

Price controls and market structure: Evidence from gasoline retail markets*

Juan Esteban Carranza Robert Clark Jean-François Houde
ICESI & Banco HEC Montréal Wharton & NBER
de la República (Colombia)

October 22, 2013

Abstract

In this paper we study the effect of price floor regulations on the organization and performance of markets. The standard interpretation of the effects of these policies is concerned with short-run market distortions associated with excess supply. Since price controls prevent markets from clearing, they lead to higher prices. While this analysis may be correct in the short-run, it does not consider the dynamic equilibrium consequences of price controls. We demonstrate that price floor regulations can have important long-run effects on the structure of markets by crowding them and creating endogenous barriers to entry for low-cost retailers. Moreover, we show that these factors can indirectly lower productivity and possibly even prices. We test this in the context of an actual regulation imposed in the retail gasoline market in the Canadian province of Québec and show that the policy led to more competition between smaller/less efficient stations. This resulted in lowered sales, and, despite the reduction in efficiency, did not increase prices.

1 Introduction

Over the last twenty years, many retail markets around the world have experienced significant restructuring, associated with the exit of small independent stores and the entry of large-scale chains. These changes were triggered by technological innovations that lowered the marginal cost of serving consumers, at the expense of higher fixed costs of operation.¹ The success of Walmart is

*This paper has previously circulated under the title "Price controls and competition in gasoline retail markets". We have benefited from excellent discussions by Matt Lewis and Jeff Prince, and from helpful conversations with Victor Aguirregabiria, Jason Allen, Allan Collard-Wexler, Steven Durlauf, Gautam Gowrisakaran, Ig Horstmann, Robert Gagné, Amit Gandhi, John Kennan, Rasmus Lentz, Greg Lewis, Salvador Navarro, Jack Porter, Andrew Sweeting, Chris Taber, Joel Waldfogel, and seminar participants at Wisconsin, IIOC-2009, UBC-Sauder Summer Conference, the SITE-2009 conference, the NBER Summer Institute (IO), and the Institute for Computational Economics conference (Chicago). Correspondence to Juan Esteban Carranza, Email: juanes@ssc.wisc.edu; Robert Clark: HEC Montreal, Phone: (514) 340-7034, Email: robert.clark@hec.ca; Jean-François Houde: Wharton School, University of Pennsylvania. E-mail: houde@wharton.upenn.edu

¹See Foster, Haltiwanger, and Krizan (2006) for an empirical analysis of these trends in the context of U.S. retail markets, and Campbell and Hubbard (2010) for an analysis of the reorganization of service stations on U.S. highways.

a well documented example (see Jia (2008) and Holmes (2010)), but similar patterns exist in other markets. For instance, the North-American retail hardware and gasoline markets experienced important shifts towards fewer and larger volume retail outlets.

In some cases, lobbying groups were able to convince local and state governments to impose various kinds of price-control regulations in order to protect small independent retailers from this reorganization.² A common example of this type of regulation is a below-cost law, also known as a “fair-trade” policy, which prevents firms from posting prices below a stated level, approximating the cost of a representative firm. This effectively imposes a minimum resale-price maintenance policy common to all stores.

The impact of this type of regulation is not well understood by economists and policy makers. The traditional textbook evaluation of price floors is concerned with short-run market distortions associated with excess supply. While this analysis may be correct in the short run, it ignores the long-run equilibrium effects of price controls on the composition of industries. The central objective of this paper is to demonstrate that price-floor regulations can have an important long-run impact on the structure of markets. We argue in particular that these policies can crowd markets with smaller/less efficient retailers and create an endogenous barrier to entry for low-cost retailers. Taken together, these factors may indirectly lower productivity and affect prices.

To describe the mechanism that drives the long-run impact of the price floor on market structure we construct a model of entry and price competition. The model shows that a price-floor regulation has two opposite effects. First, such a policy causes excess entry into and crowding of markets by raising the expected profit of being active. Second, by protecting small firms the policy blocks the entry of more efficient low marginal cost retailers who face larger fixed costs. Together these forces lead to lower average volume per station in markets subject to a price floor. This is due both to the fact that a greater fraction of small technology stations leads to lower volume per station, and to the fact that increased competition reduces the market share of each station. Furthermore, these forces can lead to higher or lower prices depending on the relative efficiency of firms and the level of the floor.

We analyze this question empirically by studying a specific below-cost price regulation instituted in 1997 in the retail gasoline market in the Canadian province of Québec. For our analysis we have constructed a rich data set at the gasoline-station level featuring close to 1600 stations observed between 1991 and 2001 in five cities in the province of Québec, and nine cities in three other Canadian provinces where the regulation was not implemented. The data, containing detailed information on individual stations’ sales volume, posted price, exact location and characteristics, allow us to study the effect of the floor on station behavior at the local-market level.

We perform a detailed econometric analysis comparing the behavior of local markets and sta-

²Throughout the paper we will refer interchangeably to these types of policies as: price controls, price floors, sales-below-cost laws, below-cost-sales laws, and unfair sales acts. All refer to legislation that limits the prices firms can set either in a particular industry, or broadly across all products.

tions in regulated and unregulated regions. The results are in line with the model predictions. We show that the policy slowed down industry reorganization, discouraging large stations from entering and allowing more smaller stations to survive. In Québec, following the policy, stations became more homogeneous in terms of amenities and sizes, and less spatially differentiated. Importantly, we find that these changes in the structure of local markets significantly reduced per-station sales volumes, a common measure of productivity in this industry. We also show that the regulation did not lead to higher prices in the long-run. This suggests that the upward pressure of the policy on prices, caused by a reduction in efficiency and the floor constraint itself, was at least offset by the long-run increase in competition that we document.

