regulation changes, and not by factors related to local demand conditions.[29]

The key identifying assumption is that the unobserved attribute of each location evolve independently of the characteristics of nearby stores, conditional on location fixed-effects. This assumption is valid if firms' entry, exit, or remodeling decisions are based on $\bar{\xi}_j$, but not on the realization of the transitory shock $\tilde{\xi}_{jt}$. This assumption is reasonable given the large sunk costs involved in adjusting the characteristics of stations, since most of these changes require installation of new pump systems and involve replacing underground storage tanks.

In the empirical analysis, I compare the robustness of the results to different measures competing product characteristics, and combine different neighborhood definitions. The instruments should be chosen such that they are associated with important changes local market shares, and therefore correlated with own and nearby stations sales. For instance, if a store adopts an amenity that does not affect the market share of its close competitors, the moment conditions will not restrict the magnitude of the transportation cost and the parameter will be imprecisely estimated. Some amenities are also likely to violate the independence assumption either because they involve small sunk costs, or are less dependent on observed factors predicted by the model (i.e., traffic, income, and commuting time).

As argued by Berry, Levinsohn, and Pakes (1995), the characteristics of local competitors are relevant instruments for prices as well, since they enter the equilibrium pricing rule in most pricing models with product differentiation. As we will see below, these variables are, in general, weak instruments for prices, since the market exhibits little cross-sectional price dispersion. In principle, wholesale cost shocks are observed in this market, and could be used to instrument for prices. However, these variables must vary significantly across products since period fixed-effects are included in the model. This rules out rack and crude oil prices.

To circumvent this problem, I exploit the fact that there exists a small amount of cross-sectional dispersion in the posted rack prices at the Quebec terminal, and that the identity of the independent terminal operator, Olco, changed in 1997 following the Ultramar/Sunoco merger (see discussion below in Section VA). I construct instrumental variables that focus only on Sunoco and Olco rack prices since unbranded retailers are more likely to buy gasoline on the spot market.[30] In particular, I construct an instrumental variable that interacts Sunoco's rack price with a dummy variable equal to one if a Sunoco station was located in the same neighborhood as station j. A similar instrumental variable is constructed for Olco, and I

[29] Alternatively, one could use observed events that change locally the distribution of consumers as instrumental variables (e.g., major construction work). In some context this type of variation might be preferable to the market structure variation exploited here, but such data is not available for Quebec City. It is an important avenue for future research to find "natural experiments" that can be used to identify the distribution of taste for differentiated products.

[30] Rack prices of major brands are only weakly correlated with retail prices, conditional on period fixed-effects.

Table 2—Description of Instrumental Variables and Main Specifications

Instrumental variables definitions	IV	Specifications (1)	(2)	(3)	(4)
Count and average competitors' size (street) Average competitors' size (distance) Distance rings: 0:1, 0:33, and 1 KM	5	✓		✓	✓
Olco and Sunoco rack prices × presence Neighborhoods: Street and 1 KM	4		✓	✓	✓
Count and average competitors' size (street) Average competitors' size (distance) Fraction of competitors with: (i) large convenience store; (ii) repair shop; (iii) full service; (iv) company-owned; (v) open 24 hours. Distance rings: 0:1, 0:33, and 1 KM	25				✓

Note: The size of stations is measured by the number of pumps.

use two neighborhood definitions. These two instruments capture two sources of variation that are correlated with price: the presence of a Olco and Sunoco in the neighborhood, and the dispersion of Olco and Sunoco rack prices.

Table 2 presents the definition of the instrumental variables used in the estimation. The first specification uses only the number and average size of competitors along the same street, and the average size of competitors in three different physical distance rings. The common street indicator variable is equal to one if two stores share at least one street and are located within at least 5 minutes driving distance of each other. The second specification uses the rack price of Sunoco or Olco, interacted with dummy variables indicating whether each store is directly competing with Olco or Sunoco. The third one mixes the first two, and the last specification incorporates a larger number of amenities. The third column is the main specification used in the supply analysis below.

IV. Discussion of the Demand Results

A. *Preference Parameter Estimates*

The parameter estimates for the multi-address demand model are reproduced in Table 3.[31] Each column represents a different combination of instrumental variables as defined above.

The estimates of the nonlinear preference parameters are generally robust to the choice of instruments. Not surprisingly, the transportation cost and the price coefficients are the most sensitive. The price coefficient leads to an average store-level price elasticity of demand between -10 and -15 depending on the specification. This high elasticity is consistent with the fact that gasoline stations are not highly differentiated from each other, and suggests that consumers are very price sensitive

[31] Tables 3 and 4 in the online Appendix present the estimates of the parameters measuring the valuation of station characteristics. The number of pumps, which proxies for the ability of a store to serve many customers, enters positively and is the most important characteristic. Most other coefficients have the expected sign, but many of them are not as precisely estimated since store characteristics do not vary significantly over time.

TABLE 3—DEMAND PARAMETER ESTIMATES FROM THE MULTI-ADDRESS MODEL

Variables	(1)	(2)	(3)	(4)
Min. consumption (c_0)	4.5746	4.5290	4.5582	4.5035
	(0.0334)	(0.0442)	(0.0299)	(0.0242)
Commuting distance (c_1)	0.2	0.2	0.2	0.2
	—	—	—	—
Income (α − log($/hour)/100)	0.0396	0.1082	0.0474	0.1677
	(0.0307)	(0.0809)	(0.0327)	(0.0482)
Long commuters (λ_0)	1.4972	1.4356	1.4768	1.3891
	(0.0326)	(0.0683)	(0.0299)	(0.0432)
Transportation cost (λ_1)	−1.2777	−0.5642	−1.0004	−0.3961
	(0.537)	(0.29)	(0.291)	(0.0909)
Price ($\overline{\alpha}$)	−0.2181	−0.1687	−0.1974	−0.1490
	(0.0767)	(0.033)	(0.0378)	(0.0233)
Travel cost − cpl/min.				
α_{25}/λ	5.880	3.387	5.091	2.717
α_{50}/λ	5.886	3.398	5.098	2.733
α_{75}/λ	5.890	3.407	5.103	2.746
Observations	14,263	14,263	14,263	14,263
Number of stores	429	429	429	429
Objective (*J*-stat)	3.07	1.22	6.06	40.5
Degrees of freedom	3	2	7	27
χ^2 Critical value (5%)	7.82	5.99	14.1	40.1

Notes: Standard errors robust to serial and spatial correlations are in parentheses. Each specification also includes location, time, and brand fixed effects, as well as a full set of time-varying station characteristics. The moment conditions used in each specification are described in Table 2.

with regard to their purchasing decisions. Notice also that the model is estimated over a period of high price volatility, which possibly overestimates the price sensitivity of consumers.

The transportation cost parameter is large in magnitude and precisely estimated, except in specification (2) where only the cost-based IVs are used. It is the smallest, in absolute value, in specification (4) where a large number of competitors' characteristics are included in the set of instruments, and largest in specification (1) where only the size and presence of local competitors is considered.

To get a sense of the economic magnitude of the parameters, it is useful to consider how large the price difference between two stores needs to be to justify a 1 minute deviation from the optimal commuting path. The value of this coefficient is obtained by taking the ratio of λ_1 over the price coefficient $\alpha_i = \overline{\alpha} + \alpha y_i$. The median value of the ratio is equal to 5 cents per liter in specification (3). Notice that this value is not very dispersed across income groups, since the estimated income elasticity is very small (i.e., α).

Assuming a purchase of 20 liters, the travel cost implies that the median value of a minute of shopping is \$0.9, or \$54 per hour. This value is four times the average hourly wage in Quebec City. This interpretation is probably too restrictive however, since large deviations from a commuting path likely imply additional time lost in traffic and the uncertainty associated with reoptimizing a new route. Both considerations are ignored from the calculation of the distance measure, and $D(r,l)$ should be considered as a lower bound on true distance.

