

Misallocation in Differentiated Product Markets: Evidence from Collusion under Uniform Pricing¹

Daniel Chaves² Robert Clark³ Marco Duarte⁴
Jean-François Houde⁵

November 2025

Abstract

We develop a method to quantify misallocation from collusion, accounting for demand-side frictions using a novel model of spatial differentiation. By estimating a non-cooperative equilibrium model, we recover marginal costs and decompose inefficiencies due to collusion versus other frictions. We apply this to Quebec City’s retail gasoline market, which transitioned to uniform-price collusion. We show that colluding on a common focal price causes substantial misallocation, accounting for 80% of allocative efficiency gains from marginal-cost pricing. Measures ignoring demand-side frictions overstate collusion’s effect by a factor of eight. Finally, uniform-price collusion introduces dynamic inefficiencies by maintaining excess capacity and reducing investment.

¹We thank Dave Byrne, Matt Mitchell, Diego Restuccia, Otto Toivanen, Nahim Zahur, and seminar participants at Penn State, Virginia, Rochester, FGV-EESP, Pontificia Universidad Catolica de Chile, and Universidad Los Andes, as well as attendees of the 2024 and 2025 Annual Meetings of the CEA, the 2024 IIOC, the 2024 meetings of the SEA, the 2025 Rotman Conference on Economic Analysis and Policy, and the 2025 Bates White conference, for their helpful comments.

²Western University

³Queen’s University

⁴University of North Carolina-Chapel Hill

⁵University of Wisconsin-Madison and NBER

1 Introduction

Traditional antitrust analysis focuses primarily on the welfare loss generated by higher collusive prices, as described by Harberger (1954). However, collusion may create additional welfare losses by distorting the allocation of production across firms. In static terms, a cartel can allocate sales inefficiently—diverting output toward high-cost producers and away from low-cost ones. Beyond this, collusion can have dynamic consequences: by muting competitive pressures, it may reduce firms’ incentives to invest, enter, or exit the market. The distortions stemming from these market structure effects can persist well after a cartel has formally dissolved. The central research question is therefore not only how much collusion raises prices, but also how it misallocates production and investment across firms.

One approach for diagnosing misallocation comes from the macroeconomics literature, particularly Hsieh and Klenow (2009) and Restuccia and Rogerson (2017).¹ Their framework evaluates the degree of misallocation by examining the dispersion of total factor productivity across plants within an industry or an economy. The extent of misallocation is then benchmarked by comparing the observed industry or economy to an undistorted or efficient benchmark, such as the U.S. economy.

Asker, Collard-Wexler, and De Loecker (2019) propose an alternative method that relies on the dispersion of marginal costs, measured using plant-level accounting data. Rather than inferring distortions indirectly from output and inputs, they assess efficiency by comparing the observed allocation of production to an ideal situation in which all output would be produced by the lowest-cost firms. This method provides direct evidence on the degree of misallocation arising from collusion, and they apply it to study the OPEC cartel.

Both approaches ignore demand-side frictions and imperfect competition that limit the scope of the efficiency gains that can be realized from eliminating collusion. Failure to consider a proper non-cooperative benchmark that accounts for market power and differentiation can overstate the extent of misallocation. A second crucial limitation is the difficulty of measuring cost and productivity in markets with differentiated products. This is especially challenging in retail markets, where the services provided are not easily measurable, and firms often sell multiple complementary products. The contribution of this paper is to develop a new approach for measuring misallocation from collusion that does not rely on accounting cost/productivity data, while taking into account the reallocation of demand under a non-cooperative pricing equilibrium with spatial differentiation.

¹See also Midrigan and Xu (2014), Restuccia and Rogerson (2017), and Edmond, Midrigan, and Xu (2023) amongst others.

We study a specific source of cost misallocation arising from the ubiquitous use of uniform prices in collusive arrangements. The reliance on focal points to coordinate pricing is well documented in cases of tacit and explicit collusion (Schelling 1960, Knittel and Stango 2003). Because cartel organization and monitoring are costly activities, most known collusive arrangements feature simple pricing rules, such as uniform pricing across locations, to reduce complexity.² However, uniform and focal prices create inefficiencies beyond standard “overcharges” by weakening the link between costs and market shares.

Our approach involves estimating an equilibrium model of demand and supply to recover retailers’ marginal costs from observed prices and quantities. The next step is to quantify the importance of cost misallocation and identify its sources. Following Olley and Pakes (1996), we compare the static allocative efficiency of the market in four environments: (i) uniform-price collusion, (ii) Bertrand-Nash pricing, (iii) perfect competition, and (iv) no demand frictions. We next examine whether low-cost firms/outlets face stronger or weaker incentives to invest, enter, or exit, shedding light on how uniform-price collusion affects the long-run dynamic efficiency of the market.

The empirical focus is on Quebec City’s retail gasoline market, for which we have access to detailed data on prices, market shares, and station characteristics for more than ten years covering the 1990s and 2000s.³ This setting has several features that make it useful for studying the misallocation effects of cartels in retail markets.

First, pricing patterns in Quebec City are consistent with the emergence of collusion. Unlike in the 1990s, when pricing was more competitive and there was significant price dispersion and lower margins, we document the transition to a stable equilibrium with high margins and near-uniform pricing in the aftermath of a market-wide price war that lasted nearly a year. These patterns persisted for several years and followed closely the patterns observed in nearby markets investigated by the Competition Bureau of Canada.

Second, there exists heterogeneity in stations’ operating costs due to stations being part of chains with networks of various sizes, having different vertical contracts with refineries, different sizes, and different amenities. Together, these differences may imply that stations have different opportunity costs of serving customers, which is a necessary condition for production misallocation.

Lastly, the industry underwent a major technological transition during this period, driven

²See, for instance, Clark, Horstmann, and Houde (2024) for an example of a uniform-price retail cartel, and Harrington (2016) for a discussion of the challenges of cartel organization and possible solutions.

³Legal disclaimer: This paper analyses any alleged cartel activity strictly from an economic point of view. To our knowledge, there has been no antitrust investigation into collusive practices in this market.

by the adoption of automatic payment and logistic technologies. As a result, in addition to the static effects just described, uniform collusive prices may also determine which firms operate and invest in the market, affecting the long-run dynamic production efficiency of the industry.

Establishing relevant markets and quantifying market power requires estimating substitution patterns between own and rival stations that capture geographical differentiation. To do so, we consider two possible demand frameworks. The first follows Houde (2012) and takes spatial differentiation seriously by incorporating the commuting patterns of consumers into a mixed-logit framework. The second builds on the flexible parametric demand models developed by Pinkse, Slade, and Brett (2002), Gandhi and Houde (2019), and Monardo (2025). Both demand models capture rich substitution patterns. The results show that in Quebec City’s gasoline market consumers are very price sensitive (own-price elasticity of -9). However, we find that the estimated mixed-logit model implies sizeable i.i.d. shocks in order to rationalize the data, which leads to invalid sales diversion and business stealing across space. Since our goal is to predict reallocation of sales under competitive pricing, we use the flexible parametric demand model as our main specification.

We combine the demand elasticities with pricing first-order conditions to recover marginal cost estimates for each station during a competitive period (late 1990s). With the cost estimates and our model, we perform a set of counterfactuals designed to measure productive inefficiency under collusion. Drawing on Olley and Pakes (1996), we measure static misallocation by analyzing the covariance between stations’ output levels and marginal costs. The sign and size of this covariance indicate the extent of static production misallocation. In a competitive environment, the covariance should be negative, as firms with lower marginal costs tend to have lower prices and higher market shares. In contrast, with collusive uniform prices, consumers will sort based solely on station characteristics and convenience, such that the sales share of high marginal cost stations increase and the covariance shrinks.

Our first finding is that a negative correlation between marginal costs and market shares exists regardless of conduct, which suggests that low-cost stations tend to be closer to consumers and have more attractive characteristics. The results show that the covariance between cost and market shares under the observed collusive allocation is 5.3% of the most cost efficient allocation. This is the extent of misallocation that would be attributed to collusion if using the approach developed in Asker et al. (2019). However, our findings suggest that much of this misallocation is in fact caused by differentiation between stations, rather than uniform price collusion directly. The Bertrand-Nash allocation corresponds to 17.3% of the

efficient allocation, indicating that eliminating collusion would improve allocative efficiency by 12%. The marginal cost pricing allocation corresponds to 20% of the efficient allocation. This is close to the Bertrand-Nash allocation, implying that market power itself leads to relatively small cost misallocation relative to perfect competition. The remaining sources of cost misallocation are the differences in preferences associated with spatial differentiation, as well as differences in the value of amenities across stations (vertical differentiation). If instead of comparing to the cost-minimizing benchmark, we consider efficiency gains from making the market competitive, our results imply that eliminating collusion would get us roughly 80% of the way there, while eliminating market power would do the rest.

These findings have implications for the calculation of antitrust damages in cartel cases. Using our model estimates and counterfactual simulations, we show that the average but-for-collusion price that is the basis of damage calculations can be adjusted to reflect both the direct price effects of collusion and the allocative inefficiencies arising from excessive production by high-cost firms. Our results suggest that ignoring the misallocation effect results in a roughly 3% underestimate of damages. The discrepancy is not large, but it illustrates that even in a market like retail gasoline—where stations may differ but the product is essentially homogeneous—nontrivial misallocation can still arise.

To analyze the dynamic misallocation effects of uniform-price collusion, we use our demand estimates and the empirical distribution of demand and marginal cost shocks to measure the expected variable profit obtained at different marginal cost levels. The ratio of the expected variable profit under uniform-price collusion with the expected variable profit under Bertrand-Nash pricing provides a measure of the expected return of uniform-price collusion. Our results imply a positive correlation between this ratio and marginal costs, indicating that uniform-price collusion favors high-cost inefficient stations as opposed to low-cost efficient stations. Furthermore, we also find that the return to a one-standard deviation decrease in marginal cost is 21% lower under collusion than Bertrand-Nash pricing. Our results imply that uniform-price collusion reduces incentives to enter and invest in markets. Consistent with this, we document a decrease in industry turnover. Following Melitz and Polanec (2015), we decompose the change in covariance between cost and market shares into within-incumbents reallocation, and entry/exit. In the decade before uniform pricing began, roughly a quarter of the observed efficiency gains were due to industry turnover, compared to nearly zero under uniform pricing.

Literature Review

Since at least Harberger (1954), the impact of higher prices resulting from collusion on welfare has been well understood. More recently, research has shown that cartels can harm society by misallocating sales to less efficient players (see Asker, Collard-Wexler, and De Loecker (2019) and Bridgman, Qi, and Schmitz (2015)). Other papers have focused more on the impact of cartels on the entry and exit of inefficient firms. Carrera (2022) studies Chile’s nitrate cartel and finds collusion may provide an umbrella under which inefficient firms can survive in a market. He finds that, under collusion, the rate of entry was higher and that entrants were of lower productivity. Model simulations suggest that had competition prevailed, a quarter of entrants would have remained on the sidelines. Finally, there is a small number of papers that estimate the overall impact of collusion on productivity. See for instance Petit, Kemp, and van Sinderen (2015) and Moreau and Panon (2023).

In addition to these papers, we are related to Asker (2010), who examines bid rigging in stamp auctions. Cartel participants met before auctions to allocate the right to bid and determine side payments. His findings suggest that the bidding ring resulted in limited misallocation and therefore led to only a mild decrease in efficiency. Röller and Steen (2006) study the inner workings of the Norwegian cement cartel and find that its production-sharing rule was inefficient as it provided incentive to overinvest in capacity and to export at a loss in an effort to increase their share of domestic sales.

We are also related to the empirical literature on the organization of cartels, especially those papers that have focused on describing their inner workings such as Genesove and Mullin (2001), Asker (2010), Harrington and Skrzypacz (2011), Clark and Houde (2013), Clark, Coviello, Gauthier, and Shneyerov (2018), Hyytinen, Steen, and Toivanen (2019), Igami and Sugaya (2022), Clark, Horstmann, and Houde (2024), and Starc and Wollmann (2025). Starc and Wollmann (2025) study the US generic drug cartel and argue that the collusive price attracted new entry, although they do not consider which types of firms enter as we do in this paper. They also point out that the cartel remained stable, despite entry, relevant for our dynamic analysis. Clark, Coviello, Gauthier, and Shneyerov (2018) study a procurement cartel to show that entry was deterred under collusion and that following the collapse of the cartel new firms entered and proceeded to win a significant share of auctions, suggesting that they were relatively efficient. Most relevant is Clark and Houde (2013), which makes use of documents submitted by the Competition Bureau to the Superior Court of Quebec containing summaries of and extracts from recorded conversations between participants in a cartel that allegedly operated in the early 2000s in the retail gasoline markets

of four cities in the province of Quebec—Magog, Sherbrooke, Thetford Mines and Victoriaville. The paper describes the communication patterns and pricing strategy employed and explains the mechanism used by the alleged participants in order to solve the enforcement and coordination problems that arise in price-posting markets with heterogeneous firms. The mechanism is based on the ordering of moves and adjustment delays during price changes that allow market share transfers to *stronger* players. The paper estimates the implicit market share gains of the stronger players, and evaluates the importance of the transfers for cartel stability.