While our focus is on the effect of the policy on market structure, most of the existing empirical analysis evaluating the effect of gasoline market regulations has focused on the impact on prices. Fenili and Lane (1985), Anderson and Johnson (1999) and Johnson (1999) evaluate the effect of sales-below-cost laws in retail gasoline markets in the U.S. They find that jurisdictions with sales-below-cost laws have higher prices and/or margins than those without. These studies are all cross-sectional and, therefore, cannot account for the unobserved heterogeneity across jurisdictions. Their conclusions, linking sales-below-cost laws with higher prices may therefore be spurious. Not all of the prior empirical literature on below-cost regulations concludes that they lead to higher prices. A recent study by Skidmore, Peltier, and Alm (2005) finds that prices tend to fall after the adoption of sales-below-cost laws in US gasoline markets. They argue informally that such regulations could affect market structure, and then use a monthly panel of state-level prices for thirty states, over a twenty year period and investigate empirically what happens to the number of stations. Closest to our analysis, Johnson and Romeo (2000) study the effect of a related policy aimed at protecting independent retailers: a ban on self-service gasoline stations in New Jersey and Oregon. They find that the bans slowed the penetration of convenience store tie-ins, but do not seem to achieve their objective of protecting smaller stations. They also find that the ban led to higher prices.

Our paper is also linked to a large literature that studies the effect of different forms of government intervention on market structure and prices. Another form of government intervention is the imposition of environmental regulations, and studies such as Brown, Hastings, Mansur, and Villas-Boas (2008), Ryan (2006) and Busse and Keohane (2009) have looked at the effect of gasoline content regulation and of the Clean Air Act on market structure. There are also studies evaluating the impact of advertising restrictions on competition and prices in various industries (Milyo and Waldfogel (1999) study the effect of a ban on price advertising, while Clark (2007) looks at the effect of a ban on advertising directed at children on competition in the cereal market and Tan (2006) considers advertising restrictions in cigarette markets). Theoretical work by Armstrong, Vickers, and Zhou (2009) points out that in markets with costly information acquisition, regulations designed to protect consumers, such as price caps or measures which enable consumers to refuse to

receive advertising, could have the unintended consequence reducing consumers' incentives to become informed, resulting in softened price competition.

The rest of the paper is organized as follows. In the next section we describe the minimum price regulation in the Québec gasoline market and relate it to other below-cost price regulations. In the Section 3 we present the model that characterizes the distortions caused by minimum price regulations. In Section 4 we describe our data and present the empirical analysis. Finally, Section 5 concludes.

2 Price floor regulations

A key feature of the retail gasoline industry is the reorganization that took place in the 1990's. During this period all North American markets underwent a major reorganization, characterized by an increase in the size and automation of stations, and the exit of nearly 30% of stations active at the beginning of the decade (see for instance Eckert and West (2005)).³ These changes took place through the entry and reconfiguration of larger and more efficient stations, and the exit of smaller stations. They were caused by technological innovations common to most retailing sectors which increased the efficient size of stores (e.g. automation of the service, better inventory control systems, etc.), by changes in the value of certain amenities (e.g. decreased use of small repair garages), and by changes in regulations regarding the environmental safety of underground storage tanks (see Yin, Kunreuther, and White (2007) and Eckert and Eckert (2008)).⁴ Most regulatory changes and technological innovations took place in the early 1990s, and were slowly implemented across Canadian markets over the decade.

These changes to retail gasoline markets were accompanied by an increase in the frequency of price wars due to excess capacity. In light of these events, the Québec government decided to regulate gasoline prices to protect small independent retailers. In the lead-up to the enactment of the law, Québec's association of independent gasoline retailers conducted a very effective lobbying operation aimed at convincing the provincial government that the exit of independent retailers would adversely affect consumers in the long-run. Several consumer protection groups

³In the US, the number of retail outlets selling gasoline declined from approximately 202 thousand to 171 thousand from 1994 to 2001. Yin, Kunreuther, and White (2007) describe the reorganization of several US gasoline markets throughout the 1990s.

⁴The Environmental Code of Practice for Underground Storage Tank Systems Containing Petroleum Products was published in Canada in 1988 and provided guidance on appropriate upgrading and removal behavior for storage tanks. It was left up to individual provinces whether they adopted these guidelines or established their own regulation. In both Québec and Ontario, regulation came into effect around 1991 regarding approval of unprotected tanks. In terms of timing these restrictions seem to be very similar. In Québec all tanks not meeting the protection standards were to be removed within two to seven years, depending on the age of the tank as of July 1991. In Ontario, no approval was to be given for unprotected tanks that had not been upgraded and they would have to be removed by 1997. Given the similarity in terms of timing, the only remaining concern would be with regard to the extent to which these regulations were enforced in the two provinces. We have found no evidence that there were any differences in enforcement across both regions, but we will nevertheless focus on the years before 1993 and after 1999 and avoid drawing conclusions from the behavior of the markets during the interim period when the policy was implemented.

also supported the policy at the time. The objective of the policy is to protect small retailers against aggressive pricing strategies from large chains and big-box retailers, such as Walmart or Costco.⁵

The *law on petroleum products* was enacted at the beginning of 1997 and administrated by the Régie de l'énergie du Québec (hereafter the "Board"). This followed the occurrence of an important price war during the summer of 1996, which was deemed by the Board to be the result of predatory pricing behavior by the major retailing chains. However, it is not actually clear that predation was actually the cause. Instead, the price wars were likely triggered by excess capacity, and by the decision of Québec's largest retailer (Ultramar) to implement a low-price guarantee (LPG) policy for all of its affiliated stations at the beginning of the summer 1996. Indeed, an investigation by the Canadian Competition Bureau never uncovered predatory behavior.