Table 4—Demand Parameter Estimate from the Single-Address Model Specifications

Variables	(1)	(2)	(3)	(4)
Min. consumption (c_0)	5.0992	5.0399	5.1246	4.9940
	(0.158)	(0.0965)	(0.0691)	(0.0706)
Commuting distance (c_1)	0.2	0.2	0.2	0.2
	—	—	—	—
Income (α − log($/hour)/100)	−0.8433	−0.6463	−0.9318	−0.4574
	(0.87)	(0.409)	(0.36)	(0.226)
Long commuters (λ_0)	3.6236	2.9233	3.9919	2.4657
	(2.53)	(1.07)	(1.06)	(0.534)
Transportation cost (λ_1)	0.5556	0.4235	0.6152	0.3181
	(0.519)	(0.235)	(0.211)	(0.126)
Price ($\overline{\alpha}$)	−0.1498	−0.0918	−0.1054	−0.0630
	(0.0543)	(0.0252)	(0.0275)	(0.0176)
Travel cost − cpl/min.				
α_{25}/λ	−3.344	−4.059	−4.981	−4.426
α_{50}/λ	−3.258	−3.933	−4.793	−4.284
α_{75}/λ	−3.197	−3.844	−4.663	−4.185
Observations	14,263	14,263	14,263	14,263
Number of stores	429	429	429	429
Objective (*J*-stat)	10.2	3.83	14.2	60.6
Degrees of freedom	3	2	7	27
χ^2 Critical value (5%)	7.82	5.99	14.1	40.1

Notes: Standard errors robust to serial and spatial correlations are in parentheses. Each specification also includes location, time, and brand fixed effects, as well as a full set of time-varying station characteristics. The moment conditions used in each specification are described in Table 2.

Perhaps more importantly, the small level of price dispersion reported in Figure 1, panel B suggests that consumers are unlikely to deviate far from their commuting path. However, this does not imply that consumers will patronize a small number of stores, since the average consumer faces 10 stores within one minute of its optimal commuting path. The set of close products also expands with commuting time, as consumers with longer commutes encounter more stations. This implies that long commuters pay lower prices on average, and can better arbitrage price differences between stations.[32]

Table 4 presents the same parameter estimates obtained for the single-address model. The most important difference between the two models is the transportation cost. The point estimate of λ_1 is positive in all specifications, albeit not precisely estimated in all specifications. This should not be surprising given the correlations presented in Section IC. The single-address model requires a small or positive transportation cost to explain observed market shares, which are larger for stores located relatively far from where people live.

In general the variables measuring the characteristics of close-by stations are relatively weak IVs with respect to prices, and are associated with small first-stage *F*-statistics.[33] Specifications (2) and (3) generate *F*-statistics that are larger than the

[32] This prediction of the model is consistent with the empirical results of Yatchew and No (2001), who studied household demand for gasoline in Canada. They found that transaction prices are endogenous: consumers with large unexplained demand residuals tend to pay lower prices.

[33] Table 1 in the online Appendix reproduces the regression results associated with a log-linear demand model to illustrate the role of the different IVs and report the weak instrumental variable tests.

usual weak IV thresholds (i.e., larger than 10), which suggests that the cost-based instrumental variables are highly correlated with prices. This statistic does not give a complete picture of the quality of the instruments, however. The *F*-statistic only measures the correlation between prices and the instruments, while the full model includes four additional nonlinear parameters.

The *J*-statistics reported in Table 3 translate into *p*-values that are all below the 5 percent level, except for the last specification. The moment conditions using the large set of instruments are weakly rejected at 5 percent. Specification (3), which combines (1) and (2), is the preferred specification for the analysis of the model in the remainder of the paper, in part because it yields more precise estimates of the key parameters of the model. Notice also that the *J*-statistics associated with the single-address model are nearly twice as large as the ones reported in Table 3 for the multi-address model, rejecting the model for all but specification (2). The moment conditions are therefore more easily satisfied with the multi-address model.

B. *Analysis of Cross-Price Elasticities*

A common criticism of the multinomial logit model is that substitution patterns are not driven by the proximity of products in characteristics space.[34] In many applications, adding random coefficients does not fix the problem since the variance terms are either not identified or quantitatively important. In this subsection, I analyze the shape of the elasticities of substitution between products to evaluate the importance of spatial differentiation at the parameter estimates (i.e., specification (3) of the multi-address model). The objective is to understand the relationship between observed measures of distance between stations, and the estimated elasticities of substitution. Since the number of elasticity pairs is very large, I consider only a 10 percent random sample.

Table 5 provides summary statistics on the distribution of elasticities and different measures of distance between stores. Looking at the percentiles of the elasticity distribution, the model predicts a very skewed distribution. Most stores have a small number of close competitors: the 1 percent largest cross-elasticities are above 0.4. This is a reflection of the spatial distribution of firms and consumers: 1 percent of store pairs (about 2 or 3 competitors per firm) are located within 1.25 minutes of each-other and have more than 70 percent of traffic in common.[35]

Table 6 investigates the relationship between the estimated substitution patterns and the physical distance between stores. Each column reports the results of a regression of cross-price elasticity pairs on four measures of distance. Columns 1 and 2 also control for the quality of the two stores in order to evaluate the contribution of the Logit distribution.[36] The regression coefficients confirm that the elasticity of substitution decreases sharply in the distance between stations, and increases in the proportion of common traffic. Stations who have a street in common also have larger elasticities of substitution, but this effect is much smaller when we control

[34] See for instance the discussion in Bajari and Benkard (2005) and Berry and Pakes (2007).

[35] The common traffic variable is measured as the proportion of simulated consumers who commute within ½ minute of two stations.

[36] In the multinomial logit model, the elasticity of substitution is proportional to the relative market shares of products.

Table 5—Summary Statistics on Elasticity of Substitution and Measures of Distance between Stations

Variables	Avg.	SD	p_1	p_{10}	p_{50}	p_{90}	p_{99}	Max
Cross-elasticity	0.025	0.094	0	0	0.0013	0.056	0.397	5.26
Common traffic	0.041	0.124	0	0	0	0.111	0.686	1
Driving time	12.7	6.59	0.952	4.82	11.7	22.3	29.5	36.1
Common street	0.0246	0.155	0	0	0	0	1	1

Notes: The sample corresponds to a 10 percent random sample of station pairs. The statistics in column p_k refers to the kth percentile of the distribution of each variable. "Common traffic" measures the traffic share that is common between two stations. "Common street" is an indicator variable equal to one if two stations have a street in common. Driving time is measured in minutes.

Table 6—Regression Results of Elasticity of Substitutions on Measures of Distance between Stations

Variables	(1)	(2)	(3)
Common street	0.182	0.0567	0.0611
	(0.0112)	(0.00779)	(0.00812)
Driving time	−0.00478	−0.00150	−0.00105
	(0.000104)	(7.16e-05)	(7.16e-05)
Share of common traffic		0.449	0.449
		(0.0173)	(0.0178)
Quality index $(\delta_{jt} \times \delta_{kt})$	0.0411	0.0409	
	(0.00158)	(0.00111)	
Constant	−0.0977	−0.154	0.0186
	(0.00599)	(0.00511)	(0.00123)
Observations	428,636	428,636	428,636
R^2	0.225	0.472	0.453

Notes: Standard-errors are clustered at the location level (in parentheses). Dependent variable: Cross-product price elasticity $\varepsilon_{ij} = \Delta^{\%} s_i / \Delta^{\%} p_j$ normalized by the standard deviation of station j's elasticities. The sample corresponds to a 10 percent random sample of station pairs.

for the share of common commuters between the two stations. Not surprisingly, the share of common traffic is the variable that explains the largest fraction of the variance. Moreover, because of the logit assumption the elasticities are increasing in the quality index of locations, but this variable explains only a small fraction of the total variance. Finally, the overall R^2 is close to 0.5, suggesting that driving distances and commuting patterns alone explain a significant fraction of the elasticity of substitution between stores.

V. Analysis of a Vertical Merger: Ultramar and Sunoco

In this section, I use an observed change of ownership between Sunoco and Ultramar stations to compare the counterfactual predictions of the model, with a retrospective analysis of the merger. Since the demand model was estimated without using any supply restrictions, we can use those predictions to evaluate the modeling choices. In the present context, the merger between Sunoco and Ultramar is treated as a natural experiment that changes exogenously the number of vertically integrated competitors in a neighborhood of each Sunoco stations. The section is organized as follows. I first describe the Sunoco-Ultramar transaction in Section VA. Then in Section VB I present the merger simulation analysis. The results of the retrospective analysis are presented in Section VC, followed by a comparative discussion of the merger impact.