In Clark and Houde (2014), a companion to their 2013 paper, the authors use micro-level price data to study the impact on margins of the collapse of the alleged cartels in the same markets following the Competition Bureau investigation. Notable is their result that not only do stations in the four target markets react to the Bureau’s investigation by reducing their margins, but similar patterns can be detected in a number of other markets in the province, including Quebec City.^{4,5} In more recent work, Clark, Duarte, and Houde (2024) study the price war that occurred in Quebec City in 2000.

In addition to the papers mentioned above, several others have examined collusion in retail gasoline markets, including Wang (2008, 2009) and Byrne and deRoos (2019) for Australia, and Chaves and Duarte (2025) for Brazil, and Sorgard and Steen (2024) for Sweden. Assad, Clark, Ershov, and Xu (2024) examine whether algorithmic pricing can facilitate collusion in retail gasoline markets. There are also been papers on competition in retail gasoline markets more generally, including some that have focused on Quebec’s retail gasoline market. See for instance, Houde (2012) and Carranza, Clark, and Houde (2015).

We are also related to a number of papers that quantify the impact of cartels on prices or the damages they impose. In addition to Clark and Houde (2014), see for instance Miller and Weinberg (2017), Clark, Coviello, Gauthier, and Shneyerov (2018), and Kim, Mansley, Miller, Remer, and Weinberg (2023).

Finally, we are related to the set of papers that have documented uniform pricing in retail markets and investigated the distributional impact on consumers. DellaVigna and

⁴See also Erutku and Hildebrand (2010) who estimate the collusion overcharges in the Sherbrooke market. Their results show that following the announcement of the Bureau investigation, prices decreased by nearly 2 cpl relative to Montreal, but not at all relative to Quebec City.

⁵These findings are suggestive, but not conclusive, of the possibility that collusion was widespread in the province. To our knowledge the Competition Bureau investigation was not extended to markets outside of Magog, Sherbrooke, Thetford Mines and Victoriaville, and as a result there exists no direct evidence of explicit coordination among players outside these four markets and other reasons can explain the similarities in pricing. However, we will present evidence of elevated margins and limited price dispersion in Quebec City, patterns that are consistent with either tacit and explicit coordination.

Gentzkow (2019) and Hitsch, Hortacsu, and Lin (2021) study pricing at grocery stores and find that chain-level pricing is largely uniform. Adams and Williams (2019) examine pricing at hardware stores and documents zonal pricing. We add to this literature by establishing a connection between uniform pricing and collusion, and highlighting the production inefficiencies implied by such pricing rules.

Outline

The paper proceeds as follows. In the next section, we describe the data used in this study and characterize Quebec City’s retail gasoline market over time. Section 3 provides a discussion of the impact of collusion on misallocation and the challenges of proper measurement. In Section 4 we present the specification and estimation methodology of the model used to simulate counterfactual prices and sales. We then present and discuss the model estimates in Section 5, and our quantification of static and dynamic misallocation in Section 6.

2 Data and Industry Background

2.1 Data

Our main data source is Kent Marketing, the leading survey company for the Canadian gasoline market.⁶ Kent surveys price and sales information for the universe of stations in many large Canadian markets (including Québec City) for one day every two months (less often in smaller cities). The survey contains accurate measures of the volumes sold for each fuel grade since the last date of read, along with information on station characteristics and addresses. It also provides a snapshot of the prices for each grade collected on the date of the survey. We have this information from August 1992 to June 2001, as well as isolated periods during 2004 and 2007 (and for certain periods for a handful of other markets). The available time range allows us to capture periods of both elevated price dispersion and near-uniform prices, as we explain in further detail below.

We complement the Kent data with three other data sets. First, we have access to weekly retail price information from the Régie de l’énergie du Québec (Québec’s energy board), which surveys a sample of stations in a number of cities across the province in an effort to monitor retail gasoline prices. We observe station-level prices from December 1997 until February 2004 for around 13 stations in Québec City (and in some other markets in

⁶Now Kalibrate Canada Inc.

the province too). The Régie also provides information on the weekly price floor for each market, which we use as a first proxy for costs (further information on the floor is provided below). Second, since the Régie price survey only runs until early 2004, we complement this information with similar survey data from MJ Ervin, a market research company.⁷ The MJ Ervin data cover the period from January 2001 to November 2012 for 15 stations in Québec City. We also have daily price and sales information from the largest independent retail chain in Québec City, Eko, for the period covering August to December 2000. These data are at the station level and allow us to precisely measure the impact of pricing on sales. There are 29 active Eko stations during this period. Finally, we use data on the ownership of underground-storage tanks to measure the vertical integration of stations. This information is provided by the Québec Ministry of Natural Resources.⁸

To estimate our model of station choice, we also make use of data on the population residency distribution and commuting patterns along with the distance / driving time for each commute route. We describe these data in more detail when we discuss the model.

2.2 Québec City’s Retail Gas Market

There are three different types of stations operating in the Québec City market: refinery-owned, and both branded and unbranded independents. Independent branded stations operate under two different types of contracts – commission-based and cross-lease. Both types of vertical contracts are negotiated every three to five years, reflecting local market conditions. During our sample period, commission-based contracts were used mainly by Ultramar, the largest chain of branded retailers, which owned the nearest refinery at the time. The contracts involved Ultramar controlling the inventory and setting the price, with the station owner receiving a commission. Rather than using this arrangement, most independent branded retailers signed long-term cross-lease contracts with a refiner. The structure of these contracts differs between stations and brands, but the idea is that the independent retailer leases the facility to a refiner, and the refiner leases it back to the retailer under various terms and conditions.⁹ Stations operating under cross-lease contracts are responsible for setting the retail price but face a non-linear compensation scheme.¹⁰

⁷MJ Ervin performed the survey at the time, but was subsequently acquired by Kent Marketing, which in turn was acquired by Kalibrate Canada Inc.

⁸Houde (2012) and Clark and Houde (2013) also use these data.


⁹The main reason for the existence of these contracts is that they allow the upstream supplier to maintain access to the site in the event that the station owner sells the property.

¹⁰These agreements specify what equipment must be installed at the facility (e.g. pumps and convenience store), and determine the structure of the retailer’s compensation. The cross-lease contract specifies a “lease”

The focus of our study will be on the period of near-uniform pricing from 2001 until at least 2006, but to understand the structure of the market during this period, it is necessary to consider its evolution starting in the early 1990s. As in most North American retail gasoline markets, Quebec City experienced an important reorganization throughout the 1990s. This transformation was a response to technological innovations such as automated service and better inventory control systems, and was characterized by an increase in the size and automation of stations that resulted from the entry and reconfiguration of larger and more efficient stations, and the exit of smaller stations.

As a result of this transition, Québec City’s retail gasoline market became steadily more concentrated over the years. In Figure 1 we present the evolution of the total number of stations and the network sizes of the five largest brands (hereafter major brands). In the first half of the 1990s, the retail market was relatively unconcentrated, with as many as 360 stations. The network of stations affiliated with a major distribution brand accounted for 64% of stations and 70% of fuel sales in 1993. However, by 2004 the total number of stations had already fallen by 22%, while the networks of the major chains expanded to 68% of the market and their share of sales increased to 79%.

Figure 1: Evolution of Number of Stations & Major Brands Network



Notes: Evolution of the total number of stations (solid line) and the network size of the five largest brands (dashed line). Major brands are Ultramar, Irving, Shell, PetroCanada, and Esso.

The transition to larger and more efficient stations greatly squeezed the margins of small independent stations, and encouraged these retailers to lobby for protection from the provin-

paid to the station operator, which is expressed on a per-volume basis (e.g. 1.6 cpl), but not the wholesale price since it is not contractible due to fluctuations in oil prices.

cial government. In response, the government enacted a floor price, set weekly, for all gasoline retailers in Quebec.¹¹ Despite the existence of the price floor a significant price war occurred in 2000, with markups at many stations falling to zero for the better part of the year. Clark, Duarte, and Houde (2024) provide an extensive discussion of the 2000 price war. For our purposes, the relevant information is that the price war was triggered by the expiration of vertical contracts and not by specific demand-shocks faced by stations.

Table 1 presents summary statistics for size, amenities, prices, and sales volume for three different periods of our sample: (i) prior to the start of the transition (1993), (ii) post-transition but before the price war (1999), and (iii) near uniform pricing (2004). The table highlights a number of key features of the market and its evolution. First, the table reveals notable asymmetries between stations. This can be seen, for instance, from the standard deviation in the number of pumps and islands. Variation in the size of stations is also reflected in the significant heterogeneity in daily volume sold. We can also see that over time, the average station gradually became larger and invested more in convenience stores rather than repair shops.

There is substantial heterogeneity in station density across different parts of Quebec City. As shown in **Table 1**, in 1999 a comparable number of stations faced more than ten nearby competitors as faced fewer than three. The map in Figure 8 of Appendix A illustrates this spatial pattern: stations with many competitors tend to be located along major commuting corridors and at the intersections of primary thoroughfares, whereas those with fewer competitors are situated in residential or mixed-use commercial areas, farther from the main highway network. Local density also evolved considerably over time. In 1993, 42% (148/351) of stations encountered more than ten competitors within a three-minute driving radius, while 13% (46/351) faced fewer than three; by 2004, these figures shifted to 19% (51/272) and 18% (50/272), respectively.

We hypothesize that differences in local station density generate meaningful heterogeneity in consumers' price responsiveness. To explore this, we draw on daily pricing data from a

¹¹The price floor is determined according to the following approximation for the average marginal cost of selling gasoline in each local market m at time t :

$$MEP_{mt} = w_t + \tau_{mt} + T_{mt}, \quad (1)$$

where w_t is the minimum wholesale price at the terminal, τ_{mt} is an estimate of the transportation cost to deliver gasoline from the refinery to market m , and T_{mt} is the sum of federal and provincial taxes. Since most stations (especially branded stations) pay a wholesale price below w_t through negotiated discounts, the MEP is an upper bound on the marginal cost of selling gasoline. The price floor is determined and posted on the Régie's website every Monday.

Table 1: Summary Stats

	Fall-1993		Fall-1999		Fall-2004	
<i>Market outcomes</i>						
Price (cents/liter)	66.18	(0.67)	66.04	(0.81)	83.96	(0.58)
Price range (90th-10th percentile)	1.33	-	1.03	-	0.18	-
Volume (ths liters/day)	3.59	(2.21)	5.19	(3.33)	6.13	(3.57)
<i>Station characteristics</i>						
Number of pumps	5.87	(4.09)	7.22	(4.96)	8.10	(5.44)
Number of islands	2.20	(1.34)	2.38	(1.40)	2.51	(1.41)
Large convenience	0.25	(0.43)	0.31	(0.46)	0.41	(0.49)
Repair shop	0.31	(0.46)	0.24	(0.43)	0.21	(0.41)
Carwash	0.20	(0.40)	0.19	(0.39)	0.19	(0.39)
Self Service	0.43	(0.50)	0.58	(0.49)	0.73	(0.45)
<i>Market characteristics</i>						
Major brands's sales share	0.70	-	0.76	-	0.79	-
Number of stations	351	-	296	-	272	-
Number of stations with > 10 local comp	148	-	69	-	51	-
Number of stations with < 3 local comp	46	-	51	-	50	-

Notes: We show the average value for the corresponding period. Standard deviation is shown in parentheses when appropriate. Major brands are Ultramar, Irving, Shell, PetroCanada, and Esso. Prices are in 2001-01 values. A local competitor is a station within 3 minutes' driving distance. The Fall refers to the fourth or fifth bimonth of the year.

single chain (Eko) covering a three-month period during the 2000 price war. Table 2 reports the correlation between log prices and log sales at Eko stations as a function of the number of rival stations within a two-minute driving distance, controlling for station fixed effects. Even within the same brand, this correlation differs by more than 2 percentage points depending on whether a station faces no nearby competitor or two nearby competitors. Capturing these differences in consumer responsiveness is crucial for mapping observed price variation into implied cost differences, and therefore will inform our choice of demand model.