The mandate of the Board is threefold:

1. Monitor the gasoline industry, and gather information on prices.
2. Determine a weekly floor price or Minimum Estimated Price (MEP).
3. Prevent the occurrence of price wars by imposing a minimum margin regulation in a designated geographic market.

The determination of the MEP is given by the following simple rule which measures the average marginal cost of selling gasoline in each local market:

$$MEP_{mt} = w_t + \tau_{mt} + T_{mt},$$

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. The MEP is calculated and posted on the website of the Board every Monday.

The role of the MEP is to set a price floor under which a firm can sue its competitor(s) for financial compensation on the basis of "excessive and unreasonable commercial practices." This new feature of the law facilitates suing procedures between competitors in the market, in a fashion similar to anti-dumping laws.

In cases where companies repeatedly fail to respect the MEP, the regulation provides the Board with the ability to impose an additional minimum margin to the MEP. It allows the Board to add \$0.03 per liter to the calculation of the MEP in a specific region after the occurrence of a period of sufficiently low prices.

The minimum margin serves two purposes. First, it establishes an implicit (or long run) price floor, under which the Board considers that stations are not covering their fixed operating costs. Second, it enables the Board to indirectly compensate stations after a price war.

⁵Other provinces have also considered implementing similar policies. During the same period that the adoption of the floor was being debated in Québec, a similar regulation was discussed in Ontario, but ultimately was rejected. Nova-Scotia, New-Brunswick and Prince Edward Island all adopted similar regulations after 2001.

The mechanics of the policy are roughly as follows. First, after the occurrence of a long enough low-price period, a gasoline retailer can ask the Board to investigate evidence of price anomalies. The Board then conducts a formal consultation of different groups (retailers, consumer protection groups), in order to evaluate the credibility of the predatory accusation. If the Board is convinced of the accusation, it can add \$0.03 per liter to the calculation of the MEP for a certain period of time in a specific geographic zone where the price war occurred. In practice the Board considers that a price is predatory if the margin (price minus the MEP) is below \$0.03 per liter for a month or more.

This minimum margin approximates the average operating cost of a representative station in the province.⁶ The geographic zone typically includes all local markets which suffered from the price war. Similarly, the length of time for which the minimum margin is applicable is proportional to the length of time of the price war.

The minimum margin has been put in effect three times in two different markets, St-Jérôme and Québec city. In St-Jérôme (North of Montréal), it was added to the MEP from April of 2002 to February of 2003, and again from December of 2003 to June of 2005. The imposition of this price floor followed the entry of Costco in St-Jérôme in 2000, which drove the market price to the MEP level for more than a year. In Québec city it was added to the MEP from July to October of 2001. Its imposition followed a severe price war in the Québec City metropolitan area, during the fall of 2000.

Finally, our experience from studying the Québec gasoline regulation is that the price floor is rarely observed to be binding, although the minimum margin has been put into effect on a few occasions. Figure 1 presents the evolution of weekly average price in Québec city, along with the price floor and the average retail margin.⁷ The dots identify weeks when the average market price was equal to or smaller than the price floor, and the two vertical lines indicate the imposition of the additional minimum margin. Over the period studied, the floor is thus binding less than 10% of the time.

2.1 Below-cost price regulations elsewhere

Currently, twenty-four states in the US have general sales-below-costs laws. In Europe, France recently strengthened a below-cost price regulation applied to all retail markets through the passage of the *Galland* law in 1997. Biscourp, Boutin, and Vergé (2008) look at the impact of the Galland law in France, focusing on the consequences of limiting intra-brand competition on prices. They find that the 1997 reform led to higher prices and softened competition from large grocery store chains.

⁶After public consultations, the Board decided that the representative station is a self-service station operating a convenience store and having annual sales of 350 million liters.

⁷We calculate the floor as the rack price (including taxes) plus transportation cost of 0.34, which reflects its 2001 value.




Figure 1: Evolution of average prices, margins and floor in Québec City between 1991 and 2001

A number of other states and countries have laws for specific industries or products. The most common are sales-below-costs restrictions in the retail gasoline market, but other markets feature similar restrictions.⁸ For instance, in Tennessee there are price floors in the markets for cigarettes, milk, and frozen desserts. In Ireland, below-invoice sales were banned in the retail grocery industry until 2005. More generally, similar policies have been enacted in other contexts: wages in most labor markets are subject to explicit floors as are prices in many agricultural markets, anti-dumping policies forbid foreign firms from setting price below average variable costs, and the entry of big-box retailers is often restricted. In each case, the policies are designed to protect particular groups of firms.