A. *A Brief History of the Merger*

In March 1996, Ultramar and Sunoco, two of the largest firms in the Canadian petrol industry, announced their intentions to exchange the ownership of their service stations in Quebec and Ontario.[37] The goal of the transaction was for Ultramar to increase its presence in Quebec by acquiring 127 Sunoco stations, in exchange for 88 sites in Ontario.[38] At the time, Sunoco did not have a refinery in Quebec and chose to concentrate its retail activities in Ontario and western Canada. Ultramar on the other hand adopted the strategy of increasing its dominance in the Quebec market, and distributing locally a larger fraction of its Saint-Romuald refinery's production (near Quebec City). The Canadian Competition Bureau studied the proposed transaction in December 1995 and "deemed [it] unlikely to result in a substantial prevention or lessening of competition."[39]

The implementation of the transaction was progressive. The exact date at which Sunoco stopped supplying gasoline to its Quebec stations is not publicly known. However, data on wholesale prices obtained from MJ Ervin & Associates reveal that Sunoco stopped supplying wholesale gasoline at its Quebec City terminal on January 1, 1997. After this date, an independent wholesaler, Olco, started offering unbranded gasoline in Quebec City, presumably using Sunoco's equipment. The dataset on station characteristics also indicates the first rebranding of a Sunoco station occurred in the second quarter of 1997. Therefore, in what follows, I will assume that the date of Ultramar's acquisition of Sunoco stations is January 1997.

Sunoco had a relatively small presence in Quebec City, with 12 gasoline stations active in December 1996. Still, combined with Ultramar's existing network, this merger represents a sizable increase in market share for Ultramar: from 50 to 62 stores and from 16 percent to 20 percent market share. Table 7 illustrates the timeline of the transaction. The first two rows show that by the end of 1999 all Sunoco stations had been rebranded to Ultramar.

B. *Merger Simulation Analysis*

In order to simulate the counterfactual impact of a merger using an estimated demand system, one must first assume a supply model and recover an estimate of firms' marginal cost functions. The simulated impact is then calculated by solving the pricing game under two or more alternative ownership structures. I describe each of these steps below, and discuss the results of the counterfactual merger between Sunoco and Ultramar using data prior to 1997.

[37] The details of the transactions are described in a newspaper article dated on March 12, 1996 ("Ultramar and Sunoco Swap Parts of Their Service Station Networks," *Business Wire* 1996).

[38] Ultramar did not divest completely in Ontario. It is still active in the eastern part of the province.

[39] See http://www.competitionbureau.gc.ca/eic/site/cb-bc.nsf/eng/00235.html.

Table 7—A Timeline of the Sunoco/Ultramar Transaction in Quebec City

	Jan. 1997	Jan. 1998	Jan. 1999	Jan. 2000
Distribution of Sunocos				
Number Sunoco brand	12	6	3	0
Number Ultramar brand	0	5	8	10
Fraction in ½ minute	0.053	0.052	0.053	0.051
Fraction in 1 minute	0.063	0.061	0.063	0.061
Fraction in 1.5 minutes	0.101	0.100	0.103	0.101
Fraction common street	0.314	0.306	0.309	0.307
Total number of stations	313	309	296	296

Supply-Side Model and Estimation of Marginal Cost.—Upstream suppliers have substantial control over retail prices. In the Quebec market, the major branded suppliers distribute their products mainly through three types of vertical agreements: (i) company-owned stores, (ii) commission, and (iii) lessee contracts. In the first two categories upstream suppliers are responsible for setting the retail prices, and station owners set prices in the latter one.[40]

In 2001, between 52 and 72 percent of branded stations were company-owned. Ultramar, the leading chain, uses exclusively commission contracts with its lessee station operators and therefore acts as a vertically integrated firm. The other firms use lessee contracts in which the wholesale price is set weekly at the station level (i.e., zone pricing). Lessee station owners also negotiate a price-support clause that ensures them of a minimum profit margin. The combination of these two clauses suggests that upstream suppliers are indirectly choosing retail prices at lessee stations. As a result, I will estimate the marginal cost functions assuming perfect resale-price maintenance.[41]

Hastings (2008) provides direct evidence on retail markups that is consistent with the idea that upstream suppliers control most retail prices. Her survey of stations in San Diego reveals that operators systematically add a constant markup to the wholesale price charged by their suppliers, which is consistent with a resale-price maintenance model.

Under resale-price maintenance, each brand supplier f behaves as a vertically integrated firm. Conditional on a vector of prices $\mathbf{p}_{-f}$ posted by competing brands, firm f chooses a price for each of its retail outlets:

$$\max_{\mathbf{p}_f} \sum_{j \in \mathcal{J}_f} (p_j - c_j) Q_j(\mathbf{p}_f, \mathbf{p}_{-f}), \tag{8}$$

where $\mathcal{J}_f$ denotes the partition of stores selling gasoline of brand f. If J denotes the number of retail stores, a Nash equilibrium is described as a set of J first-order conditions simultaneously solving this problem for all firms. Conditional on given

[40] Section 4 in the online Appendix describes the most common vertical contracts used.

[41] In principle, one can use other vertical pricing models to recover the sum of upstream and retail marginal cost by changing the first-order condition. Villas-Boas (2007), and Bonnet and Dubois (2010), for instance, propose non-nested tests to discriminate between resale-price maintenance, two-part tariffs, and linear pricing models. See also Brenkers and Verboven (2006), and Asker (2005) for related empirical analysis of the competitive impact of vertical contracts.

Table 8—Comparison of Estimated Markups under Alternative Supply and Demand Models

Model	Store	Owner	RPM	Collusion	Observed MK
Multi-address	0.09703	0.101	0.1036	0.2684	0.0898
	(0.0106)	(0.0117)	(0.012)	(0.0799)	(0.0551)
Single-address	0.1414	0.1504	0.1576	0.858	0.0898
	(0.0123)	(0.0168)	(0.0163)	(0.0934)	(0.0551)

Notes: Markup = $(p - c)/p$. Top numbers represent sales weighted averages, and numbers in parentheses represent sales weighted standard-deviations (weight = market share). The observed markups are calculated using the observed rack prices. Supply models: (i) store = stations set price independently, (ii) owner = store-owners set prices jointly, (iii) RPM = brand suppliers set prices jointly, (iv) collusion = joint-profit maximization. The sample excludes the years 2000 and 2001, and the summer of 1996.

ownership structure $\mathbf{\Omega}$, we can express the solution of this Betrand-Nash equilibrium in matrix form:

$$\mathbf{Q}(\mathbf{p}) + (\mathbf{\Omega} * \mathbf{\Delta}(\mathbf{p}))(\mathbf{p} - \mathbf{c}) = 0, \tag{9}$$

where $*$ denotes the element-by-element product operator, $\Delta(p)$ is a Jacobian matrix of the demand system (i.e., $\Delta_{ij} = \partial Q_j / \partial p_i$), and $\mathbf{\Omega}$ is a $J \times J$ ownership matrix with element ij equal to one if store i and j sell the same brand of gasoline.

The marginal cost of each store is estimated by inverting the previous equilibrium conditions at the observed prices and ownership structure for each period t, denoted by Ω_t^0 and p_t^0. Table 8 compares the implied markups associated with the third specification of the two demand models and four supply assumptions: (i) individual store pricing, (ii) owner-store pricing, (iii) resale-price maintenance, and (iv) joint profit maximization (i.e., collusion). The last column of the table also reports a common measure of markups used in this industry, calculated using the average Quebec rack price (removing consumption taxes).

Under the RPM assumption, average markups are estimated to be 10 percent. This implies that the average station's marginal cost is one percentage point lower than the posted rack price. The demand and supply models therefore provide very realistic markup estimates.

The difference between the RPM and store-pricing markups, suggests that the market is much closer to the most competitive structure than the collusive one. If stores were setting their prices independently of each other, the demand model would imply a 0.4 percentage point reduction in markups. In comparison, the collusive assumption would imply markups that are 200 percent higher than the RPM model.