2.3 The Evolution of Conduct

Table 1 also highlights the changes over time in the degree of price dispersion in the market. In the 1990s prices were relatively dispersed: the difference between the 90th and 10th percentile price was over 1 cpl in both 1993 and 1999. In contrast, in 2004, the range fell to just 0.18 cpl. Note that this decrease is despite the increase in dispersion of station characteristics.

This can also be seen in Figure 2, which characterizes the evolution of price dispersion and the median margin (relative to the local price floor) in Quebec City. In the early 1990s, price

Table 2: Price elasticity per number of local competitors
– Eko stations

	log(Sales)
	1
log(Price) of closest Regie station	0.538 (0.821)
log(Price)x(N. stations within 2min=0)	-5.717 (0.734)
log(Price)x(N. stations within 2min=1)	-7.029 (0.824)
log(Price)x(N. stations within 2min=2)	-7.836 (0.976)
Data Source	Eko
Time FE	Week+Day of Week
Station FE	Yes
Adj. R ²	0.292
Num. obs.	1442


Notes: Regie station refers to stations that are survey by Régie de l'énergie du Québec. N. stations within 2 min refers to the number of stations that are less than 2 min driving distance apart. There are 3, 8 and 5 stations with 0, 1 and 2 competitors in a 2 min driving distance, respectively.

dispersion and margins were high before falling in the mid-1990s as the market transitioned. The 2000 price war saw price dispersion and margins shrink dramatically, but by 2001 margins had already recovered to a level above 5 cents per liter, and retailers sustained margins at least this high until 2008. At the same time, price dispersion in Quebec City diverged significantly from what was observed during the 1990s. The within-week standard deviation fell sharply after 2001 and by 2003 reached a point where price surveys could not capture any differences in prices throughout the market.¹²

To demonstrate that the pricing changes observed after 2000 result from a shift in firm conduct, in Figure 2 we contrast the evolution of margins and price dispersion in Quebec City with the evolution of these variables in two other markets in the province, Sherbrooke and Montreal. These two cities provide useful benchmarks when studying margins and price dispersion in Quebec City. First, even taking into account level differences stemming from supply chain heterogeneity, the costs in these three cities fluctuate in a similar fashion. Second, as mentioned, Sherbrooke was one of the cities targeted by the Competition Bureau of Canada in its investigation into price fixing, while there is no evidence that Montreal was ever suspected of any wrongdoing. We can see from the figure that before 2000 margins

¹²One might be concerned that this is driven by the fact that the Régie and MJ Ervin only surveyed a small number of stations (around 10). But the drastic reduction in price dispersion is also observed in the single-day market snapshots from the more comprehensive Kent dataset.

Figure 2: Evolution of Price Margin and Price Dispersion across Markets



(a) Weekly Median Margin (cpl)

(b) Weekly Price Standard Deviation (cpl)

Source: Régie de l'Énergie du Québec (1998-2001), and MJ Ervin (2001-2012). We aggregate weekly statistics using the median value in a year. The average number of stations survey in a week is 7, 10 and 5 for Montreal, Quebec City and Sherbrooke, respectively. Margins are measured relative to the local price floor.

and price dispersion in Quebec City were similar to what is observed in Montreal. However, starting in 2001 there was an important divergence, with margins in Quebec City much higher than those in Montreal, and price dispersion significantly lower. In contrast, during this later period, the pricing pattern observed in Quebec City closely resembles that observed in Sherbrooke.¹³

As further evidence that prices may have been coordinated in Quebec City we consider the findings presented in Clark and Houde (2014) regarding pricing patterns in the four markets targeted by the Competition Bureau investigation as compared to other markets in the province. In particular, they noted that several markets in the eastern part of Quebec, including Québec City, had pricing patterns similar to those in the four cities targeted by the Bureau and in stark contrast to the pattern observed in the western part of the province. Specifically, they document that the markets in the target cities had high margins and asymmetric pricing adjustments (i.e., price changes $\neq$ cost changes). That being said, as mentioned in Footnote 1, there have been no allegations leveled by the Competition Bureau of Canada of collusion in Quebec City's retail gasoline market. For our purposes, what matters is that prices are nearly uniform during the early 2000s.

¹³The convergence in margins between Montreal and Québec starting in 2007 is consistent with the findings in Clark and Houde (2014) and Erutku and Hildebrand (2010) regarding the impact of the collapse of the alleged cartel in nearby cities following the Competition Bureau's investigation.

3 Uniform Collusive Price and Misallocation

In addition to the potential damages caused by higher prices, the collusive equilibrium with uniform prices that started in 2001 could also have distorted the sales distribution between low- and high-cost stations. In this section, we point out how the uniform collusive price can lead to static and dynamic productive inefficiencies, and we explain our approach to measure these inefficiencies.

3.1 Static and Dynamic Misallocation

Under competition, market share is allocated as a function of firms' efficiencies, and persistent differences in production costs should drive inefficient firms out of the market. In contrast, in situations of elevated market power, such as collusive arrangements, there is a potential for inefficient firms to survive and earn market shares at the expense of their more efficient rivals. In static terms, this market reallocation increases the aggregate cost of supplying goods during any period the collusive arrangement is in place and creates the productive inefficiency discussed in Asker, Collard-Wexler, and De Loecker (2019). In dynamic terms, by distorting firms' incentives to enter, exit and invest, a collusive arrangement can have a lasting effect on prices and production efficiency.

Because of the challenges of implementing and sustaining collusive arrangements in retail markets, participants often use simple pricing rules such as uniform (or nearly uniform) prices to facilitate organization, which may amplify production misallocation. Uniform collusive pricing can yield sales allocations that are different from those that would be observed if prices reflected the underlying cost and demand-side characteristics of firms operating in the market. In retail, there may be important cost asymmetries between firms based on features such as store characteristics, management structure, and network size. Under competition, larger low-cost players will post prices inferior to those of their competitors and, as a result, sell more. In contrast, under collusive uniform prices consumers can opt to shop at more convenient outlets with more appealing characteristics, e.g. closer to their homes. Therefore, the sales distribution resulting from uniform prices can lead to greater misallocation depending on the level of product differentiation.

Uniform collusive prices can also affect the set of firms operating in the market, generating a dynamic production inefficiency that may have a lasting effect on market prices even after the collusive arrangement collapses. There are three main economic forces leading to dynamic effects. First, the collusive arrangement increases profits flowing to high-cost firms,

which allow them to continue operating instead of exiting the market. Second, by increasing overall industry profit, collusion facilitates the entry of new firms. However, if more efficient firms have higher entry/fixed costs, and if they do not believe they will be able to disrupt the cartel by undercutting prices and gaining market share, then a uniform collusive price distorts entry incentives favoring entry of relatively inefficient firms. Lastly, uniform collusive pricing weakens the incentives of incumbent firms' to lower operational costs. Under uniform pricing, firms cannot lower their price to increase market share following a cost reduction, so the benefits flowing to firms from cost decreases are limited relative to competitive settings, reducing the incentive to invest.

3.2 Methods for Measuring Misallocation

The current approaches that have been developed for measuring production inefficiencies due to misallocation abstract from important features that are present in differentiated-products industries such as retail. Hsieh and Klenow (2009) and much of the literature that follows have used revenue-based estimates of total factor productivity (TFPR) dispersion as the relevant statistic to measure production misallocation. In this framework, the extent of misallocation is obtained by comparing the economy of interest with a benchmark economy, which is assumed to be undistorted. Despite providing a reasonable approximation for many settings, this approach is not well suited for retail industries. Product differentiation is modeled using an industry CES aggregator. Doing so restricts the substitution patterns implied by the model and limits the role of demand frictions, such as spatial and brand differentiation, in determining firms' market power. Furthermore, the estimation of TFPR often assumes single-product firms, which abstracts from the fact that pricing power varies across retailers depending on their store network size.

Asker, Collard-Wexler, and De Loecker (2019) propose an alternative approach for measuring misallocation that focuses on dispersion in marginal costs. In their setting – the global oil extraction industry – marginal cost of production can be obtained by combining accounting data for each oil field with assumptions about the production function. The authors then measure misallocation from collusion by comparing the total cost implied by the realized oil extraction under the OPEC cartel with the total cost implied by a counterfactual efficient benchmark extraction level that would be realized if the distribution of production distribution were to perfectly follow the firm (country) cost ranking up to capacity constraints. Despite providing a creative alternative for measuring misallocation that does not depend on TPFR, this approach has two limitations that make it inadequate for application to re-

tail settings. First, good measures of production costs are typically not available, and even cost proxies, such as the wholesale price of the goods being studied, often do not capture retailers’ opportunity costs of selling the good. Second, the efficient benchmark in Asker, Collard-Wexler, and De Loecker (2019) abstracts from demand-side frictions, such as horizontal differentiation and shipping costs. As we show in our empirical exercise, accounting for these frictions can be important for obtaining the competitive counterfactual benchmark.

We depart from Asker, Collard-Wexler, and De Loecker (2019) and the earlier literature by developing a demand-based approach for measuring production misallocation that can be applied to industries with multi-product firms operating in differentiated goods markets. Applications include, but are not limited to, retail markets where product and spatial differentiation are key features. We estimate outlet-level elasticities from a demand model for the industry and infer marginal cost through a conduct assumption and equilibrium pricing conditions, focusing on a more competitive period. Importantly, the quality of the marginal cost estimates depends on how well our demand model approximates the actual substitution patterns of consumers in the market. One problem with the typical logit-based approaches employed in the literature to estimate demand is that elasticities are, to varying degrees, proportional to market share. In our application we address this problem by considering (i) the mixed logit specification developed in Houde (2012) that allows elasticities to also depend on traffic flows, and (ii) the flexible parametric demand models of Pinkse, Slade, and Brett (2002), Gandhi and Houde (2019), and Monardo (2025). We discuss these approaches in more detail in the next section.

With reasonable marginal cost estimates in hand, the basis of the misallocation analysis is the Olley and Pakes (1996) decomposition of the sales-weighted aggregate marginal cost:

$$c_t = \bar{c}_t + \sum_{j=1}^{J_t} (c_{jt} - \bar{c}_t)(s_{jt} - 1/J_t), \quad (2)$$

where $\bar{c}_t$ is the unweighted average of firm-specific marginal cost and $\sum_{j=1}^{J_t} (c_{jt} - \bar{c}_t)(s_{jt} - 1/J_t)$ is the covariance between market shares and marginal costs. The first component provides a summary statistic on the industry technology. The second provides a summary statistic on the co-movement between costs and market shares, reflecting the allocation of sales among plants with varying degrees of production efficiency.

To assess the degree of misallocation and its underlying causes we can compare the covariance terms implied by different counterfactual pricing scenarios. To evaluate misallocation under uniform pricing, we compare the observed misallocation measure with the misalloca-

tion generated by a competitive scenario where prices partially reflect the underlying cost differences. Finally, by applying the Olley and Pakes (1996) decomposition to different periods of our sample we are able to unpack how investment, entry, and exit affect productive efficiency.

4 Empirical Model of Demand and Supply

In this section, we describe our approach for estimating an equilibrium model of demand and supply from observed prices and quantities, and for recovering marginal cost estimates that will later be used in the decomposition defined in [Equation 2](#).

Demand

We model gas station demand as a discrete choice problem. For each bimonth t , consumer i chooses between purchasing gasoline from one of the J stations or the outside option (e.g. public transit, bicycle, etc). For station j , the aggregation of individual choices is represented by a hedonic demand equation,

$$s_{jt} = \sigma_j(\mathbf{X}_t, \mathbf{L}_t, \mathbf{p}_t, \xi_t),$$

where in our context $\mathbf{L}_t$ is the matrix of station locations (in other contexts it could denote the matrix of locations of product characteristics in characteristics space), $\mathbf{X}_t$ is a matrix of station characteristics, $\mathbf{p}_t$ is the vector of posted prices, and ξ_t is a vector of unobserved (to the econometrician) station characteristics. The set of observed station characteristics includes the number of pumps and the number of islands, a dummy variable indicating the presence of a repair shop, a dummy variable indicating the presence of a carwash, a dummy variable indicating the presence of a convenience store, a dummy variable indicating if the station is self-service, and brand fixed effects.