The debate over whether to adopt or overturn sales-below-cost restrictions is ongoing in many jurisdictions. The advocates of these policies typically associate aggressive pricing with predatory or loss-leading behavior. On the other hand, detractors argue that they protect inefficient firms and lead to higher prices and, more generally, to welfare loss – arguments that are consistent with the short-run distortions predicted by the textbook evaluation of price floors. Antitrust authorities typically view such legislation as unnecessary, and they point out that state governments may be too easily convinced by accusation of predation made by various interest groups. When asked to evaluate the merit of proposed below-cost sales legislation in Virginia and North Carolina in 2002 and 2003 respectively, the Federal Trade Commission argued that anticompetitive below-cost pricing

⁸Currently nine U.S. states and three Canadian provinces have sales-below-cost laws in the retail gasoline market.

ing rarely occurs, and that such legislation could harm consumers.⁹ Similarly, in February 2009 following a lawsuit brought by gasoline retailer *Flying J*, the Wisconsin Supreme Court ruled as unconstitutional a local statute which guaranteed a 9.18% markup over the average posted terminal price for gasoline retailers.¹⁰ In Canada, the Competition Bureau has stated that regulation of this type results in higher average prices, and that it does not provide for the highest quality products and the most efficient production, relative to competitive markets.¹¹

3 Theoretical model

We construct a model of entry and price competition, and evaluate the impact of a price floor on the structure of the market. Specifically, we study the effect of the policy on market structure as the industry evolves from a state with only “traditional” gasoline stations, to a state with both traditional stations and a new set of potential entrants: large capacity stations with automated self-service technologies. Since we think of the governmental intervention as occurring *after* the technology innovations took place, we focus on comparing the “post reorganization” state of the market with and without a price floor.

To capture the features of the retail gasoline industry we consider a market in which two types of horizontally differentiated firms compete in prices. The first type of firm only has access to the traditional “small” technology, characterized by a cost function $c_s(q) = c_s q + F_s$. The second type has access to both a “large” and a “small” technology. The large technology is characterized by a low constant marginal cost and high fixed cost (i.e. $c_l < c_s$ and $F_l > F_s$). The assumption that large volume stores represent the more efficient technology is supported by the data, since the marginal cost of selling gasoline is decreasing in the number of pumps (see Houde (2012)).

This results in a simultaneous game of complete information, in which the sets of available actions for both types of firms are summarized by $\{A_1, A_2\} = \{(0, s), (0, s, l)\}$, where 0 , s , and l are out, small, and large respectively. There are six possible market structures denoted by ω : $(0, 0)$, $(0, s)$, $(0, l)$, $(s, 0)$, (s, s) , and (s, l) . The first element of ω refers to the choice of the type 1 firm, and the second to the choice of the type 2 firm. In the appendix we present a numerical example in which we derive the outcomes predicted in the model developed in this section. For completeness, in the example we allow both firms to have access to both technologies.

Pricing in our model is not predatory in the sense that firms are pricing below cost; however, since the l technology has lower marginal cost but offers the same product quality as the s technology, when both types compete in equilibrium the most efficient stores will post the lowest prices.

⁹Virginia Senate Bill No. 458, “Below-Cost Sales of Motor Fuel”, <http://www.ftc.gov/be/V020011.shtm>; and North Carolina House Bill 1203 / Senate Bill 787 (proposed amendments to North Carolina’s Motor Fuel Marketing Act), <http://www.ftc.gov/os/2003/05/ncclsenatorclodfelter.pdf>.

¹⁰The statute in question was s.100.30. See http://www.datcp.state.wi.us/trade/business/unfair-comp/unfair_sales.act.jsp.

¹¹See <http://www.cb-bc.gc.ca/eic/site/cb-bc.nsf/eng/00892.html>.

Depending on its fixed costs, the firm with the s technology may therefore prefer to stay out of the market rather than enter and compete against the l technology. If the s -type firm stays out of the market, the entrant is free to act as a monopolist. The purpose of the model is to investigate whether a price floor can prevent this outcome, and lead to a market structure in which the small technology is present and competition is enhanced.

Before firms commit to their technology and entry decisions, a regulator imposes a price floor constraint p_f . We assume that the floor is always set such that $c_s < p_f < p_l^m$, where p_l^m is the monopoly price for the large technology. In other words, we assume that the floor potentially affects the equilibrium pricing game only in oligopoly markets. If $D_j(p_1, p_2|\omega)$ denotes the demand of a firm of type j in market structure ω , the Bertrand-Nash equilibrium is characterized by the following Kuhn-Tucker first-order conditions:

$$D_j(p_1, p_2|\omega) + \frac{\partial D_j(p_1, p_2|\omega)}{\partial p_j}(p_j - c_j) + \lambda_j = 0 \quad (1)$$

$$\lambda_j(p_j - p_f) = 0, \quad (2)$$

for each $j \in \{1, 2\}$ and $\lambda_j \geq 0$. Since when both types compete we have $p_l \leq p_s$, the price floor will generate three possible outcomes: (i) neither price is constrained ($\lambda_2 = \lambda_1 = 0$), (ii) both prices are constrained ($p_s = p_l = p_f$), or (iii) only the large firm is constrained ($\lambda_1 = 0$ and $p_l = p_f$).

We consider two types of long-run distortions affecting the equilibrium market structure. In the first case, the floor binds in all oligopoly markets and induces excessive crowding, relative to the unconstrained situation. In the second case, the price floor distorts the market by blocking the entry of the most efficient firm.