The last two lines of Table 8 present the same estimates for the single-address model. The estimated markups appear to be too large relative to the observed markups over the rack prices. The RPM model corresponds to markups of 15 percent; 50 percent larger than the ones predicted by the multi-address model.

I now decompose the estimated marginal cost of stations between retail and wholesale costs. Under RPM, the estimated marginal cost of a station is the sum of the upstream marginal cost w_{ft}, which is common across stores for a given time period, and an idiosyncratic component that reflects the retail marginal cost of the store. This last component can be expressed as a function of observed characteristics

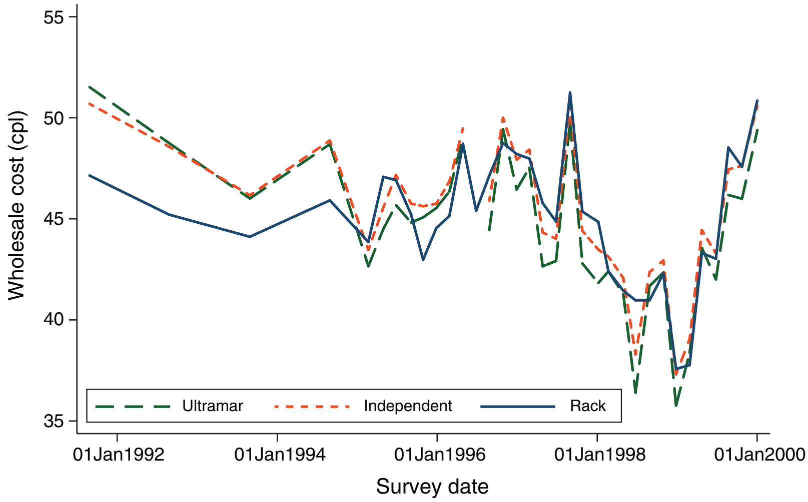

Figure 3. Evolution of the Estimated Upstream Marginal Costs and Rack Price

of the stations Z_{jt} and unobserved cost shocks denoted by η_{jt}. More specifically, we can express the marginal cost of store j as

$$(10) \qquad c_{jt} = p_{jt}^0 + \left(\boldsymbol{\Omega}_t^0 * \boldsymbol{\Delta}(\mathbf{p}_t^0)\right)^{-1} Q(\mathbf{p}_t^0)_{jt}$$

$$= \mathbf{Z}_{jt}\gamma + w_{f_j,t} + \eta_{jt}.$$

The vector of store characteristics $\mathbf{Z}_{jt}$ includes the number of pumps, the type of service, and a set of dummies describing the amenities offered (repair shop, car wash, convenience store, etc.). The regression line also includes indicators for non-company-owned stores, defined by the ownership of the underground storage tank. This last variable captures the idea that lessee stations might add a constant extra margin to the upstream wholesale price.

Equation (10) is estimated by OLS, controlling for supplier and time specific fixed-effects.[42] The coefficients associated with the fixed-effects, denoted by $\hat{w}_{f,t}$, estimate the sum of the upstream marginal cost $w_{f,t}$ and any brand specific intercept. If the latter part is close to zero (i.e., the marginal cost of a representative retailer is equal to the wholesale price), the magnitude of these coefficients should be comparable to the minimum rack price. Figure 3 shows that the estimated upstream cost indeed tracks the posted rack price very well.[43] The estimates also suggest that

[42] The cost function is estimated in the subsample prior to 2000, since the price floor was binding more frequently in 2000 and 2001. I also exclude the 1996 price war episode.

[43] The important gap between $\hat{w}_{f,t}$ and the rack price before 1994 suggests the supply model is not appropriate in the earlier part of the sample. After 1993, average retail margins decreased significantly in the province. Explaining this change is beyond the scope of this paper, but it appears to be related with Ultramar's desire to expand its market share, and possibly a change in the nature of its vertical contracts as this period coincides with the exit of a large number of independently owned branded retailers (e.g., elimination of double marginalization).

independent stations face higher wholesale costs than Ultramar, which is the only company operating a refinery in the city.[44]

Equation (11) presents the remaining estimated coefficients of the combined cost function:

$$
\begin{aligned}
(11)\qquad c_{jt} = {} & \underset{(0.0061)}{-0.0231} \times \text{Reg. pumps}_{jt} + \underset{(0.062)}{117} \times \text{No conv. store}_{jt} \\
& + \underset{(0.08)}{0.171} \times \text{Full service}_{jt} - \underset{(0.062)}{0.152} \times \text{Carwash}_{jt} \\
& + \underset{(0.082)}{0.236} \times \text{Repair shop}_{jt} + \underset{(0.072)}{0.088} \times \text{Lessee dealer}_{jt} \\
& + \hat{w}_{fj,t} + \hat{\eta}_{jt}, \qquad R^2 = 0.92.
\end{aligned}
$$

Overall, station characteristics enter the marginal cost function of stations in a very intuitive way. Self-service stations and large capacity stations (in terms of number of pumps) have significantly lower marginal costs. Also, stores that have an attached convenience store or a car wash have a lower marginal cost. This suggests that these additional lines of business are complementary to gasoline sales, and thus reduce the effective marginal cost of serving an additional customer (i.e., gasoline is a loss-leader product). Repair shop stores however tend to have higher costs, which suggests that their services are not complementary to gasoline sales and/or that the presence of a repair shop proxies for the age of the station. Finally, lessee stations do not have marginal costs statistically different from the ones of company-owned stores selling the same brand. This estimate is sensitive to the sample choice, however. The point estimate is roughly doubled and statistically significant when the 1996 price war episode is included in the sample; suggesting that company-owned stores undercut their rivals by a larger margin.

The previous decomposition can be used to calculate the efficiency gains or losses associated with the ownership change. For instance, we can use the estimated upstream marginal cost of Sunoco and Ultramar to infer what would be the marginal cost of Sunoco stations if the merger occurred in 1996 or 1995 rather than 1997 (i.e., assuming that the merger does not affect the downstream component η_{jt}).

Merger Simulation Results.—In order to compare the simulation and retrospective merger effects, I present the average predicted price differences induced by the merger in four competitive neighborhoods around Sunoco stations (i.e., treatment groups), relative to the average price differences outside of these neighborhoods (i.e., control groups). This difference between the average price changes in the treatment and control groups corresponds to the "difference-in-difference" estimate.

This first set of results is obtained by solving the pricing game under resale-price maintenance, holding the marginal costs fixed to their estimated values. More specifically, for every time period between 1995 and 1997, except for June/July 1996, I calculated the vector predicted prices under the counterfactual merger between Sunoco and Ultramar. The effect of the merger is obtained by subtracting the observed prices,

[44] The upstream costs of other companies lie in between these two points. The difference between the wholesale cost of independents and Ultramar is statistically different from zero and slightly larger than 1 cpl.

Table 9—Aggregate Effects of the Simulated Merger on Prices as a Function of Distance to Sunoco Stations

		Competitive neighborhoods			
		½ Minute	1 Minute	1.5 Minutes	Common street
Aggregate effects	Treatment	0.406	0.343	0.219	0.101
		(0.236)	(0.261)	(0.260)	(0.186)
	Control	0.019	0.019	0.019	0.011
		(0.056)	(0.056)	(0.056)	(0.034)
	Difference	0.388	0.325	0.200	0.090

Notes: Each entry corresponds to the average price effects (in cpl) of the simulated merger between January and December 1996, excluding June, and holding marginal cost fixed. Numbers in parentheses correspond to the standard deviation of the simulated outcomes. The direct effects are calculated using price differences at Ultramar, Sunoco, and Sergaz stations.

which are assumed to be generated by a Nash equilibrium. Table 9 presents the average and standard-deviation price changes for the treatment and control groups.[45]

Recall that the key assumption in the estimation of the treatment effect is that stations outside of Sunoco's competitive neighborhoods are unaffected by the ownership change. In the model, this is not entirely true since some stores (especially Ultramars) react to the merger, but the simulation results suggest that this effect is small (less than 0.02 cpl). Stations located close to a Sunoco station increase their price by a much larger margin, and the difference implies a net change of 0.38 cpl for the smallest neighborhood. The simulated effect also decreases sharply with distance reflecting the fact that competing stations are not greatly affected by the merger.