Location plays an important role in retail gasoline markets, and we assume that gasoline stations are horizontally differentiated and that consumers have homogeneous preferences for station characteristics (including prices):

$$\delta_{jt} = X_{jt}\beta + \alpha p_{jt} + \xi_{jt}.$$

The demand equation can thus be expressed as

$$s_{jt} = \sigma_j(\boldsymbol{\delta}_t, \mathbf{L}_t).$$

Assuming that consumer preferences satisfy the invertibility conditions in Berry, Gandhi, and Haile (2013), this demand system admits the following inverse representation:

$$\sigma_j^{-1}(\mathbf{s}_t, \mathbf{L}_t) = \delta_{jt} = X_{jt}\beta + \alpha p_{jt} + \xi_{jt}. \quad (3)$$

Next, we discuss different specifications of the inverse demand function and highlight the extent to which each captures spatial differentiation.

Logit: The first specification we consider is the standard logit model. In this model, there is a simple analytic expression for the inverse demand given by

$$\sigma_j^{-1}(\mathbf{s}_t, \mathbf{L}_t) = \log(s_{jt}/s_{0t}),$$

where s_0 is the market share of the outside option. Note that the logit model abstracts from spatial differentiation by ruling out the possibility of stations' locations influencing the inverse demand. Furthermore, this model is well known for generating substitution patterns that are proportional to market shares. Nonetheless, we estimate this model since it is a useful benchmark in a differentiated goods setting.

Mixed Logit: We follow Houde (2012) and model spatial differentiation in a mixed logit framework. In this model, we define a distance function that captures the extra distance consumers incur when they deviate from their commuting path to shop at a particular station. Specifically, every consumer i can be fully characterized by an origin-destination pair (o_i, d_i) . Let $r(\cdot, \cdot)$ be the function that determines the shortest driving time between two points. We define the distance from consumer i to gas station j as the additional driving time incurred by i from its commuting path when driving to j , denoted by $D_i^{ml}(L_j)$. Then our measure of distance between consumer i and station j is given by

$$D_i^{ml}(L_j) = r(o_i, L_j) + r(L_j, d_i) - r(o_i, d_i).$$

The conditional indirect utility for consumer i of choosing alternative j during time period

t is:

$$u_{jit} = \begin{cases} \delta_{jt} + \lambda D_i^{ml}(L_j) + \epsilon_{ijt} & j = 1, \dots, J \\ \epsilon_{i0t} & j = 0, \end{cases}$$

where λ measures the travel cost incurred by consumers and ϵ_{ijt} is an i.i.d Type 1 Extreme Value shock. The demand system takes the form

$$\sigma_j(\delta_t, \mathbf{L}_t) = \sum_{i \in \mathcal{I}} R_i \frac{\exp(\delta_{jt} + \lambda D_i^{ml}(L_j))}{1 + \sum_k \exp(\delta_{kt} + \lambda D_i^{ml}(L_k))}, \quad (4)$$

where $\mathcal{I}$ is the set of all origin-destination pairs and R_i is the relative mass of consumers located on a given origin-destination path. The distribution of consumers across commute paths is computed using census track population data from the census and traffic flow information from the origin/destination survey by the Quebec government for flow between 950 geographical points across Quebec City Census Metropolitan area.

There is no analytical expression for the inverse demand in our mixed logit framework. Instead, we obtain [Equation 3](#) numerically by applying the fixed point algorithm proposed in Berry, Levinsohn, and Pakes (1995) to invert the demand system defined by [Equation 4](#).

Flexible Parametric Demand: As discussed in Gandhi and Houde (2019) and Monardo (2021), the permutation invariance property of discrete-choice models of demand with exchangeable random-utility shocks allows us to write the demand function as follows

$$\sigma_j^{-1}(\mathbf{s}_t, \mathbf{L}_t) = G(s_{jt}, \{s_{kt}, D_{jk}^{fpd}(\mathbf{L}_t)\}_{k \neq j}) + \kappa_t, \quad (5)$$

where G is a function of the own share s_j , the set of competitors' shares $\{s_k\}_{k \neq j}$, and D_{jk}^{fpd} refers to the geographic distance between stations j and k . To ease notation, we drop the reference to $\mathbf{L}_t$ from now on.

In order for G to be consistent with consumer theory and the invariance conditions as defined in Monardo (2021), it must satisfy three requirements. First, it must be permutation and translation invariant, so that the labeling or ordering of stations does not affect the outcome. Second, it must be log-homogeneous, which ensures that proportional changes in market shares scale the function in a consistent way. Third, the derivative matrix of G must be symmetric and positive semi-definite.

Log-homogeneity implies that G can be expressed as a sum of the log of the own share and a residual component that depends on the relative shares of competitors and the spatial

distances. Formally, this can be written as

$$G(s_{jt}; \{s_{kt}, D_{jk}^{fpd}\}_{k \neq j}) = \ln s_{jt} + \tilde{G} \left(\{s_{kt}/s_{jt}, D_{jk}^{fpd}\}_{k \neq j} \right). \quad (6)$$

To approximate the term $\tilde{G}$, we adopt a body-order expansion that is truncated at the first order. This means that we only keep the additive contributions of competitors individually, rather than modeling higher-order interactions across multiple stations. The resulting expression is

$$\tilde{G} \left(\{s_{kt}/s_{jt}, D_{jk}^{fpd}\}_{k \neq j} \right) = V_0 + \sum_{k \notin \{j, 0\}} V_1(s_{kt}/s_{jt}, D_{jk}^{fpd}), \quad j \neq 0, \quad (7)$$

where V_0 is a constant term and V_1 is a function describing how each competitor k influences the inverse demand of station j .

Following Pinkse, Slade, and Brett (2002) and Pinkse and Slade (2004), we choose a simple and empirically tractable functional form for the first-order term V_1 :

$$V_1(s_{kt}/s_{jt}, D_{jk}^{fpd}) = \sum_{h \in \mathcal{H}} \gamma^h w^h(D_{jk}^{fpd}) \ln \left(\frac{1}{1 + s_{kt}/s_{jt}} \right). \quad (8)$$

In our approximation, the effect of competitor k depends both on the interaction between its relative share s_k/s_j and the distance. Note that the symmetry condition of the G function is imposed by the functional form. The functions $w^h(D_{jk}^{fpd})$ are weights that determine how competition decays with distance. Intuitively, nearby competitors are given greater weight than more distant ones, reflecting the fact that consumers are more likely to substitute across geographically close stations. The index $h \in \mathcal{H}$ allows for multiple weighting functions, providing flexibility in how distance shapes substitution patterns and the coefficients γ^h govern the relative importance of these different weighting functions. In our application, the set of weights is a partition of $\mathbb{R}_+$, with driving time distance thresholds of 1, 3, and 5 minutes.

To complete the specification, we impose a normalization that sets

$$\tilde{G} \left(\{s_{kt}/s_{0t}, D_{0k}^{fpd}\}_{k \neq 0} \right) = 0 \quad \text{and} \quad \delta_0 = 0,$$

which implies that the constant term satisfies $\kappa = -\ln s_0$. Bringing together these elements,

we obtain the following expression for the inverse demand function:

$$\ln s_{jt}/s_{0t} = X_{jt}\beta + \alpha p_{jt} + \xi_{jt} - \sum_{h \in \mathcal{H}} \sum_{k \notin \{j, 0\}} \gamma^h w_{jkt}^h \ln(1/(1 + s_{kt}/s_{jt})). \quad (9)$$

This equation highlights how the inverse demand depends not only on the own and outside shares, but also on the competitive pressure exerted by nearby stations, with the magnitude of this pressure governed by the distance-based weights.

Discussion: The mixed logit model approximates any random utility model (McFadden and Train 2000) and provides a framework that facilitates the incorporation of consumer taste heterogeneity, theoretically allowing for rich substitution patterns. However, as discussed in Lu and Saito (2022), this approximation result depends critically on the scale of the i.i.d. idiosyncratic random-utility shock relative to the mixing distribution. In empirical applications, a non-negligible i.i.d. shock can lead to undesirable predictions for sales diversion and business stealing. First, despite the inclusion of random coefficients (e.g., on travel cost), substitution patterns may remain effectively global (proportional to market shares) rather than local (substitution toward products with similar characteristics), resulting in unreliable markup and marginal cost estimates. Second, while the i.i.d. component facilitates the invertibility of the demand system, it conflates true taste heterogeneity with model misspecification. Consequently, relying on the i.i.d. shock can bias the predicted re-allocation of market shares when simulating counterfactual non-cooperative equilibria.

By directly approximating the inverse demand, the flexible parametric demand model allows us to remain agnostic about the true heterogeneity distribution, and explicitly incorporate the role of distance in a parsimonious parametric specification. However, the functional form used for the approximation can introduce other misspecification biases. Ultimately, the scope of the misspecification is an empirical question, and we estimate all three models. Our model selection procedure relies on three features: (i) analysis of substitution patterns; (ii) pricing responses implied by each model; and (iii) out-of-sample predictive power.

Supply

Consider the problem of a firm f that owns a subset of the J gas stations ($\mathcal{J}_f$) operating in Quebec City in a given period t . Conditional on a vector of prices $\mathbf{p}_{-ft}$ posted by the competing firms, the station chain f sets prices $\mathbf{p}_{ft}$ to maximize profits:

$$\max_{\mathbf{p}_{ft}} \sum_{j \in \mathcal{J}_f} (p_{jt} - c_{jt}) Q_{jt}(\mathbf{p}_{ft}, \mathbf{p}_{-ft}), \quad (10)$$

where $Q_{jt}(\mathbf{p}_{ft}, \mathbf{p}_{-ft}) \equiv s_j(\boldsymbol{\delta}_t, \mathbf{L}_t)$ denotes the demand function for station j . The system of J first-order conditions obtained when simultaneously solving the profit maximization problem for all firms characterize the equilibrium of the pricing game. Expressing the system of first-order conditions in matrix notation, we obtain:

$$\mathbf{q}_t(\mathbf{p}_t) + (\boldsymbol{\Omega}_t * \boldsymbol{\Delta}_t)(\mathbf{p}_t - \mathbf{c}_t) = 0, \quad (11)$$

where $\boldsymbol{\Omega}_t$ is a $J \times J$ matrix capturing the ownership structure in the market ($\Omega_{jkt} = 1$ if $j, k \in \mathcal{J}_f$ and 0 otherwise), $*$ denotes the element-wise product, and $\boldsymbol{\Delta}_t$ is the Jacobian matrix of the demand system ($\Delta_{jkt} \equiv \partial q_{kt}(\mathbf{p}_t) / \partial p_{jt}$).

We recover the marginal cost for each gas station j by inverting the system of equations in (11). Cost estimates are a function of market shares, the Δ_t implied by the demand model, and any assumption about the ownership structure Ω_t . Assuming that a Bertrand-Nash pricing game reasonably approximates the competitive environment observed during the period before the 2000 price war, the observed price and sales data are sufficient to construct marginal cost measures for each station-period in our sample. In our empirical results, we recover costs only for the bi-month periods of 1999, as we are more confident about the competitive conditions.¹⁴

In order to compute misallocation measures for time periods with uniform pricing collusion, a model is required to predict marginal costs. We project the recovered marginal cost onto observed cost shifters and make distributional assumptions about unobserved cost components. Specifically, we use a two-level mixed-effects specification:

$$c_{jt} = W_{jt}\phi + \psi_t + u_{g(j)} + \omega_{jt}, \quad (12)$$

where the set of observable cost shifters W_{jt} contains proxies for the level of retail service costs (indicators for convenience store, carwash, repair shop, and self-service) and proxies for the level of wholesale price discounts (number of pumps and type of vertical contract). The term ψ_t is a time fixed effect that captures common cost shifters such as changes in the international oil price. $u_{g(j)} \sim N(0, \sigma_g)$ are random effects that capture unobserved shocks

¹⁴Our assumption relies on the fact that we observe price dispersion, and the margins measured as the price difference from the price floor are relatively small.

that are common for a subset of stations according to a partition g of the station set, such as shocks in the distribution cost for a given zip code or shocks at the refinery level. Finally, $\omega_{jt} \sim N(0, \sigma_\omega)$ is a mean-zero cost shock that is idiosyncratic to each station-time period.

Identification

In this section, we discuss the assumptions necessary to identify the parameters of the demand models - α , λ and γ^h - that are key to determining own- and cross-price elasticities. In our empirical exercise, we decompose ξ_{jt} into three components: $\xi_{jt} \equiv \xi_j + \xi_t + \Delta\xi_{jt}$. The component ξ_j captures stations' unobserved characteristics fixed over time, such as the side of the street, size of the parking lot, and road congestion. ξ_t captures time-varying factors common across all stations, such as the price of public transportation, weather, and economic factors. Lastly, $\Delta\xi_{jt}$ is the structural error term that captures deviations from ξ_j and ξ_t , such as switches in the direction of traffic due to road construction.