Case 1: Excessive crowding

Consider an example in which the price floor is high and binds even in the (s, s) market structure. Since firms are symmetric, their profits in the unconstrained and constrained cases are given by:¹²

$$\begin{aligned} \pi_j^u(s, s) - F_s &= D_j(p_s^u, p_s^u)(p_s^u - c_s) - F_s \\ \pi_j^c(s, s) - F_s &= D_j(p_f, p_f)(p_f - c_s) - F_s, \quad \forall j = \{1, 2\}. \end{aligned}$$

Notice that unless aggregate demand elasticity is very high, the constrained profits are increasing in p_f even for a reasonably large floor level. This is a reasonable assumption in gasoline retail markets, since store-level elasticity is large, but aggregate demand is inelastic. In this range, the presence of the price floor can increase profits sufficiently to justify staying active with the small technology. In particular, the regulated market will be more crowded after the policy change if F_s

¹²We use superscripts c and u to indicate that the market is regulated and unregulated respectively.

is in the following range:

$$D_j(p_s^u, p_s^u)(p_s^u - c_s) < F_s < D_j(p_f, p_f)(p_f - c_s). \quad (3)$$

As a result, in this case the policy attracts firms to the market. The equilibrium market structure with the floor is (s, s) rather than $(s, 0)$ or $(0, l)$ without. The floor will be binding and so prices will increase or decrease relative to the unconstrained case depending on how high the floor is. A similar result can be found in the real option literature, since a binding price floor provides a form of insurance that allows more firms to enter and survive (see for instance Dixit and Pindyck (1994)).

Case 2: Blockaded entry

There are two ways the policy can block the use of the more efficient technology. As in the first example, the price floor can be set high enough such that it binds in all cases and makes the selection of the large technology less profitable. For instance, consider the case in which (s, l) is an equilibrium in the unregulated market:

$$\pi_1^u(s, l) - F_s > 0, \quad \pi_2^u(s, l) - F_l > 0, \quad \text{and} \quad \pi_2^u(s, l) - F_l > \pi_2^u(s, s) - F_s,$$

where $\pi_2(s, l)$ is the variable profit of the large firm in state (s, l) , and $\pi_1(s, l)$ is the variable profit of type s .

When the price floor binds for both types, the type 1 firm is strictly better off and therefore s is a dominant strategy in the regulated market. However, the type 2 firm might prefer to enter with the small technology since the market is now split in half (i.e both firms charge p_f). In particular, if the fixed-cost F_l is large relative to F_s , it is likely that the type 2 firm's best-response to s is now to opt for the small technology:

$$D_2(p_f, p_f)(p_f - c_l) - F_l < D_2(p_f, p_f)(p_f - c_s) - F_s.$$

In this example, the equilibrium under the price-floor regulation will therefore be (s, s) instead of (s, l) . The policy thus induces efficiency losses and yields higher prices by blockading the entry with the large technology.

When the price floor is set to a lower level, it is possible that the regulation prevents the adoption of the large technology, without raising prices. To see this, consider a situation in which the selection of the large technology by the type 2 firm causes the type 1 firm to stay out in the unregulated market:

$$\pi_1^u(s, l) - F_s < 0 \quad \text{and} \quad \pi_2^u(s, l) - F_l > 0.$$

Assume further that the price-floor is low-enough and binds only for the large technology in state (s, l) . In this case, the regulation acts as a subsidy for the type 1 firm, and reduces the market share of the large-technology firm. It is therefore possible for the regulator to set p_f such that the type 1 firm will revise its decision to stay out of the market:

$$\pi_1^u(s, l) - F_s < 0 \text{ and } \pi_1^c(s, l) - F_s > 0.$$

If this condition is satisfied, the policy prevents the least efficient firm from staying out. This in turn can block the use of the large technology, provided that F_l is large enough. The equilibrium under the price-floor is therefore $(s, 0)$ or (s, s) instead of $(0, l)$. The equilibrium price in the regulated market can therefore be lower than the price in the unregulated market if the resulting market structure is (s, s) and if the price in (s, s) is less than the price in $(0, l)$. In this case the competition enhancing effect of the policy will dominate the inefficiency losses present in the previous example.¹³

Importantly, this last example shows that the policy can distort market structure without actually binding in equilibrium. As long as F_l is large enough, $(0, l)$ will not be an equilibrium and the price will be higher than the floor. That is, (s, s) or $(s, 0)$ will be the resulting equilibrium depending on the parameters. This is important since the price floor in the retail gasoline market in Québec is rarely observed to bind. The policy is unlikely to generate the first blockaded entry case. Most of the time, the floor is set close to the wholesale price and so is not expected to bind unless a low-cost retailer enters a very crowded area. For this reason we would also not expect to observe the result from the excessive crowding case in Québec.

Finally, note that these results are not specific to the simple two-period model presented here. In an online Appendix we show using numerical simulations that the two types of distortions characterized in the simple model can be generated by a dynamic equilibrium model of entry and product choice in which firms are infinitely lived.¹⁴ We use two parametric examples to show that a price floor policy can prevent the exit of smaller firms and block entry of low-cost retailers. In one example the policy clearly tends to raise prices because competition between efficient stores in the unregulated market easily compensates for the larger number of firms in the regulated one. In the other example, the price floor is almost never binding in equilibrium, but successfully keeps the large stores out of the market and lowers prices.

One feature of the policy that is left out of the model is the fact that the floor can be dynamically adjusted to compensate firms for price wars, to allow medium-size stations to recoup their fixed costs. We did not add this feature to the model, but it very likely would exacerbate the third

¹³In a recent paper Asker and Bar-Isaak (2010) argue that minimum resale price maintenance policies can have exclusionary effects by blocking the entry of more efficient manufacturers. Although the set-up is different, the mechanism at work in their model is very similar to ours: minimum resale price maintenance can be used by incumbent manufacturers to increase retailer profits when they deter entry.