The merger generates two effects: (i) a direct price increase by Sunoco and Ultramar stations due to the externality that is internalized by the merger, and (ii) an indirect price increase by competitors due to the fact that firms likely behave as strategic substitutes (see, for instance, Deneckere and Davidson 1985). The size of the externality is expected to vary as a function of the distance between Ultramar and Sunoco stations, and the relative isolation of Sunoco stations.

Table 10 decomposes the merger effects into a direct and indirect effect for the smallest competitive neighborhood. The first two specifications (chain) correspond to the resale-price model with and without adjusting for cost efficiencies induced by the merger, as described above. The first line shows that the average price increase at Sunoco stations alone is much larger than the averages presented in Table 9 (i.e., 0.526 cpl versus 0.40 cpl). The price increase is also heterogeneous across stations; the predicted price increase at Sunoco stations goes from 0.2 cpl to nearly 1 cpl. The second column shows that adjusting for cost efficiencies reduces the merger effect only slightly (i.e., 0.05 cpl), since Sunoco is estimated to be almost as efficient as Ultramar. The third column performs the same exercise but assuming that Sunoco stations were setting their prices independently of each other before the merger. This supply model (i.e., owner pricing) is more competitive than resale-price maintenance, and therefore leads to a slightly larger direct effect (i.e., 0.607 cpl).[46]

[45] The standard deviation does not take into account the sampling variance. It is a measure of the heterogeneous reaction of firms to the merger.

[46] For the owner-pricing simulations, I use different marginal-cost estimates obtained from the appropriate first-order condition. That is, Ω_t corresponds to the ownership matrix rather than the brand identity matrix.

TABLE 10—DIRECT AND INDIRECT EFFECTS OF THE SIMULATED MERGER ON PRICES FOR ALTERNATIVE SUPPLY MODELS

		Chain		Owner	Restricted choice set	
		Status quo	Cost adj.	Status quo	Chain	Owner
Sunoco price changes		0.526	0.486	0.607	0.452	0.516
		(0.173)	(0.351)	(0.202)	(0.219)	(0.234)
Direct effect	In ½ m.	0.460	0.429	0.523	0.409	0.459
		(0.200)	(0.331)	(0.243)	(0.212)	(0.237)
	Out ½ m.	0.102	0.103	0.101	0.087	0.086
		(0.105)	(0.107)	(0.104)	(0.132)	(0.131)
	Difference	0.357	0.326	0.422	0.323	0.373
Indirect effect	In ½ m.	0.026	0.025	0.026	0.064	0.066
		(0.105)	(0.107)	(0.104)	(0.132)	(0.131)
	Out ½ m.	0.002	0.002	0.000	0.002	0.001
		(0.003)	(0.003)	(0.009)	(0.004)	(0.008)
	Difference	0.023	0.022	0.025	0.062	0.066

Notes: Each entry corresponds to the average effect of the simulated merger on prices between January and December 1996 (in cpl). Numbers in parentheses correspond to the standard deviation of the simulated outcomes. The direct effects are calculated using price differences at Ultramar, Sunoco, and Sergaz stations. The chain simulations correspond to the model with resale-price maintenance, and the owner simulations correspond to a supply model in which the stations' owners set prices independently.

The next three rows calculate the difference between the average price change at Sunoco and Ultramar stations in the treatment group, relative to Ultramar stations located outside of the ½ minute neighborhood. This yields an estimate of the difference-in-difference coefficient bias if Ultramar stations located far from Sunoco's were used as controls. The model predicts a bias because on average Ultramar stations in the control group increased their prices after the simulated merger (i.e., $0.46 - 0.102 = 0.357$).

The last three rows of Table 10 present the indirect effect of the merger. The model predicts that nearby stores should increase their prices by a small amount relative to other stations in the market (i.e., 0.026 cpl), which is again quite heterogeneous. Intuitively, the reaction of competitors should be much larger. Everything else being equal, competing stations should not be able to maintain a ½ cent price difference due to the large store-level price elasticity. Instead, this small effect is more likely caused by the presence of the logit error term, which introduces another dimension of differentiation that might not be realistic for a market like gasoline. When a store decreases its price, it loses market share to every other option available (including the outside good), which leads to a diffused substitution effect.

One way of attenuating this effect is to reduce the choice-set of consumers, and simulate the equilibrium prices with and without the merger. The results in columns 4 and 5 were obtained by assuming that consumers can purchase only from the 10 closest stores relative to their commuting path.[47] The direct effect of the merger

[47] To simulate the merger under this alternative demand model, I solved the Nash equilibrium price with and without the merger. This is because the marginal cost of stations is recovered using the multinomial model with J_t options, and therefore the prices do not satisfy the equilibrium conditions when consumers are restricted to purchase from a restricted choice-set.

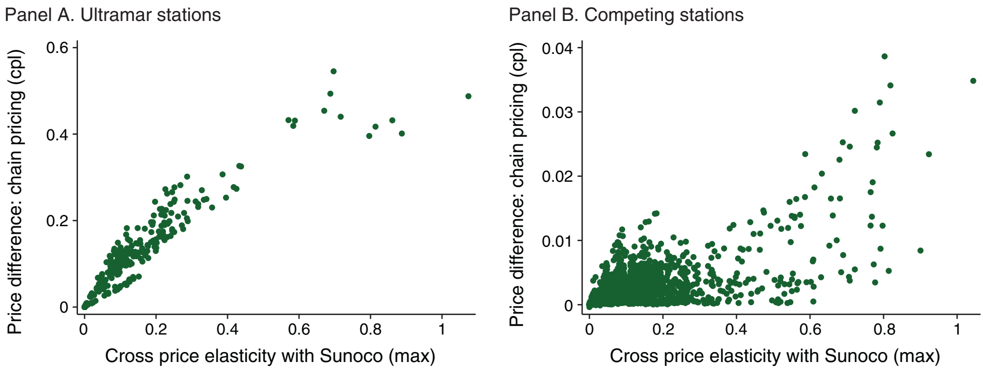

Figure 4. Simulated Price Change Post Merger as a Function of the Elasticity of Substitution with Respect to a Sunoco Station

is almost unchanged, but as expected the average price increase of competitors is doubled (0.064 instead of 0.026). The indirect effect under this alternative assumption is still small in magnitude relative to the direct effect (i.e., about 15 percent). Nonetheless, this result suggests that multinomial logit models that incorporate choice-set restrictions have better pricing predictions, especially when the market includes a large number of products (see, for instance, Goeree 2008).

Figure 4, panels A and B illustrate the direct and indirect effects of the merger as a function of the proximity to Sunoco, measured by the maximum elasticity of substitution with a Sunoco station. Similar figures can be constructed using only physical distance, but the scatter plots would be "noisier" since distance is only one factor determining the elasticity of substitution as discussed in Section IVB. Ultramar stations that compete closely with a Sunoco station react the most, while those with little direct contact experience almost no price increase. The relationship is almost linear, suggesting that the heterogeneity observed in the simulated merger effect is driven mainly by variation in the proximity to Sunoco stations.

The price increase at competing stations shows the same upward sloping relationship, but is a lot flatter. Notice that this figure also shows that the relationship is not strictly monotonic and exhibits a larger variance. A fraction of competing stations increase their prices slightly without directly competing with a Sunoco station, while others located closely do not react to the merger. This heterogeneity represents two phenomena. First, a few stations located near an Ultramar station, but relatively far from a Sunoco station, increase their price following the merger. Since the scale of the merger is small (i.e., only 12 Sunoco stations), this effect is not important and stations in the control group do not increase their prices significantly after the merger. Second, single-owner stations tend to react differently to the merger than multiproduct firms, and in particular have a more pronounced response to the price increase of a close competitor.

C. *Retrospective Analysis*

In order to estimate the observed price effect of the Sunoco/Ultramar merger, I conduct a difference-in-difference analysis similar to one performed by Hastings (2004).