Let $\mathbf{Z}$ be a set of instruments formed from a given transformation of the observed station attributes $(\mathbf{X}, \mathbf{L}, \mathbf{W})$. The key identification assumption is that, conditional on location fixed effects, the transitory unobserved demand shocks of each station, $\Delta\xi_{jt}$, evolve independently of its own attributes and the attributes of competitors:

$$E[\Delta\xi_{jt}|\mathbf{Z}] = 0,$$

In other words, the validity of the conditional moment restriction relies on the assumption that transitory demand shocks are uncorrelated with market structure. This assumption is valid if, for example, stations make entry/exit and remodeling decisions based on the location's time-invariant characteristics. Given the magnitude of the entry costs, we believe that this is a good approximation.

In an oligopolistic setting, equilibrium prices and quantities depend on the unobserved demand shocks. This dependence leads to an endogeneity problem when estimating the price coefficient of demand, α . To overcome this endogeneity issue, we draw from retrospective merger analysis literature (Hastings 2004) and use the Sunoco-Ultramar merger in 1996 to construct instruments designed to capture exogenous price variation.¹⁵ Specifically, we build extra driving-time rings of 90 seconds around Ultramar/Sunoco stations and interact a dummy variable that is equal to one if the station is within the 90-second ring with a dummy that is equal to one for the post-merger period. The merger affects the price of the

¹⁵For more details on the merger see Houde (2012).

stations owned by the merging firms because it allows them to internalize cannibalization effects, and affects the price set by non-merging nearby firms through a competitive effect. Because the merger involved every Ultramar and Sunoco station in the province of Quebec, the change in ownership is unlikely to be correlated with the demand shocks of particular stations in Quebec City once we condition on the location fixed effects.

The other relevant parameters are the ones capturing the degree of spatial differentiation - λ in the mixed-logit model and γ^h 's in the FPD model. Heuristically, since we include location fixed effects, the covariance between changes in market shares and changes in measures of local competition should be informative about these parameters. For instance, if the entry of a station has a strong negative impact on the market share of only nearby competitors, then the travel cost parameter (λ) should be large, and only the γ^h parameters for small distances should be relevant. In contrast, if the entry of a station has a strong negative impact on the market share of both nearby and far-away competitors, then the travel cost parameter should be small and the γ^h parameters of nearby and far-away stations will be relevant.

Building on this intuition, we construct two sets of instruments that measure the extent of local competition faced by each station j . The first set of instruments is constructed by counting the number of stations within a given driving distance bandwidth that have the same status as j (self-service, full-service):

$$Z_j^a = \sum_{k \neq j} \mathbf{1}(D_{jk} \in b) \times \mathbf{1}(x_j = x_k),$$

where b indexes the driving distance bandwidth ($[0, 250m]$, $(250m, 500m]$, $(500m, 750m]$), and D_{jk} is the driving distance in meters between stations k and j . The variation in this set of instruments is driven by the entry/exit of stations, which we demonstrated in Figure 1 was meaningful during our sample period.¹⁶

The second set of instruments captures the degree of local competition faced by station j during the 2000 price war. Specifically, for each station j during the year 2000, we count the number of stations within a driving distance bandwidth that charge a price lower than does j :

$$Z_j^b = \mathbf{1}(t = 2000) \sum_{k \neq j} \mathbf{1}(D_{jk} \in b) \times \mathbf{1}(p_j - p_k > 0),$$

¹⁶To account for the specific functional form chosen, in the FPD estimation we also include as instruments the number of stations within 1, 1 to 3, and more than 3 driving minutes.

where b and d_{jk} are defined as in the first set of instruments. The validity of the second set of instruments relies on two assumptions: (i) that the price war observed in 2000 is independent of the transitory demand shocks stations face, and (ii) that during the war, prices were only a function of supply-side factors. These assumptions seem reasonable in our context. For instance, Clark, Duarte, and Houde (2024) argue that the war was caused by the expiration of the exclusive dealing contracts between stations and refineries, which were signed years in advance. Furthermore, stations often set prices below the MEP during the war, suggesting that their costs and not demand were the relevant factor affecting pricing decisions.

5 Estimates

Table 3 presents the estimates for the three demand models. The price coefficients imply a similar median own-price elasticity of approximately -9.3 , reflecting both the strength of our instruments and the fact that consumers are highly price sensitive when choosing where to purchase fuel. The coefficients of the station characteristics are likewise broadly consistent across specifications and somewhat difficult to interpret because of the inclusion of station fixed effects.

Table 3: Demand estimates

	Logit	FPD	Mixed Logit
Price (cents/liter)	-0.150 (0.041)	-0.115 (0.030)	-0.15 (0.03)
Travel Cost			-0.423(0.10)
Weighted log share ratio - <1min range		0.106 (0.034)	
Weighted log share ratio - (1,3) min range		0.049 (0.012)	
Weighted log share ratio - >3 min range		0.001 (0.011)	
Number of pumps	0.011 (0.004)	0.011 (0.004)	0.01 (0.003)
Number of islands	0.003 (0.034)	-0.001 (0.029)	-0.004 (0.018)
Repair shop indicator	-0.305 (0.111)	-0.226 (0.106)	-0.251 (0.061)
Carwash indicator	-0.102 (0.063)	-0.081 (0.051)	-0.058 (0.042)
Conv. store indicator	-0.004 (0.034)	0.010 (0.029)	-0.092 (0.040)
Self service indicator	-0.070 (0.045)	-0.046 (0.035)	-0.028 (0.029)
Med. Own price elast.	-9.84	-9.27	-9.30
Station FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Observations	14,938	14,938	14,938


Notes: Standard errors are robust to spatial correlation. Prices are in 2001-1 values. All specifications use the Sunoco-Ultramar merger, local competition, and price war instruments. The FPD specification also includes the local competition instrument version with 1, (1,3), and >3 driving minutes bandwidth.

In the Mixed Logit model, the travel cost parameter is precisely estimated and indicates

that consumers are willing to forgo savings of roughly 4.5% of their purchase value to avoid a one-minute increase in driving time, which is similar in magnitude to Houde (2012)'s results.¹⁷

In the FPD model, the coefficients of the weighted share ratio show a quick decline in the co-movement of market shares as the distance between stations increases. The coefficient that captures the co-movement in shares between stations that are between 1 to 3 minutes driving distance is half the value of the coefficient for stations that are less than 1 minute apart. The coefficient for stations more than 3 minutes apart is closer to zero and statistically insignificant.

Figure 4: Demand Models Comparison



Notes: In Panel (a), lines refer to the linear regression model. In Panels (b) and (c), bars refer to the median value, and the lower and upper brackets refer to the 25th and 75th percentiles, respectively. In Panel (c), the deviator refers to the individual station that reduces price by 3 cents.

¹⁷To obtain this number we calculate the price difference that would offset a one-minute deviation from the commuting path. This is given by the ratio λ/α , which is around 3 cpl, evaluated relative to the average price in our sample of 67 cpl.

In Figure 4 we contrast the implications that the different model estimates have for the price elasticity of demand and, consequently, for stations' price responses. As discussed in Section 2, spatial proximity is a key source of differentiation among gasoline stations in Quebec City. Consequently, the number of nearby competitors directly affects the price elasticity of demand faced by each station. Figure 4 (a) illustrates the relationship between own-price elasticity and the number of local competitors in the first bi-month of 1999 for the three demand models. In the logit model, the relationship is essentially flat: stations face an elasticity of roughly -8.3 regardless of how many competitors are within the defined radius. The mixed logit model exhibits a modest pattern of local competition, with demand becoming slightly more elastic as competition increases — from around -7.5 when a station has almost no competitor within three minutes, to approximately -8.3 when the number of competitors exceeds ten. In contrast, the FPD model generates substantially more variation, with implied elasticities ranging from about -7 to -11 depending on local market density. The patterns derived under FPD are consistent with the empirical evidence presented above for Eko stations that revealed that stations operating in denser competitive environments face more elastic demand.

The differences in substitution patterns implied by the three demand models are evident in panels (b) and (c) of Figure 4. Panel (b) presents the distribution of cross-price elasticities across station pairs at varying driving distances. Under the FPD model, cross-elasticities exhibit discrete jumps at the distance thresholds embedded in its functional-form assumptions. In contrast, the mixed logit model produces a smoother, more continuous decline in cross-elasticity as distance increases. To further illustrate the implications of the substitution patterns implied by each model, panel (c) presents the distribution of stations' price response following a price cut by one station.¹⁸ Although the magnitudes of the resulting price changes remain small, the FPD model predicts considerably stronger responses to price reductions by nearby competitors relative to the mixed logit model.

The last criterion in which we evaluate the goodness of fit of the FPD model relative to the mixed logit is by measuring their out-of-sample predictive power. To this end, we implement a cross-validation exercise based on the prediction error of market shares. For each model, we re-estimate the demand parameters while leaving one time period out of the sample, and then use the estimated parameters to predict the market shares in the omitted period. In generating these out-of-sample predictions, we set the market fixed effect

¹⁸Specifically, we cut station k 's price by 3 cpl and compute the best response of all other stations following the price cut. We repeat this procedure for all stations and collect the resulting price difference and distance to the station k that is initiating the price cut.

for period t equal to that of the preceding period ($\xi_t = \xi_{t-1}$) and impose $\Delta\xi_{jt} = \Delta\xi_{jt-1}$ for all j . We assess predictive performance using the root mean squared error, defined as $RMSE_m = \sqrt{\frac{1}{N_m} \sum_{j \in \mathcal{J}_m} (s_{jm} - \hat{s}_{jm})^2}$. Out of 27 out-of-sample predictions, FPD generates a lower RMSE in 20 cases.¹⁹ If we focus only on the prediction of the inside shares, ignoring the misfit from the prediction of total sales, then the FPD performs better in all 27 cases.²⁰

Marginal cost estimates: Marginal cost estimates are obtained under the assumption that prices are set competitively at the brand-chain level (multi-station pricing) and use only data from 1998 and 1999, which are the years we are most confident that neither price wars nor uniform-price collusion occurred.

Table 4: Marginal Cost Regression - 1998/1999

	Marginal Cost	
	FPD	Mixed Logit
<i>Observable shifter param.</i>		
log(Number of pumps)	−0.142 (0.055)	−0.056 (0.043)
Conv. store - small	−0.122 (0.073)	−0.076 (0.058)
Conv. store - medium	−0.028 (0.088)	−0.047 (0.070)
Conv. store - large	−0.273 (0.088)	−0.043 (0.069)
Repair shop indicator	0.614 (0.087)	0.221 (0.067)
Carwash indicator	0.121 (0.065)	−0.002 (0.052)
Self service indicator	−0.147 (0.075)	−0.167 (0.061)
Split service indicator	−0.476 (0.109)	−0.290 (0.087)
Constant	53.801 (1.070)	53.832 (0.884)
<i>Unobservable shock param.</i>		
3 digit zip code x Brand	0.521 (0.066)	0.129 (0.019)
3 digit zip code	0.448 (0.133)	0.073 (0.026)
Brand	0.887 (0.457)	0.663 (0.274)
Residual	0.753 (0.019)	0.566 (0.014)
Time FE	Yes	Yes
Median Markup	7.736	7.631
Median Marginal cost	53.16	53.203
Median Retail Marginal cost	10.575	10.77
Avg. within market Std(mc)	1.43	1.03
Observations	3,603	3,603

Notes: Small, medium, and large convenience stores have a footprint between 100-800 sqft, 800-1799 sqft, and greater than 1799 sqft, respectively.

Table 4 displays the summary statistics for marginal costs, markups and the pricing

¹⁹We dropped markets that are needed to construct instruments (1996, 1997, and 2000) from the out-of-sample predictions, along with very early markets that contain only one bimonth (1992, 1993, 1994).

²⁰We provide the mean, max, and min RMSE in Table 9 of Appendix A.

regression under different demand models. Median marginal costs are estimated to be very similar under the FPD and mixed logit models – 53.17 and 53.21 cpl, respectively. To isolate the part of the marginal cost that is due to wholesale and distribution costs, we also consider the retail marginal cost, derived by subtracting the price floor, which captures wholesale and transportation costs.²¹ We estimate the median retail marginal cost to be roughly 10.6 and 10.8 under FPD and mixed logit, respectively. Note that the average standard deviation of marginal costs is greater under FPD than under mixed logit, consistent with the differences in substitution patterns between the two models presented above. We discuss marginal cost dispersion in greater detail in the next section. The model also implies a median markup of 7.74 cpl, which is approximately 2.5 cpl higher than the median markup we obtain when using the price floor as the only marginal cost measure. The difference between these two margins likely reflects the fact that stations are able to obtain wholesale price discounts that more than compensate for the operational cost included in the calculation of the price floor.