¹⁴The Appendix is available at <http://www.jeanfrancoishoude.info/WordPress/>.

prediction, and create further crowding. Therefore, even if the lowest of the two floors is not actually binding in equilibrium, the presence of a minimum margin raises the option value of staying in the industry for small and medium size stations.

To summarize, according to our theory, a price-floor policy can distort the post reorganization equilibrium structure of retail markets in two ways: excessive crowding, and blockaded entry.¹⁵ Two testable implications follow:

H_1 . The fraction of stations with the newer technology is smaller in markets subjected to a price floor.

H_2 . The number of competitors is larger in markets subjected to a price floor.

Together these predictions yield the following corollary:

C_1 . The average volume per station is lower in markets subjected to a price floor.

This third prediction is the result of both selection and composition effects: a greater fraction of “small” technology stations leads to lower volume per station, and increased competition reduces the market share of each station.

The effect of the policy on prices is unclear. Prices can increase in the long run following the introduction of a floor if it is frequently binding or if it keeps firms with higher marginal costs active. On the other hand, if the competition enhancing effect of the policy is important and the marginal cost difference between type l and s stations is small, prices can be lower in the long-run. We test these predictions of the model in the following section of the paper.

4 Empirical analysis

4.1 Estimation strategy

The objective of our empirical analysis is to test the predictions of our theory by studying the impact of Québec’s policy on long-run market structure and stations’ outcomes. Since price-control policies are not put in place randomly, and we do not observe the evolution of a mirror image of the Québec market without the policy, we follow a quasi-experimental approach, and compare local markets in Québec with markets in comparable cities in the rest of Canada where similar policies were not put in place. Furthermore, rather than comparing the post-policy steady-states in treated and control markets, we use a difference-in-difference model to estimate the effect of the policy on the *changes* in the structure of markets and stations’ outcomes. This eliminates part of the endogenous selection problem by controlling for time-invariant market structure characteristics that might have led to the regulation of the the Québec markets.

¹⁵Note that pre reorganization, the floor would have no effect on the equilibrium structure, which would be (s, s) regardless of the floor since the l technology is not available.

Moreover, since we are interested in long-run changes, we study the evolution of control and treated markets over an eleven year period, between 1991 and 2001. During this period, the North American gasoline industry underwent major changes that satisfy the two conditions required for the policy to possibly distort markets: (i) newer stations exhibit larger capacity which yields larger volume discounts, and (ii) newer stations require fewer labor inputs due to automatized services. However, it remains an empirical question whether or not these technological changes were important enough and the price floor policy implemented at the right time and the right level to affect entry/exit and investment decisions.

We consider two main econometric specifications. First, we use the full panel to measure the (potentially) differential evolution of markets and stations in Québec after the implementation of the policy in 1997, relative to other markets in the rest of Canada. Second, we estimate the effect of the policy on the *changes* in market and station characteristics between the two end points of our panel: 1991 and 2001. These two years approximate the pre- and post-policy industry steady-states. In this specification, we control for trends in other relevant variables before and around the time of the implementation of the policy, which could be confounded with the policy implementation.

We label the first specification **full-panel**, since it uses the information contained in the full, yearly transition of markets from one steady-state to the other. We label the second specification **long-difference**, since it focuses solely on the long-run changes in markets. The long-difference specification uses less information than the full-panel specification, but is more in line with the model implications, which refer to the long-run equilibrium outcomes before and after market reorganization.

We estimate the long-difference specifications using both OLS and matching estimators (described in detail below). The policy effect is identified from differences in the long-term trends over the full decade. Note however that like in the full-panel fixed-effect specification, the estimation focuses on changes initiated from 1997 onward since we control for pre-policy changes in local market structure and demographics.

This section has five parts. In section 4.2 we provide a description of the data. In section 4.3 we discuss in detail the treatment and control groups on which our empirical analysis is based. In section 4.4 we provide descriptive evidence of the reorganization of the gasoline industry in Québec and in the rest of Canada consistent with the presence of market distortions. Sections 4.5 and 4.6 contain the center-piece of our analysis in which we use econometric techniques to compare the markets and stations in Québec with control markets and stations, before and after the policy was implemented, controlling for the observed and unobserved factors that may have driven the observed changes. Finally, in section 4.7 we investigate the effect of the policy on prices.

4.2 Description of the data

The gasoline station data used in this study were collected by Kent Marketing, the leading survey company for the Canadian gasoline market. The survey offers accurate measures of sales and station characteristics, since each site is physically visited at the end of each quarter. Volume sold over the quarter is measured by reading the pumps' meters. Station characteristics at that point in time are recorded, as is the price posted on the date of the survey. The panel spans eleven years between 1991 and 2001, and includes all 1601 stations in fourteen selected cities of Québec and three other Canadian provinces. For our analysis we take the sales volume data collected at the end of the third quarter of each year, along with the prices and station characteristics at this date.

The observed station characteristics include the type of convenience store, a car-repair shop indicator, the number and size of the service islands, the opening hours, the type of service, and an indicator for the availability of a car-wash. Brand indicator variables are also added to the set of characteristics to reflect the fact that consumers might view gasoline brands as having different qualities. We also have detailed information about the geographic location of each station in the sample.