The identification strategy relies on the idea that the merger between two networks of retail stores creates sharp changes in the structure of local markets. In Hastings' case, the acquisition of Thrifty by ARCO implied the loss of an independent competitor for stores directly competing with a Thrifty station. Moreover, because this type of merger is negotiated nationally, the change of ownership is not endogenous to the unobserved local demand conditions after controlling for store fixed-effects.

In the present context, the loss of a Sunoco did not occur on January 1997, because rebranding took place over a two-year period. Instead, the immediate change occurred vertically since Ultramar became the sole supplier of all Sunoco stations. If this new upstream supplier was able to fully or partially control the retail price, through wholesale price discrimination or some form of resale-price maintenance, the post merger prices at Sunoco stations should be set less competitively than before.

The acquisition might also have created other changes to the cost structure of Sunoco stations. In particular, as we saw above, Ultramar has a lower marginal cost of supplying gasoline because of the proximity of its refinery. The objective of the retrospective analysis is then to quantify the magnitude of the anti-competitive effect, net of any efficiency gains or losses associated with the merger.

Assuming, like Hastings (2004), that the ownership change affected prices only locally, a natural treatment group includes Sunoco stations and their immediate competitors. The control group is defined as stations outside of these neighborhoods. The last four rows in Table 7 describe the distribution of active stations located in different neighborhoods around each Sunoco station that was merged with Ultramar. Importantly, the competitive neighborhoods are the same as the ones used previously in the merger simulation: three driving distance time buffers and a discrete measure based on the connectivity of stations on the road network.

Notice from Table 7 that former Sunoco stations that were progressively rebranded are in the treatment group. In the empirical analysis, I pool these two types of stations rather than estimating separate treatment effects because of sample size restrictions.

The fraction of observations in the treatment group is smaller than in Hastings (2004). However, contrary to Hastings' sample, I observe the universe of prices in the market for a longer period of time. This allows me to separate the merger effect into a direct effect on Sunoco and Ultramar stations, and a competitive effect on other stations.

I estimate the following two fixed-effect regression models using station level prices p_{jt} as dependent variable:

$$p_{jt} = \gamma N_j^d \times T_t + \mathbf{Z}_{jt}\beta + \mu_j + \tau_t + u_{jt}, \tag{12}$$

$$p_{jt} = \gamma_0 N_j^d \times T_t \times S_{jt} + \gamma_1 N_j^d \times T_t \times (1 - S_{jt}) + \mathbf{Z}_{jt}\beta + \mu_j + \tau_t + u_{jt}, \tag{13}$$

where N_j^d is a dummy variable equal to one for stations located in a neighborhood d of a Sunoco station, T_t is a dummy variable equal to one after the ownership change, and S_{jt} is an indicator variable for stations that are supplied by Ultramar after the merger. All specifications include location and time-fixed effects (i.e., μ_j and τ_t). The

Table 11—Direct and Competitive Effects of the Sunoco/Ultramar Merger

Variables	95/97 (1)	95/97 (2)	95/97 (3)	95/97 (4)	95/98 (5)	95/98 (6)
Merger ½ m.		0.501***		0.202**		0.211***
× Direct		(0.121)		(0.077)		(0.079)
Merger ½ m.		0.0384		−0.106		0.325*
× Competitive		(0.178)		(0.164)		(0.169)
Merger ½ m.	0.433***		0.156**		0.229***	
	(0.119)		(0.079)		(0.074)	
Constant	63.40***	63.41***	62.42***	62.43***	54.91***	54.90***
	(1.414)	(1.415)	(1.293)	(1.294)	(1.248)	(1.248)
Observations	5,778	5,778	5,457	5,457	7,287	7,287
R^2	0.785	0.785	0.840	0.840	0.917	0.917

Notes: Standard errors are clustered at the location level. Dependent var.: Store prices. Merger var.: Sunoco in neighborhood × Year $\geq$ 1997. The direct effect includes post-merger Ultramar brands: Sunoco, Sergaz, and Ultramar. Columns 2–6 exclude observations from the 1996 price war.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

average effect of the merger is denoted by γ, while the direct and competitive effects are denoted by γ_0 and γ_1 respectively. The control variables Z_{jt} include time-varying station characteristics (number of pumps, type of service, amenities, brand, etc.) and sales volume.[48]

Table 11 presents the regression results corresponding to equation (12) (columns 1, 3, and 5) and equation (13) (columns 2, 4, and 6) for the smallest neighborhood definition (i.e., ½ minute). The table includes estimation results over three different subsamples: (i) 1995–1997 including the 1996 price war (columns 1–2), (ii) 1995–1997 excluding the 1996 price war (columns 3–4), and (iii) 1995–1998 excluding the 1996 price war (columns 5–6). The 1997–1998 period includes a larger fraction of rebranded stations, which complicates the interpretation of the results. It is also interesting to compare the results with and without the 1996 price war, because Ultramar was largely responsible for causing the war. As a result, the merger effect tends to be larger when the sample includes the price war episode.

The aggregate results in columns 1, 3, and 6, reveal that prices were indeed higher after the merger in the treated neighborhoods, suggesting that the anti-competitive effect of the merger dominated any efficiency gains. Depending on the sample, the average price difference ranges from 0.15 cpl to 0.45 cpl. Since the average margins were about 4 cpl during this period, this represents at most a 11 percent increase in retail margins, or a 0.7 percent price increase. By comparison, these point estimates correspond to a price increase of at most 1.75 cents per gallon, which is smaller than the results reported in Hastings (2004), but larger than what Taylor, Kreisle, and Zimmerman (2010) found (i.e., 5 and 1 cents per gallon, respectively).

[48] The difference-in-difference specification will estimate consistently the effect of the merger if the announcement does not coincide with other events affecting the price of stations in the treatment groups. In the current example, one such assumption is that the price floor regulation, which was implemented in 1997 as well, affected all stations symmetrically (i.e., treated and controlled).

The disaggregate results clearly show that most of this effect is due to Sunoco and Ultramar posting higher prices after the merger. The direct effect of the merger ranges between 0.2 cpl and 0.5 cpl, while the competitive effect is imprecisely estimated and sometimes negative (see column 4). In general, the competitive effect is found to be significantly different from zero only when the 1998 sample is included. Although the sample is not rich enough to clearly identify the rebranding effect, this result suggests that the rebranding of a larger number of stations might be causing a bigger competitive reaction.

The estimated treatment effect of the merger is robust across neighborhood definitions. Table 12 illustrates the robustness of the results to different assumptions about the size of the competitive neighborhoods. The magnitude of the merger effect is, in general, decreasing in the size of the neighborhood, reflecting the fact that competing stations were less affected by the merger. As hinted by the results in Table 11, however, the decline in the effect is much less pronounced when the sample includes the rebranding periods (see columns 5–8).

D. *Discussion*

The results of the retrospective exercise to a large extent confirm the merger simulation results. The change in ownership between Sunoco and Ultramar caused a significant price increase of up to ½ cpl in the neighborhood of Sunoco stations, which corresponds to a 10 percent increase in retail margins. As the structural model shows, most of this effect is attributed to the fact that the new merged entity internalized the cannibalization of market shares caused by independent pricing behavior. Since this externality is larger when Ultramar and Sunoco stations compete for the same consumers, the price increase is increasing in the degree of substitutability between Ultramar and Sunoco stations.

This is not to say that the reduced-form and the structural estimates fully agree. At least in one subsample (i.e., 95/98), the retrospective analysis suggests a much larger indirect price effect than the merger simulation. The observed merger effect also does not diminish in distance as fast as what the model predicts. The merger simulation results show that the net effect should go from 0.38 to 0.09 cpl in the largest neighborhood, while the difference-in-difference model suggests a decline from 0.43 to 0.26 (i.e., columns 1–5). These differences are explained in part by the fact that the simulated reaction of competitors is small, due to the distributional assumption on consumers' idiosyncratic tastes.

Similarly, if we drop the 1996 price war from the difference-in-difference estimation, the net effect of the merger is smaller than what the model predicts (i.e., 0.2 versus 0.37). Part of this might be due to the fact that the efficiency gains were larger than the ones that were estimated.