The pricing regression coefficients are broadly consistent across demand models, but not precisely estimated when considering the mixed logit model. Using the number of pumps as a measure of station capacity, we find that larger stations have a lower marginal cost, consistent with our intuition. The negative coefficients on the different convenience store dummies indicate a complementarity between gasoline sales and the presence of a convenience store, which is consistent with anecdotal evidence that stations with convenience stores often use gasoline as a loss-leader product. This complementarity is not observed in the case of repair shops, which is in line with the overall trend in the industry of newer and more productive stations focusing on size and high-margin services, such as convenience stores.

Determining the extent of misallocation caused by the collusive uniform price observed in Quebec City requires knowing the distribution of costs for this period. Unfortunately, the presence of uniform prices poses a challenge for recovering the marginal cost during this period. Because we do not know how stations set prices during this period, we cannot rely on the standard approach of using the first-order conditions characterizing the equilibrium of a Bertrand-Nash pricing game to recover marginal costs after 2001. Instead, we use the cost function estimates from [Equation 12](#), based on 1998 and 1999 data, to predict what would have been the marginal costs of the stations during the 2004 September-October period. When doing so, we use the 2004 characteristics of all stations active in that period, and we update the rack price accordingly. To account for the cost variation implied by the

²¹Specifically, we regress the time fixed effect estimates on the observed price floor: $\hat{\psi}_t \sim \beta_0 + \beta_1 price\_floor_t$, and define the retail marginal cost as the difference between the marginal cost estimate and $\hat{\beta}_1 price\_floor_t$.

unobserved cost shock ω_{jt} , we draw 100 samples of cost shocks from its empirical distribution. The counterfactual and misallocation results we report are obtained by averaging the relevant variables across these samples. We present these results in the misallocation section that comes next.

6 Collusion and Misallocation

In this section, we use a model-based approach to evaluate the extent of misallocation caused by the collusive uniform price observed in Quebec City after 2001. Our analysis focuses on understanding two dimensions of misallocation. One is static, driven by the contemporaneous allocation of sales to high-cost stations. The other is dynamic, driven by retailers' investment decisions, such as entering and exiting the market and spending on cost-reduction technologies.

A necessary condition for misallocation is the presence of meaningful differences in marginal costs across stations. Therefore, our starting point is to establish the existence of dispersion in marginal costs. [Table 5](#) presents marginal cost estimates for 2004 Sep-Oct and reveals important dispersion: the difference between the 10th and 90th percentile marginal costs is 4.4 cpl, ranging from 70.36 to 74.76 (or from 7.68 to 12.08 for the retail marginal cost). There is slightly less dispersion under the mixed-logit model. For the reasons discussed in the previous section, we focus here on results derived under the flexible parametric demand model, but misallocation results for the mixed-logit can be found in the Appendix.

6.1 Counterfactual Scenarios

Using the marginal cost estimates for 2004 Sep-Oct, we simulate five counterfactuals to assess the contemporaneous misallocation of production under the observed collusive uniform prices. The first set of counterfactuals is based on four supply assumptions: (i) marginal cost pricing, (ii) single-station pricing, (iii) multi-station pricing, and (iv) joint-profit maximization. The counterfactual allocation based on these assumptions assess the potential factors driving the production misallocation, such as frictions arising from horizontal differentiation and the market power implied by these frictions.²² The fifth counterfactual provides the cost-minimizing benchmark. To compute this benchmark, we use the commuting pattern

²²As with the marginal cost estimates, multi-station pricing is computed at the brand-chain level.

data and an allocation algorithm that assigns production to the most efficient seller while accounting for spatial differentiation and the limited capacity of each station. The algorithm proceeds as follows:

- a. Let C consumers arrive to the market needing to buy gasoline. Each station k can serve up to $200 \times S_k$, where S_k is the number of pumps in station k .²³ As before, we define each consumer’s choice set by their commuting path, e.g. stations within x minutes of extra driving time).
- b. We draw a consumer from the C consumers who arrived to the market. The consumer goes to the lowest-cost station in their choice set and purchases their desired quantity. The number of available pumps in that station is updated to $200 \times S_k - 1$. If no station is within the extra driving time bandwidth, that consumer does not purchase gasoline. If the lowest-cost station is already at capacity, the consumer goes to the next lowest-cost station.
- c. We keep drawing consumers until the total quantity sold equals the total quantity we observe in the data under the collusive uniform price.

We repeat the process described in a) to c) 100 times for different sequences of consumer arrival and report the averages across simulations. We consider two different consumer choice sets: (i) all stations in the market, and (ii) 3.3 minutes extra driving time from the commute path. The former yields the minimum production cost and corresponds to the maximal production efficiency. The latter is closer to the average extra driving time observed in the data.

Table 5 displays summary statistics on markup and sales implied by the different conduct assumptions. The markup summary statistics highlight two relevant patterns. First, markups implied by multi-station pricing (7.99 cpl) are substantially closer to markups implied by single-station pricing (7.28 cpl) than markups implied by joint profit maximization (15.68 cpl). Therefore, even when brands internalize consumer diversion within their network of stations, the market is still closer to the competitive benchmark. Second, markups implied by the collusive uniform price (11.4 cpl) are 43% higher than markups under multi-station pricing but substantially lower than markups under joint profit maximization. This result illustrates that, despite the considerable profitability of the collusive equilibrium reached in

²³We obtain the factor of 200 by assuming that a station is open 14 hours per day and that each customer takes approximately four minutes to get in, fill the tank, and exit.

Quebec City, stations were not able to fully exploit potential profits. The third panel shows that the markup increase is followed by a reduction in the total volume sold. This reduction in aggregate volume indicates that a non-negligible fraction of consumers reacted to the increase in price levels from the collusive uniform price equilibrium by reducing consumption, which creates another source of inefficiency besides misallocation.

Table 5: Counterfactual summary statistics for 2004

	Mean	Std.	P10	P90
Marginal Cost (cents per liter)	72.56	1.71	70.36	74.76
Marginal Cost Retail (cents per liter)	9.87	1.71	7.68	12.08
<i>Markups (cents per liter)</i>				
Collusive Uniform Prices	11.40	1.82	9.19	13.63
Single-station Bertrand	7.28	0.92	6.00	8.40
Multi-station Bertrand	7.99	1.04	6.61	9.27
Joint Profit Maximization	15.68	0	15.68	15.68
<i>Quantity (ths liters/day)</i>				
Collusive Uniform Prices	6.67	3.93	2.85	11.50
P=MC	9.92	6.49	3.34	18.12
Single-station Bertrand	8.05	4.79	2.97	14.24
Multi-station Bertrand	7.82	4.66	3.02	13.83
Joint Profit Maximization	5.33	3.48	1.79	9.73

Notes: Marginal cost and markups are in 2001-1 value. Marginal cost retail is computed according to the steps described in footnote 21.

6.2 Static Misallocation

The covariance between market shares and marginal costs is informative about the sales share accrued by high-cost firms. This correlation is positive if high-cost stations have relatively large market shares (strong misallocation) and negative if low-cost stations have higher-than-average market shares. Since there is limited vertical differentiation between stations in our market, we expect this covariance to be negative under competitive pricing, as firms with lower marginal costs tend to have lower prices and, all else equal, higher sales. However, if prices are uniform, consumers will sort to stations based solely on station characteristics and convenience. Therefore, the share of sales that goes to high marginal cost stations increases, and the covariance term reflects the correlation between consumer preferences for stations and cost.

Table 6 displays the Olley and Pakes marginal cost decomposition for the uniform collusive price and the counterfactual scenarios. Here we drop the multi-station Bertrand Nash counterfactual to focus just on the single-station counterfactual, since results are very similar and the single-station counterfactual has the advantage of eliminating whatever market power stems from having a network of stations.

We also express our measure in relative terms and divide the covariance implied by each pricing counterfactual by the covariance obtained from the cost-minimizing benchmark. Interestingly, the covariance between market shares and marginal cost is negative for all counterfactuals, indicating that more efficient, low-cost stations tend to have valuable amenities and/or locations. However, uniform-price collusion significantly weakens this relationship as evidenced by how close its covariance is to zero, and how small it is compared to the efficient benchmark (i.e., 5.3% of the covariance obtained under the cost-minimizing benchmark: -0.136/-2.555).

Table 6: Marginal cost decomposition (based on FPD estimates)

	c_t	$\bar{c}_t$	$\sum_j \Delta c_{jt} \Delta s_{jt}$	Ratio	Extra time (min.)
<u>Cost Minimizing</u>					
Choice set: all stations	6.598	9.870	-3.272	1.280	23.42
Choice set: 3.3 min. range	7.315	9.870	-2.555	1.000	1.74
<u>Market Allocation</u>					
Collusive Uniform Price	9.734	9.870	-0.136	0.053	1.73
Price = Marginal Cost	9.363	9.870	-0.507	0.200	
Single-station Bertrand	9.428	9.870	-0.442	0.173	

Notes: See Equation 2 for definitions of c_t , $\bar{c}_t$ and $\sum_{j=1}^{J_t} \Delta c_{jt} \Delta s_{jt}$. Values are in cents per liter.

Similar to Asker, Collard-Wexler, and De Loecker (2019), the comparison between collusive uniform price outcomes and the cost-minimizing benchmark provides the extent of production misallocation in the market. However, this comparison abstracts from demand frictions and the wedge between the decentralized market allocation and the cost-minimizing allocation that these demand frictions may create. As such, comparing with the cost-minimizing benchmark is likely to overstate production misallocation. To better understand the role of demand frictions in the misallocation of production, we compare to the marginal-cost pricing and single-station Bertrand-Nash pricing counterfactuals.

In the cost-minimizing benchmark, the allocation is fully determined by the marginal cost of each station, while accounting for capacity and spatial differentiation. In contrast, in the marginal cost pricing counterfactual, relative prices fully reflect stations' efficiencies, but

consumer preferences also determine stations' production. Therefore, the covariance ratio implied by marginal cost pricing captures the misallocation caused by product differentiation. Furthermore, in the single-station pricing counterfactual, relative prices reflect marginal costs and markups, which depend on the degree of product differentiation. As such, comparing the covariance ratios implied by single-station pricing with those implied by marginal cost pricing allows us to assess the misallocation caused by market power. Lastly, in the observed uniform price setting, firms exert market power, prices convey no information about stations' efficiencies, and production is determined solely by consumer preferences. Therefore, the difference in the covariance ratio between the single-station pricing counterfactual and the observed uniform price allocation captures the allocative inefficiency implied only by the collusive uniform price strategy.

As discussed, the comparison to the cost-minimizing benchmark does not account for demand-side frictions, and doing so may overstate the actual contribution of collusion towards misallocation. Indeed, our findings suggest that product differentiation is the most relevant source of production misallocation. The single-station pricing allocation corresponds to 17.3% of the cost-minimizing allocation, indicating that uniform price collusion explains 12% of the productive inefficiency. The marginal-cost pricing allocation corresponds to 20% of the cost-minimizing allocation. Because this is virtually identical to the single-station pricing allocation, we have that the market power following from product differentiation leads to a small production inefficiency relative to perfect competition. The remaining sources of cost misallocation are the differences in preferences for stations' amenities and spatial differentiation.²⁴ Put differently, if instead of comparing to the cost-minimizing benchmark, we consider the possible efficiency gains from making the market competitive, our findings suggest that eliminating collusion would get us roughly 80% of the way there (0.3 cpl) while eliminating market power would take us the rest of the way (0.06 cpl).

It is important to point out that these results hinge on the underlying demand model. The logit and mixed logit models imply nearly constant markups. As a result, nearly 100% of the misallocation from imperfect competition will be due to collusion. Results for alternative demand specifications are presented in the Appendix.

²⁴Comparing our benchmark with the unrestricted cost-minimizing benchmark in which a consumer is assigned to a station regardless of the driving distance, we find that eliminating all sources of product differentiation would raise production efficiency by 28%.

6.3 Implications for Damage Calculations

Antitrust authorities typically define cartel damages as the part of the consumer surplus that the cartel appropriates by raising prices. These damages are calculated by determining the average price that would arise in the absence of collusion (but-for-collusion price), subtracting it from the average observed collusive price, and then multiplying by the total quantity sold under collusion. In this section we provide evidence that this measure of damages understates the true harm from collusion in differentiated-product settings by failing to account for misallocation. Specifically, we use our model estimates and counterfactual simulations to show that the average but-for-collusion price can be adjusted to account not only for the damages caused by higher prices, but also for the inefficiencies resulting from the allocation of production to high-cost firms.