The Kent data contain no information on wholesale prices and so in our analysis we use data collected by the market research firm MJ Ervin to measure wholesale prices. MJ Ervin collects prices directly from the wholesalers and we have access to the regular unleaded wholesale (Rack) price. These data are at the company (marketer) level, and for most of the analysis we compute either an average or a minimum wholesale price.¹⁶

Since there is important variation in gasoline prices over time, there is some concern that the annual price data from Kent are inadequate. Therefore, we have also collected weekly retail price data from MJ Ervin (in the same provinces, but for a different set of markets) to complement the Kent data. We will use this supplementary data set in Section 4.6 to show that our price results are robust to the use of these higher frequency data.

In our analysis, markets correspond to metropolitan areas. Nevertheless, since retail gasoline markets are spatially differentiated and stations face competition from a set of local competitors, we are also interested in quantifying the degree of competition more locally. For this we will proceed in two ways. First, we construct a station-specific neighborhood, defined as a two-minute driving distance radius around each active station.

Second, we also construct neighborhood boundaries that define a set of spatially homogeneous locations. To do so we use a clustering algorithm that groups store locations according to two criteria, related to the distance between stores and whether they share a common street. Given these criteria, a neighborhood approximates an intersection or a major street segment.¹⁷ Importantly,

¹⁶Wholesale prices can often include station- or firm-level discounts which we do not observe.

¹⁷ A common way of specifying neighborhoods is to use existing definitions, for instance census-tracts or zip codes. While these definitions typically allow researchers to get accurate measures of population characteristics, their boundaries are arbitrary and do not necessarily reflect competition between stores. Although we use these measures to define

our definition is time-invariant since the set of possible locations is defined as all locations ever active throughout our sample. We describe in greater detail our procedure in the Appendix.¹⁸ The median neighborhood size is three stations per neighborhood in the whole sample, but some cities are clearly more dense than others. For instance, in Hamilton (Ontario) the median neighborhood size is five stations, while Chicoutimi (Québec) neighborhoods have a median size of two stations. Overall, the algorithm constructs neighborhoods that are very comparable across all the regions, since the size distributions are very similar.

Table 11 in the Appendix presents a set of descriptive statistics for some of the key variables used in the analysis, for both the pre- and post-policy periods and in both the cities inside and outside Québec (we will denote those outside Québec as Rest). The table shows, for instance, that: (i) the number large stations with more than four service islands and the average number of pumps both increased by about 20% in Québec and by more than 50% in the rest of the cities, (ii) the number of neighboring competitors decreased by 17% in Québec and by 29% elsewhere. We come back to a description of these trends below.

4.3 Comparability of markets and control variables

The fourteen Canadian cities we study include five cities in Québec, which are our treatment cities. Our control cities, which were not subjected to the policy include seven cities in Ontario, and one city in each of Nova Scotia and Saskatchewan. We will refer to our control cities as “Rest of Canada,” despite the fact that they are located mostly in Ontario. We observe all stations in all cities throughout the years 1991-2001.

We exclude from the sample the biggest metro areas of Toronto and Montreal. As shown in Table 1, the selected metropolitan areas are all comparable in terms of size, population growth, volume of gasoline sold per capita and growth of volume per capita. Furthermore, there is considerable overlap in the set of major players active in each market. Each market has some subset of six chains that are integrated with the refinery sector: Shell, Esso/Imperial Oil, Ultramar, Irving, Sunoco, and Petro-Canada.

As indicated in subsection 4.1, we base our econometric analysis on two different specifications, based on the full panel and the long-run changes between 1991 and 2001, respectively. The two specifications are based on different information and have slightly different sets of regressors. We label the controls in our full-panel specifications $Z_{j,t}$, where j is the cross-sectional component

markets, below in our market-level analysis we will use information at the forward sortation area (FSA) level to construct demographic trends that we match with the markets we construct. The FSA is the first half of a postal code. This unit of aggregation corresponds to about 4 to 6 census-tracts in urban areas (or between 10,000 and 40,000 households), or one small town in more rural areas. The median population size per FSA is about 16,000. There are over 1,300 FSAs in Canada, and over 850,000 postal codes.

¹⁸The size and composition of the neighborhoods is affected by the parameters used in the clustering algorithm. We pick the bandwidth parameter in order to obtain average neighborhood sizes around 3 stations, and avoid having neighborhoods bigger than 15 stations.

Table 1: Market characteristics

Market names	Population	Δ Population (%)	Volume per cap. (litres/day)	Δ Volume per cap. (%)
Cornwall (On)	45726	0.29	2.59	3.57
Guelph (On)	101163	1.67	2.31	-0.91
Hamilton (On)	483981	0.91	2.03	0.02
Kingston (On)	145090	1.10	2.10	1.25
St Catharines (On)	129144	0.41	2.22	3.37
Chicoutimi (Qc)	162410	-0.42	1.84	1.77
Drummondville (Qc)	64241	0.99	2.54	-0.38
Quebec (Qc)	512746	0.33	3.03	1.22
Sherbrooke (Qc)	138957	0.70	2.38	1.40
Trois-Rivieres (Qc)	140847	0.05	2.03	0.72
Halifax (NS)	352548	0.95	1.95	2.04
Saskatoon (Sk)	224014	0.74	2.15	2.08
Brantford (On)	89615	0.79	1.90	-1.33
Windsor (On)	233709	1.49	1.91	3.08

Population and *Volume per capita* are market (i.e. city) averages taken over the period 1991 – 2001. The change variables are averages of year-to-year log-changes taken over the same period and expressed in percentage ($\times 100$).