On the other hand, the retrospective analysis does not provide a complete picture of the merger effect. Although all specifications suggest that the merger induced higher prices, the actual magnitude of the change is sensitive to the sample choice. This sheds some doubts on the validity of the identifying assumptions: the difference-in-difference estimate might be capturing the end of the price war rather than the ownership change (i.e., stations competing indirectly with Ultramar received two "treatments").

Table 12—Aggregate Effect of the Merger for Alternative Neighborhood Definitions

Variables	95/97 (1)	95/97 (2)	95/97 (3)	95/97 (4)	95/98 (5)	95/98 (6)	95/98 (7)	95/98 (8)
Merger in ½ min.	0.433*** (0.119)				0.229*** (0.074)			
Merger in 1 min.		0.419*** (0.110)				0.216*** (0.064)		
Merger 1.5 min.			0.290*** (0.108)				0.223*** (0.059)	
Merger in street				0.266*** (0.084)				0.175*** (0.067)
Constant	63.40*** (1.414)	63.41*** (1.413)	63.40*** (1.411)	63.57*** (1.379)	54.91*** (1.248)	54.91*** (1.247)	54.89*** (1.245)	62.48*** (1.206)
Observations	5,778	5,778	5,778	5,778	7,287	7,287	7,287	7,287
R^2	0.785	0.785	0.785	0.785	0.917	0.917	0.917	0.917

Notes: Standard errors are clustered at the location level. Dependent var.: Store prices. Merger var.: Sunoco in neighborhood $\times$ Year $\geq$ 1997. Columns 5–8 exclude observations from the 1996 price war. Competitive neighborhoods are measured in driving minutes and common street.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

A more general challenge for the difference-in-difference approach is the identification of a valid control group. As the simulation results showed, the treatment effect underestimates the true direct effect of the merger, since Ultramar stations in the control group increased their prices by a sizable margin. Overall, the results suggest that this bias is not too large in the current case, because Sunoco only operated a small fraction of stations in the market. However this bias should be important in larger merger cases involving spatially differentiated firms. When consumers are commuting between locations, the cannibalization of market shares is expected to spill-over in larger neighborhoods, and the number of stations unaffected by the merger can be very small.

Finally, the interpretation of the retrospective analysis through the lens of the merger simulation analysis shed light on the recent debate surrounding the replicability of retrospective merger analysis. For instance, Taylor, Kreisle, and Zimmerman (2010) recently failed to replicate the magnitude of the results obtained by Hastings (2004) concerning the merger between ARCO and Thrifty. Hastings' results suggested that the merger induced price increases of 5 cents per gallon, while Taylor et al. obtained a much smaller estimate of about 1 cent per gallon. Both papers used the same estimation methodology, but relied on different price data: Hastings used quarterly information on a 20 percent random sample of stations, while Taylor et al. used average monthly price data from a larger but less representative cross-section of stations. Importantly, neither dataset included Thrifty stations (i.e., the independent chain), and the estimation relied only on the effect of the merger on ARCO and competing firms' prices.

The results of the merger simulation analysis shows that small differences in the cross-section of stations can have large consequences on the estimated price change. The impact of the merger on prices is highly heterogeneous across stations and local markets. Therefore, the estimated differential effect varies greatly depending on the choice of controls, and the availability of information on stores that were directly affected by the merger (i.e., Sunoco and Ultramar in this case).

For instance, since the model predicts that most of the price increases are associated with the two merging firms, any sampling scheme that leaves out the most competitive player pre-merger will underestimate the effect of the merger. In the current example, dropping Sunoco stations from the averages reported in Table 10 would decrease the estimated direct effect by more than half. Similarly, the point estimate is expected to be very sensitive to the presence of Ultramar stations in the treatment group. A dataset that over-samples Ultramar in the control group would lead to a downward bias, while a dataset that over-samples Ultramar in the treatment group would overestimate the effect.

VI. Conclusion

In this paper I develop and estimate a novel model of demand for spatially differentiated products, applied to retail gasoline. My approach contributes to the literature on spatial differentiation by formally modeling commuting paths as the "locations'' of consumers. This extension of the standard single-address model generates substitution patterns that depend in an intuitive way on the structure of the road network and the direction of traffic flows.

The methodology combines computing tools from the transportation Geographic Information System literature, and econometric methods developed to estimate discrete choice models of demand. The model is estimated using panel data and covers an important period characterized by a large North American reorganization of retail networks. This feature of the data enables me to identify the structural parameters of the model, even after controlling for important unobserved characteristics of store locations.

The results validate the modeling choices in several ways. The distribution of gasoline sales within the market is shown to be poorly correlated with the distribution of local population, and significantly more so with the distribution of work commuters. This directly translates into a small and positive estimate of the transportation cost parameter in the traditional model, which implies very little spatial differentiation. Since the degree of substitutability between locations feeds directly into predictions of mark-ups and prices, the results from the multi-address model differ sharply in terms of the evaluation of market power. Importantly, the estimated markups coming from the model match very well the observed profit margins over the spot wholesale price for gasoline.

The demand specification is also validated by comparing the pricing implications of the model with the estimated impact of a consummated merger between Ultramar and Sunoco. The difference-in-difference estimate roughly equates the average price increase predicted by the counterfactual simulation of the merger. This result is important as it provides support for the merger simulation methodology.

The comparative analysis also highlights the relative strengths of reduced-form and structural methods. While the treatment effect estimates are shown to be sensitive to definition of control group and sample size, it also serves as a useful and transparent point of comparison to validate the assumptions of structural models. In the present context, the main failure of the merger simulation is the inability of the model to predict a sizable price reaction from competing firms post merger. This problem is linked with the distributional assumption of consumer tastes, which tends to underestimate the elasticity of substitution between close competitors despite a large estimated transportation cost. This issue is common to most empirical models of product differentiation that introduce idiosyncratic utility shocks. Further research is therefore needed to relax these assumptions, while keeping the tractability of the framework developed by Berry, Levinsohn, and Pakes (1995).

REFERENCES

Ackerberg, Daniel, Lanier Benkard, Steven Berry, and Ariel Pakes. 2007. "Econometric Tools for Analyzing Market Outcomes." In *Handbook of Econometrics*, Vol. 6A, edited by James J. Heckman and Edward E. Leamer, 4171–276. Amsterdam: North Holland.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2010. "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives* 24 (2): 3–30.

Ashenfelter, Orley, Daniel Hosken, and Matthew Weinberg. 2009. "Generating Evidence to Guide Merger Enforcement." *Competition Policy International* 5 (1): 67–85.

Asker, John. 2005. "Diagnosing Foreclosure due to Exclusive Dealing." NYU-Stern Working Paper EC-04-36. http://pages.stern.nyu.edu/~jasker/Paper1_ExclusiveForeclosure.pdf.

Bajari, Patrick, and C. Lanier Benkard. 2005. "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach." *Journal of Political Economy* 113 (6): 1239–76.

Baker, Jonathan B., and Timothy F. Bresnahan. 1985. "The Gains from Merger or Collusion in Product-Differentiated Industries." *Journal of Industrial Economics* 33 (4): 427–44.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2): 242–62.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.

Berry, Steven, and Ariel Pakes. 1993. "Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis." *American Economic Review* 83 (2): 247–52.

Berry, Steven, and Ariel Pakes. 2007. "The Pure Characteristics Demand Model." *International Economic Review* 48 (4): 1193–1225.

Bonnet, Céline, and Pierre Dubois. 2010. "Inference on Vertical Contracts between Manufacturers and Retailers Allowing for Nonlinear Pricing and Resale Price Maintenance." *RAND Journal of Economics* 41 (1): 139–64.

Brenkers, Randy, and Frank Verboven. 2006. "Liberalizing a Distribution System: The European Car Market." *Journal of the European Economic Association* 4 (1): 216–51.

Carranza, Juan Esteban, Robert Clark, and Jean-François Houde. 2009. "Price Controls and Competition in Gasoline Retail Markets." University of Wisconsin-Madison Working Paper. http://www.ssc.wisc.edu/~houdejf/floor_v6.pdf.

Claycombe, Richard J. 1991. "Spatial Retail Markets." *International Journal of Industrial Organization* 9 (2): 303–13.

Dafny, Leemore. 2009. "Estimation and Identification of Merger Effects: An Application to Hospital Mergers." *Journal of Law and Economics* 52 (3): 523–50.