Taking misallocation into account, cartel damages can be expressed as:

$$Q^{cartel} \times \left[\sum_j w_j^{cartel} (p_j^{cartel} - p_j^{comp}) + \sum_j (w_j^{cartel} - w_j^{comp}) p_j^{comp} \right], \quad (13)$$

where the superscript “cartel” and “comp” index outcomes during collusive and competitive (but-for-collusion) scenarios, respectively, and where w_j^{cartel} and w_j^{comp} are j ’s share of Q^{cartel} and Q^{comp} , respectively. The first term inside the square brackets in [Equation 13](#) is the measure of damage normally used in antitrust cases. This measure holds the quantity (share) of each product fixed at the collusive quantity and multiplies this quantity by the usual measure of cartel overprice. The second term in square brackets captures the damages caused by the reallocation of production to high-cost firms. From our model and counterfactual simulations, we have all the components required to compute each element in [Equation 13](#).

It is well-known that the typical damage calculation also ignores the deadweight loss arising from the reduction in sales stemming from higher prices. We can also calculate this to get a sense of the relative contribution of each distortion towards total damages. We quantify the deadweight loss component as the difference between the change in total consumer surplus and the change in revenue captured by our proposed measure of cartel damage. Monardo (2025) shows that consumer surplus in the FPD setting for a given vector p of prices is equal to $-\log(s_0(p))/\alpha$.

Using the observed market outcomes for 2004-5 and the competitive prices and shares implied by the FPD demand and single-product Bertrand-Nash pricing model, we calculate total damages in the order of C\$84,745 per day where the effect of misallocation corresponds


to 2% and the deadweight loss component to 11.5%.

6.4 Investment Decisions and Dynamic Misallocation

In addition to causing static misallocation, the collusive uniform price can also increase inefficiencies by distorting incumbents' incentives to invest in cost-reducing technologies, entrants' technology choices, and inefficient firms' incentives to exit the market. These market structure distortions can have a lasting effect on prices even after the cartel collapses.

In Figure 6, we use our model estimates to demonstrate the impact of conduct on the incentives to invest in cost-reducing technologies. In Panel (a), we show the percentage increase in profits from a one-cent reduction in marginal cost for one station in the Bertrand-Nash and collusive uniform price scenario. The increase in profits under the competitive scenario is, on average, 5 p.p. higher than in the collusive uniform price one. More interestingly, this difference in profit effect is increasing with the number of nearby competitors, where the business-stealing effect is the greatest and there is less incentive to pass through cost changes. With the collusive uniform price, the benefit of investing in a technology with a lower cost is to increase profits in proportion to the quantity sold. Under competition, this effect should be greater, because not only is the quantity sold larger for most stations but also because price adjustments allow more efficient firms to attract more customers and increase profits. This indirect price effect, is even more important for stations that face intense local competition.

Figure 6: Return to investment and incentives to remain in / enter the market



(a) Profit effect of a 1 cent cost reduction

(b) Profit ratio $\times$ Marginal cost draws

Notes: Marginal cost is in 2001-1 values. In Panel (b), the black line refers to the linear regression line. Marginal costs are drawn from a normal distribution according to Equation 12.

Because entry and exit decisions depend on variable profits, the difference in variable

profits across various levels of marginal cost is indicative of the exit incentive and the type of technology that entrants would pick. Panel (b) in [Figure 6](#) displays the scatter plots of the variable profit ratio between Bertrand-Nash and collusive equilibrium implied by different levels of marginal cost. The positive correlation between the expected profit ratio and marginal costs indicates that high-cost stations benefit more from the cartel than do low-cost incumbents. Assuming that the fixed cost of operating an efficient station is at least as high as for operating an inefficient station, we interpret this model prediction as saying the following: incumbent stations are more likely to continue operating and entrants prefer less efficient technologies under uniform price coordination than under competition.

Empirical Evidence on Industry Turnover The preceding analysis suggests that, under our model assumptions, a collusive arrangement involving uniform pricing may significantly affect industry turnover. Specifically, it can reduce firm exit, distort the operational costs chosen by potential entrants, and influence incumbents' incentives to invest in cost-reducing technologies. Next, we assess the importance of each channel for the observed evolution of the sales allocation in Quebec City.

To examine the dynamic drivers of industry efficiency, we follow Melitz and Polanec (2015) and decompose the changes in sales-weighted average cost into two components: changes driven by incumbent firms and changes driven by industry turnover. For a given period t , we classify firms into three groups: survivors (S), which are active in both t and $t + 1$; exiters (X), active in t but not in $t + 1$; and entrants (E), active in $t + 1$ but not in t . Applying the Olley-Pakes decomposition to surviving firms only, the sales-weighted average cost for two consecutive periods can be expressed as:

$$\begin{aligned}\hat{c}_t &= \bar{c}_{s,t} + cov(s_{j/s,t}, c_{j,t}) + s_{X,t}(\hat{c}_{X,t} - \hat{c}_{S,t}) \\ \hat{c}_{t+1} &= \bar{c}_{s,t+1} + cov(s_{j/s,t+1}, c_{j,t+1}) + s_{E,t+1}(\hat{c}_{E,t+1} - \hat{c}_{S,t+1})\end{aligned}$$

Taking the difference between the two periods yields:

$$\Delta_t \hat{c} = \underbrace{\Delta_t \bar{c}_s + \Delta_t cov(s_{j/s}, c_j)}_{\text{Surviving firms}} + \underbrace{s_{E,t+1}(\hat{c}_{E,t+1} - \hat{c}_{S,t+1}) - s_{X,t}(\hat{c}_{X,t} - \hat{c}_{S,t})}_{\text{Industry turnover}} \quad (14)$$

In [Equation 14](#), we decompose the change in the sales-weighted average cost into four components: (i) the change in the average marginal cost of surviving firms, (ii) the change in the covariance between sales shares and marginal costs among survivors, (iii) the difference

in average costs between entrants and survivors, and (iv) the difference in average costs between exiters and survivors. The first component captures technological changes faced by firms operating in both time periods. The second component accounts for sales reallocation between incumbent stations. Both changes are possibly driven by differential incentives to investment in cost-reducing technologies faced by the surviving firms. The third and fourth terms capture the role of industry turnover in the evolution of industry efficiency.

Table 7: Dynamic Olley and Pakes decomposition in Quebec City

	# Stations		$\Delta\hat{c}$	Determinants industry efficiency		
	t_0	t_1		$\Delta\bar{c}$	$\Delta cov(s_s, c_s)$	Turnover
<i>Panel A: 1992 to 1998</i>						
All stations	369	303	-0.273	-0.021	-0.155	-0.097
Areas of						
low density	44	48	-0.206	0.010	-0.179	-0.037
mid density	150	125	-0.233	-0.006	-0.131	-0.096
high density	175	130	-0.285	-0.042	-0.152	-0.091
<i>Panel B: 2001 to 2007</i>						
All stations	281	258	-0.024	-0.052	0.021	0.007
Areas of						
low density	44	44	-0.099	-0.024	-0.032	-0.043
mid density	119	111	-0.048	-0.090	0.037	0.004
high density	118	103	0.070	-0.021	0.043	0.048

Notes: Low-density refers to the set of stations that face ≤ 3 stations in a 3-minute driving range, while mid-density is stations facing $n \in (3, 10)$ competitors, and high density $n \geq 10$.

Table 7 presents the results of the dynamic decomposition for Quebec City across two periods: the pre-collusion era from 1992 to 1998 and the collusive era from 2001 to 2007. During the pre-collusion period, the sales-weighted average marginal cost declined by 0.237 cents per liter. Our decomposition indicates that 56% of these efficiency gains stem from the reallocation of production toward more efficient surviving firms (0.155/0.273), approximately 8% from technological improvements experienced by those firms (0.021/0.273), and the remainder from industry turnover—specifically, the entry of large-scale, low-cost stations and the exit of inefficient ones, typically those lacking automation or service capacity.

In contrast, between 2001 and 2007, the decline in the sales-weighted average marginal cost was only 0.024 cents per liter. If technological change had been the only active force, the decrease would have reached 0.052 cents per liter. Instead, both the reallocation of sales

among stations and industry turnover contributed positively to marginal costs. Consistent with the mechanisms emphasized by our empirical model, less efficient stations expanded their sales shares between 2001 and 2007, and entrants (exiters) were on average less (more) efficient than incumbents.

To further assess the predictions of our model, we compute the dynamic Olley–Pakes decomposition separately for stations facing different levels of competition. In particular, we estimate the decomposition for stations facing fewer than three competitors, between three and ten competitors, and more than ten competitors within a three-minute driving radius. A direct comparison of the sales-weighted marginal-cost trends between high- and low-density areas reveals a difference in trend between these two groups of $0.248 = (0.07+0.285)-(-0.099+0.206)$ cents per liter. Looking at the decomposition of this shift, we can see that all density areas experienced a decline in the covariance between sales shares and marginal costs. However, while low-density areas continued to exhibit reductions in misallocation among incumbents between 2001 and 2007, mid- and high-density areas experienced increases. Turnover effects also declined across all density groups relative to the 1992–1998 period, but the deterioration was substantially larger in high-density areas. Whereas exiters/entrants played a key role in lowering the sales-weighted average marginal cost during the 1990s, they became a source of upward pressure in the 2000s. We interpret these patterns as evidence that dynamic misallocation worsened disproportionately in high-density areas relative to low-density ones. This is consistent with the findings presented in [Figure 6](#) that highlighted the fact that the difference in profit change from a one-cent reduction in marginal cost in Bertrand-Nash vs collusion is increasing with the number of nearby competitors.

7 Conclusion

This article proposes a method for quantifying misallocation that recovers marginal costs from an equilibrium model of supply and demand in a competitive setting, using only observed prices and quantities. This approach allows us to measure productive inefficiency even when cost data are unavailable, products are differentiated, and firms sell multiple products. Counterfactual simulations reveal how collusion distorts allocation. Static inefficiency is quantified following Olley and Pakes (1996), and we compare outcomes to an efficient benchmark that incorporates production capacity and consumer heterogeneity. We apply the framework to Quebec City’s retail gasoline market, which transitioned from competitive pricing to a uniform-price regime, and decompose total misallocation into components arising

from collusion, market power, and demand frictions. While collusion contributes appreciably to misallocation, product differentiation accounts for the bulk of the static inefficiency, and uniform pricing further undermines dynamic efficiency by dampening investment and entry incentives.

References

- Adams, B. and K. R. Williams (2019, February). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics* 11(1), 124–56.
- Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review* 100(3), 724–762.
- Asker, J., A. Collard-Wexler, and J. De Loecker (2019, April). (Mis)allocation, market power, and global oil extraction. *American Economic Review* 109(4), 1568–1615.
- Assad, S., R. Clark, D. Ershov, and L. Xu (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy* 132(3), 723–771.
- Berry, S., J. Levinsohn, and A. Pakes (1995, July). Automobile Prices in Market Equilibrium. *Econometrica* 63(4), 841.
- Berry, S. T., A. Gandhi, and P. A. Haile (2013). Connected substitutes and invertibility of demand. *Econometrica* 81(5), 2087–2111.
- Bridgman, B., S. Qi, and J. Schmitz (2015). Cartels destroy productivity: Evidence from the New Deal sugar manufacturing cartel, 1934–74. *Staff Report, Federal Reserve Bank of Minneapolis* (519).
- Byrne, D. and N. deRoos (2019). Learning to coordinate: A study in retail gasoline. *American Economic Review* 109, 591–619.
- Carranza, J. E., R. Clark, and J.-F. Houde (2015). Price controls and market structure: Evidence from gasoline retail markets. *The Journal of Industrial Economics* 63(1), 152–198.
- Carrera, F. (2022). Cartels, entry, and productivity: Evidence from the Chilean nitrate cartels. Working paper.
- Chaves, D. and M. Duarte (2025). The inner workings of a hub-and-spoke cartel in the automotive fuel industry. *American Economic Journal: Microeconomics* 17(1), 41–65.
- Clark, R., D. Coviello, J.-F. Gauthier, and A. Shneyerov (2018). Bid rigging and entry deterrence in public procurement: Evidence from an investigation into collusion and corruption in Quebec. *The Journal of Law, Economics, and Organization* 34(3), 301–363.

- Clark, R., M. Duarte, and J. Houde (2024). Cartel instability and price wars: Retail gasoline in Canada. In *Cartels Diagnosed: New Insight on Collusion*.
- Clark, R., I. Horstmann, and J. Houde (2024). Hub-and-spoke cartels: Theory and evidence from the grocery industry. *American Economic Review* 114, 783–814.
- Clark, R. and J.-F. Houde (2013). Collusion with asymmetric retailers: Evidence from a gasoline price fixing case. *American Economic Journal: Microeconomics* 5(3), 97–123.
- Clark, R. and J.-F. Houde (2014). The effect of explicit communication on pricing: Evidence from the collapse of a gasoline cartel. *Journal of Industrial Economics* 62, 191–228.
- DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in US retail chains. *The Quarterly Journal of Economics* 134(4), 2011–2084.
- Edmond, C., V. Midrigan, and D. Y. Xu (2023). How costly are markups? *Journal of Political Economy* 131(7), 1619–1675.
- Erutku, C. and V. Hildebrand (2010). Conspiracy at the pump. *Journal of Law & Economics* 53, 223–237.
- Gandhi, A. and J.-F. Houde (2019). Measuring substitution patterns in differentiated-products industries. Technical Report Working Paper No. 26375, National Bureau of Economic Research.
- Genesove, D. and W. P. Mullin (2001, Jun.). Rules, communication, and collusion: Narrative evidence from the sugar institute case. *American Economic Review* 91(3), 379–398.
- Harberger, A. C. (1954). Monopoly and resource allocation. In *American Economic Review, Papers and Proceedings*, pp. 77.
- Harrington, J. and A. Skrzypacz (2011). Private monitoring and communication in cartels: Explaining recent collusive practices. *American Economic Review* 101, 2155–2191.
- Harrington, Joseph E., J. (2016). Heterogeneous firms can always collude on a minimum price. *Economics Letters* 138(0), 46–49.
- Hastings, J. S. (2004, February). Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California. *American Economic Review* 94(1), 317–328.
- Hitsch, G. J., A. Hortacsu, and X. Lin (2021). Prices and promotions in US retail markets: Evidence from Big Data.

- Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review* 102, 2147–2182.
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics* 124(4), 1403–1448.
- Hyttinen, A., F. Steen, and O. Toivanen (2019). An anatomy of cartel contracts. *The Economic Journal* 129, 2155–2191.
- Igami, M. and T. Sugaya (2022). Measuring the incentive to collude: The vitamin cartels, 1990–1999. *Review of Economic Studies* 89(3), 1460–1494.
- Kim, M., R. Mansley, N. Miller, M. Remer, and M. Weinberg (2023). Price-fixing allegations in the canned tuna industry: A look at the data. *Anitrust Bulletin* 68, 154–163.
- Knittel, C. R. and V. Stango (2003). Price ceilings as focal points for tacit collusion: Evidence from credit cards. *American Economic Review* 93(5), 1703–1729.
- Lu, J. and K. Saito (2022, feb). Mixed logit and pure characteristics models. SSWP 1461, California Institute of Technology.
- McFadden, D. and K. Train (2000). Mixed mnl models for discrete response. *Journal of applied Econometrics* 15(5), 447–470.
- Melitz, M. J. and S. Polanec (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *The RAND Journal of Economics* 46, 362–375.
- Midrigan, V. and D. Y. Xu (2014, February). Finance and misallocation: Evidence from plant-level data. *American Economic Review* 104(2), 422–58.
- Miller, N. and M. Weinberg (2017). Understanding the price effects of the Miller/Coors joint venture. *Econometrica* 85, 1763–1791.
- Monardo, J. (2021). Measuring Substitution Patterns with a Flexible Demand Model. *SSRN Electronic Journal*.
- Monardo, J. (2025, March). Measuring substitution patterns with a flexible demand model. Technical report, SSRN Electronic Journal. Available at SSRN: <https://ssrn.com/abstract=3921601> or <http://dx.doi.org/10.2139/ssrn.3921601>.
- Moreau, F. and L. Panon (2023). How costly are cartels? Working paper.
- Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6), 1263–1297.

- Petit, L. T., R. G. Kemp, and J. van Sinderen (2015). Cartels and productivity growth: An empirical investigation of the impact of cartels on productivity in the Netherlands. *Journal of Competition Law & Economics* 11(2), 501–525.
- Pinkse, J. and M. E. Slade (2004). Mergers, brand competition, and the price of a pint. *European Economic Review* 48(3), 617–643.
- Pinkse, J., M. E. Slade, and C. Brett (2002). Spatial price competition: A semiparametric approach. *Econometrica* 70(3), 1111–1153.
- Restuccia, D. and R. Rogerson (2017, August). The causes and costs of misallocation. *Journal of Economic Perspectives* 31(3), 151–74.
- Röller, L.-H. and F. Steen (2006, March). On the workings of a cartel: Evidence from the Norwegian cement industry. *American Economic Review* 96(1), 321–338.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.
- Sorgard, L. and F. Steen (2024). Coordinated rebate reductions and semi-collusion: Retail gasoline in Sweden. Forthcoming in, *Cartels Diagnosed: New Insight on Collusion*.
- Starc, A. and T. G. Wollmann (2025). Does entry remedy collusion? evidence from the generic prescription drug cartel. *American Economic Review* 115(5), 1400–1438.
- Wang, Z. (2008). Collusive communication and pricing coordination in a retail gasoline market. *Review of Industrial Organization* 32(1), 35–52.
- Wang, Z. (2009). (mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation. *Journal of Political Economy* 117(6).


A Appendix - Additional Tables & Figures

Table 8: Data set coverage

Dataset	Frequency	Time window	Time periods	Station coverage
Kent	Bi-monthly	1995-1 to 2001-6	42	306
Regie	Weekly	1997-12-08 to 2004-02-23	325	13
MJervin	Weekly	2001-01-16 to 2012-11-06	617	15
Eko	Daily	2000-08-29 to 2000-11-27	91	29

Notes: Station coverage is the average number of stations surveyed during each period.

Figure 8: Quebec City - Station Location and Local Competition - Fall 1999




Note: Stations facing more than 10 local competitors are indicated in red and those facing fewer than three in green.

Table 9: Model Fit

	Share			Inside share		
	ML	FPD	Diff	ML	FPD	Diff
mean	0.842	0.443	0.399	0.716	0.256	0.459
max	2.289	2.134	1.604	1.621	0.429	1.348
min	0.370	0.188	-0.318	0.285	0.171	0.038
# Mkts FPD<ML	20/27			27/27		

Figure 9: Marginal Cost Distribution implied by FPD demand model



B Appendix - Mixed Logit Results

In this appendix, we replicate the static misallocation analyzes using the estimates from the mixed-logit demand model and contrast these results with those obtained with our main specification, the FPD model.

Table 10 displays summary statistics on the marginal costs obtained with the mixed-logit demand model, as well as summary statistics on the markup and sales implied by the different conduct assumptions recovered from these costs. Comparing the mixed logit results with those in Table 5, we find that the mixed logit and FPD imply similar levels of marginal cost, but under mixed logit, marginal costs are less dispersed, implying a reduced scope for misallocation. The reduced scope for misallocation in the mixed-logit is a direct consequence of the substitution patterns implied by the model. Both models predict similar levels for own-price elasticities. However, in the mixed-logit, the i.i.d random utility shocks attenuate the role of spatial differentiation, and cross-price elasticities remain proportional to market shares. In contrast, with FPD, cross-price elasticities are higher for nearby competitors and decay sharply as stations become farther apart. Consequently, the FPD’s local market structure has a stronger effect on markups and, by extension, on marginal costs.

Table 10: Counterfactual summary statistics obtained with mixed-logit demand

	Mean	Std.	P10	P90
Marginal Cost (cpl)	72.63	1.25	71.02	74.25
Marginal Cost Retail(cpl)	9.90	1.25	8.28	11.51
<i>Markups</i>				
Collusive Uniform Prices	11.32	1.38	9.70	12.98
Single-station Bertrand	7.31	0.41	7.03	7.72
Multi-station Bertrand	8.04	0.74	7.21	8.89
Joint Profit Maximization	16.68	3.05	13.76	22.22
<i>Quantity (1000L)</i>				
Collusive Uniform Prices	6.67	3.92	2.83	11.54
P=MC	9.53	5.94	3.38	17.48
Single-station Bertrand	7.89	4.75	2.91	14.11
Multi-station Bertrand	7.69	4.68	2.99	13.83
Joint Profit Maximization	5.30	3.28	1.89	9.64

Notes: Marginal cost and markups are in 2001-1 value. Marginal cost retail is computed according to the steps described in footnote 21.

Table 11 displays the covariance components of the OP decomposition and the relative measure of allocative efficiency - ratio with the denominator equal to the covariance implied by the cost-minimizing benchmark - for two cases: (i) FPD; and (ii) Mixed-logit. Overall, the patterns implied by the mixed-logit are similar to the patterns implied by the FPD model. The most noticeable differences are in the scope for misallocation predicted by each model and in the role of product differentiation in driving productive inefficiencies. The result that the scope for misallocation under FPD is greater than under mixed-logit is not surprising, as the marginal cost dispersion implied by the former is higher than that implied by the latter. Interestingly, the role of product differentiation in creating production misallocation is higher under mixed-logit than FPD. With mixed logit, the allocative efficiency under marginal-cost pricing is 15.6% of the cost-minimizing benchmark, whereas with FPD, it is 20%. This result is consistent with the i.i.d preference shocks embedded in the mixed-logit model having an impact on driving misallocation and provides further support for the adequacy of the FPD model as our main specification.

Table 11: Static misallocation - FPD and mixed-logit

	FPD		Mixed Logit	
	$\sum_j \Delta c_{jt} \Delta s_{jt}$	Ratio	$\sum_j \Delta c_{jt} \Delta s_{jt}$	Ratio
<u>Cost Minimizing</u>				
Choice set: all stations	-3.272	1.28	-2.338	1.289
Choice set: 3.3 min. range	-2.555	1.000	-1.813	1.000
<u>Market Allocation</u>				
Collusive Uniform Price	-0.136	0.053	-0.083	0.046
Price = Marginal Cost	-0.507	0.200	-0.283	0.156
Single-station Bertrand	-0.442	0.173	-0.275	0.152

Notes: FPD replicates the numbers in Table 6. The mixed-logit results are obtained using the mixed-logit demand system and the implied marginal costs.

To further evaluate the role of i.i.d random utility shocks in the mixed logit model, we analyze how each element of the random utility function specification determines consumer's choice. In the mixed logit model, consumers choose stations based on stations' vertical (δ_j) and horizontal ($\lambda D_{ij} + \epsilon_{ijt}$) attributes. As such, one potential concern is that the distributional assumption on the random-utility shocks may be a key driver of misallocation in this model. To investigate whether this is the case, we compute the Olley and Pakes decomposition while controlling for various features of our demand model: (i) only vertical differentiation ($\lambda = 0$ and $\epsilon = 0$ for all i, j), with sorting determined only by observed vertical

attributes, and (ii) no random utility shock ($\epsilon = 0$) with sorting determined by observed vertical attributes and spatial differentiation. To consider the role of price levels and relative prices, we do so for the marginal-cost pricing, single-station pricing and observed uniform prices

Table 12 displays the measure of allocative efficiency in relative terms. To this end, we divide the covariance between market-share and marginal costs for each scenario by the covariance implied by the cost-minimizing benchmark (-1.813). The first and second rows display the covariances implied by our elements-of-choice decomposition exercise. The third row displays the measures implied by our full model (same as in Table 11).

To assess how the different preference attributes affect the measure of production misallocation implied by the mixed logit model, let us move from row to row as we add a different choice component. By comparing the outcomes in the first row (“vertical”) with those in the second row (“vertical and spatial”), we measure how spatial differentiation influences production misallocation. As this comparison shows, spatial differentiation is an important source of misallocation. However, the results appear to be influenced by price levels and dispersion. By comparing the outcomes in the second row with those from the full model displayed in the last row, we assess the contribution of the i.i.d. random utility shock (logit shock) to production misallocation. As the comparison indicates, the random utility component is an important source of misallocation regardless of the pricing assumption. This result provides further evidence of the inadequacy of the mixed logit model in our setting.

Table 12: Misallocation and elements of choice

	$\sum_j \Delta c_{jt} \Delta s_{jt}$		
	P=MC	Single-station pricing	Observed prices
Vertical differentiation ($\lambda = 0$ and $\epsilon = 0$)	0.296	0.587	0.01
Vertical and spatial differentiation ($\epsilon = 0$)	0.356	0.351	0.113
Full model	0.156	0.152	0.046

Notes: Each entry in the table is defined as the covariance between market share and marginal costs obtained in each setting divided by the covariance obtained from the cost-minimizing benchmark under the mixed-logit model.