Dafny, Leemore S., Mark G. Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a Premium on Your Premium? Consolidation in the U.S. Health Insurance Industry." *American Economic Review* 102 (2): 1161–85.

Davis, Peter. 2006. "Spatial Competition in Retail Markets: Movie Theaters." *RAND Journal of Economics* 37 (4): 964–82.

Deneckere, Raymond, and Carl Davidson. 1985. "Incentives to Form Coalitions with Bertrand Competition." *RAND Journal of Economics* 16 (4): 473–86.

DMTI-Spatial. 2004. "CanMap RouteLogistics, version 8.2." [Electronic resource]. Markham, Ontario. Queens University Library. http://db.library.queensu.ca/geospatial_data/ (accessed 2004).

Doornik, Jurgen A. 2007. *Object-Oriented Matrix Programming Using Ox*. London: Timberlake Consultants Press.

Dube, Jean-Pierre, Jeremy Fox, and Che-Lin Su. 2011. "Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation." Unpublished.

Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52 (2): 345–62.

Goeree, Michelle Sovinsky. 2008. "Limited Information and Advertising in the U.S. Personal Computer Industry." *Econometrica* 76 (5): 1017–74.

Hanemann, W. Michael. 1984. "Discrete-Continuous Models of Consumer Demand." *Econometrica* 52 (3): 541–61.

Hastings, Justine S. 2004. "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California." *American Economic Review* 94 (1): 317–28.

Hastings, Justine. 2008. "Wholesale Price Discrimination and Regulation: Implications for Retail Gasoline Prices." Unpublished.

Hastings, Justine S., and Richard J. Gilbert. 2005. "Market Power, Vertical Integration and the Wholesale Price of Gasoline." *Journal of Industrial Economics* 53 (4): 469–92.

Hausman, Jerry A., and Gregory K. Leonard. 2002. "The Competitive Effects of a New Product Introduction: A Case Study." *Journal of Industrial Economics* 50 (3): 237–63.

Hausman, Jerry, Gregory Leonard, and J. Douglas Zona. 1994. "Competitive Analysis with Differenciated Products." *Annales d'Économie et de Statistique* (34): 159–80.

Hortaçsu, Ali, and Chad Syverson. 2007. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *Journal of Political Economy* 115 (2): 250–301.

Houde, Jean-François. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.5.2147.

Kent Marketing. 1991–2001. "Petroleum Volume Data Survey." Bimonthly data available for purchase. http://www.kentmarketingservices.com/dnn/Default.aspx?tabid=75 (accessed 2004).

Lach, Saul. 2002. "Existence and Persistence of Price Dispersion: An Empirical Analysis." *Review of Economics and Statistics* 84 (3): 433–44.

LaLonde, Robert J. 1986. "Evaluating the Econometric Evaluations of Training Programs with Experimental Data." *American Economic Review* 76 (4): 604–20.

Lewis, Matthew. 2008. "Price Dispersion and Competition with Differentiated Sellers." *Journal of Industrial Economics* 56 (3): 654–78.

Manuszak, Mark D. 2001. "The Impact of Upstream Mergers on Retail Gasoline Markets." Unpublished.

Manuszak, Mark D., and Charles C. Moul. 2009. "How Far for a Buck? Tax Differences and the Location of Retail Gasoline Activity in Southeast Chicagoland." *Review of Economics and Statistics* 91 (4): 744–65.

McManus, Brian. 2007. "Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee." *RAND Journal of Economics* 38 (2): 512–32.

Metaxoglou, Konstantinos, and Christopher R. Knittel. 2008. "Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings." Unpublished.

Ministère des Transports du Québec (MTQ). 2002. *Mobilité des personnes dans l'agglomération de Québec: Sommaire des résultats de l'enquete origine-destination 2001.* Gouvernement du Québec. Data available for 1991, 1996, and 2001 [in French]. http://www.mtq.gouv.qc.ca/portal/page/portal/ministere/ministere/recherche_innovation/modelisation_systemes_transport/enquetes_origine_destination/quebec (accessed 2004).

MJ Ervin. 1991–2001. "Weekly Retail and Wholesale Price Survey." http://www.kentmarketingservices.com/dnn/Default.aspx?tabid=134 (accessed 2004).

Nevo, Aviv. 2000. "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry." *RAND Journal of Economics* 31 (3): 395–421.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307–42.

Nevo, Aviv, and Michael D. Whinston. 2010. "Taking the Dogma out of Econometrics: Structural Modeling and Credible Inference." *Journal of Economic Perspectives* 24 (2): 69–82.

Noel, Michael D. 2007. "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market." *Journal of Industrial Economics* 55 (1): 69–92.

Oppenheim, Norbert. 1995. *Urban Travel Demand Modeling*. Hoboken, NJ: Wiley-Interscience.

Peters, Craig. 2006. "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry." *Journal of Law and Economics* 49 (2): 627–49.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4): 705–29.

Pinkse, Joris, Margaret E. Slade, and Craig Brett. 2002. "Spatial Price Competition: A Semiparametric Approach." *Econometrica* 70 (3): 1111–53.

Raith, Michael. 1996. “Spatial Retail Markets with Commuting Consumers.” *International Journal of Industrial Organization* 14 (4): 447–63.
Regie de l’energie du Quebec. 1997–2001. “Evolution des prix de détail de l’essence ordinaire.” [in French]. http://www.regie-energie.qc.ca/energie/petrole_tarifs.php (accessed 2004).
Simpson, John, and Christopher Taylor. 2008. “Do Gasoline Mergers Affect Consumer Prices? The Marathon Ashland Petroleum and Ultramar Diamond Shamrock Transaction.” *Journal of Law and Economics* 51 (1): 135–52.
Smith, Howard. 2004. “Supermarket Choice and Supermarket Competition in Market Equilibrium.” *Review of Economic Studies* 71 (1): 235–63.
Statistics Canada. 1991, 1996, and 2001. “Aggregate Statistics.” University of Toronto. http://datalib.chass.utoronto.ca/major/censusag.htm (accessed 2004).
Statistics Canada. 1991–2001. “Labour Force Survey.” Monthly statistics between 1991 and 2001. http://www23.statcan.gc.ca:81/imdb/p2SV.pl?Function=getSurvey&SDDS=3701&lang=en&db=imdb&adm=8&dis=2 (accessed 2004).
Taylor, Christopher T., and Daniel S. Hosken. 2007. “The Economic Effects of the Marathon-Ashland Joint Venture: The Importance of Industry Supply Shocks and Vertical Market Structure.” *Journal of Industrial Economics* 55 (3): 419–51.
Taylor, Christopher T., Nicholas M. Kreisle, and Paul R. Zimmerman. 2010. “Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California: Comment.” *American Economic Review* 100 (3): 1269–76.
Thériault, Marius, Marie-Hélène Vandersmissen, Martin Lee-Gosselin, and Denis Leroux. 1999. “Modelling Commuter Trip Length and Duration Within GIS: Application to an O-D Survey.” *Journal of Geography Information and Decision Analysis* 3 (1): 40–56.
Thomadsen, Raphael. 2005. “The Effect of Ownership Structure on Prices in Geographically Differentiated Industries.” *RAND Journal of Economics* 36 (4): 908–29.
Tourisme Quebec. 2001. “Portrait statistique de l’hotellerie au Quebec” [in French]. http://www.tourisme.gouv.qc.ca/publications/categorie/hebergement-quebec-55.html (accessed 2004).
Villas-Boas, Sofia Berto. 2007. “Vertical Relationships between Manufacturers and Retailers: Inference with Limited Data.” *Review of Economic Studies* 74 (2): 625–52.
Weinberg, Matthew C., and Daniel Hosken. 2009. “Using Mergers to Test a Model of Oligopoly.” Unpublished.
Werden, Gregory J., and Luke M. Froeb. 1994. “The Effects of Mergers in Differentiated Products Industries: Logit Demand and Merger Policy.” *Journal of Law, Economics, and Organization* 10 (2): 407–26.
Yatchew, Adonis, and Joungyeo Angela No. 2001. “Household Gasoline Demand in Canada.” *Econometrica* 69 (6): 1697–1709.