(market, neighborhood, station) and t is the time period, and use \mathbb{W}_j to label the controls in our long-difference specifications for each cross-sectional component j .

Full-panel control variables

We include in $\mathbb{Z}_{j,t}$ variables that measure: (i) FSA-level demographic characteristics (average income, population size and unemployment rate),¹⁹ (ii) characteristics of the regional upstream markets (i.e. rack price and number of companies at nearest terminal) and (iii) provincial taxes on gasoline.²⁰ In the full-panel specifications we will also control for year and location fixed-effects.

Furthermore, when studying neighborhood- and station-level outcomes, we include indicators for the presence of both Ultramar and Sunoco within a 2-minute driving distance of each station, before and after 1996, to control for two events that occurred near 1997 and affected a fraction of stations in Québec and in the rest of Canada. First, in March 1996, Ultramar and Sunoco, two of the largest firms in the Canadian petrol industry, announced their intentions to swap their service stations in Québec and Ontario.²¹ Second, in a further attempt to increase its share of

¹⁹See footnote 17 for a definition of FSA.

²⁰Tax differences between Québec and the other regions were more important from 1995 on since the Québec government decreased the consumption tax and increased the excise tax on gasoline.

²¹The transaction involved Ultramar acquiring 127 Sunoco stations in Québec, and in exchange Sunoco acquired 88 Ultramar sites in Ontario. At the time, Sunoco did not have a refinery in Québec and chose to concentrate its retail activities in Ontario and Western Canada. Ultramar on the other hand adopted the strategy of increasing its dominance in the Québec market, and distributing a larger fraction of its Saint-Romuald refinery's production (near Québec City) locally.

the Québec retail market, in the summer of 1996 Ultramar instituted a low-price guarantee policy (ValeurPlus or ValuePlus). Since Ultramar has a greater presence in Québec than elsewhere, we might worry that any effects on prices and market structure that we attribute to the price floor policy are actually the result of one of these events.²² The exact list of variables included varies across specifications as a function of the unit of analysis (i.e. stations or local markets). Each results table includes a detailed footnote describing the list of control variables.

Long-difference control variables

We include in \mathbb{W}_j a set of market and station characteristics that measure pre-policy market and station characteristics, as well as changes in local demographic characteristics. More specifically, we control for the 1991 levels, as well as changes between 1991 and 1996 and between 1996 and 2001, in demographics characteristics of each FSA (i.e. employment, population and income). We also control for the presence of Sunoco and Ultramar stations in a 2-minutes neighborhood around each locations. Finally, we include in \mathbb{W}_j the 1991 levels and changes between 1991-1995 in the characteristics of local markets (i.e. local monopoly indicator, average number of pumps and islands, and number of competitors).²³ This last group of variables controls for location-specific pre-policy observable trends in market structure, and ensures that the policy effect is not identified from changes that were initiated before the policy was first announced in 1996. Notice that some variables in $\mathbb{Z}_{j,t}$ are excluded from \mathbb{W}_j because they are collinear (or close to) with the Québec indicator variable (i.e. treatment).²⁴

We use the control variables in \mathbb{W}_j in two ways. First, we use them as explanatory variables in our long-difference OLS regressions. Second, in order to control for systematic observable differences between control and treated markets we also use \mathbb{W}_j to estimate the probability that a station or local market is located in the treated region, and estimate the policy effect via a difference-in-

²²We are not aware of other events of similar importance to the Ultramar and Sunoco episodes that would have affected differently the treatment or control provinces during our sample period. This is not to say that the markets are completely identical, but rather that there were no other important events, such as mergers, that affected the wholesale or retail gasoline markets in our sample markets during our sample period. Among the other observed differences, the Atlantic Provinces in Canada experimented with various forms of price controls in the past, but the only city in our data from this region is Halifax in Nova Scotia where prices were regulated (maximum retail and wholesale price) until mid 1991 by the Public Utilities Commission. Our results are robust to omitting 1991, and also to omitting Halifax altogether. Also, as mentioned above, the impetus to reconfigure stations came in part from environmental regulations surrounding underground storage tanks (UST). Starting in 1991, the Canadian government imposed a common set of restrictions on the type of UST that gasoline stations must use in order to reduce leakage hazards. As pointed out by Eckert and Eckert (2008), provinces implemented this policy by setting up slightly different time-lines for the renovation of non-compliant stations. Québec and Ontario, our main control province, imposed similar implementation deadlines of 1998 and 1997 respectively. The last relevant deadlines in Nova Scotia and Saskatchewan were later: in 2003 and 2005 respectively. In the empirical analysis, we present the results of a long-difference analysis, in which we compare changes between 1991 and 2001. This helps us avoid drawing conclusions from what happened in the markets around 1997-1998.

²³We choose 1995 as the end date for these changes since this allows for the possibility that the arrival of the policy was anticipated starting in 1996.

²⁴This procedure excludes taxes and rack prices, since both variables do not vary sufficiently within province.

difference matching estimator. In particular, for each station or local market j , we use a Logit model to estimate the probability of belonging to the treatment group (i.e. Québec):

$$\Pr[\text{QC}_j = 1 | \mathbb{W}_j] = \Pr(\mathbb{W}_j), \quad (4)$$

where QC_j is a dummy equal to one for markets or stations located in Québec.²⁵

Comparability of treated and control locations

In Table 2 we show the validity of the matching approach, by measuring the treatment/control differences in the pre-policy outcome variables: station-level sales volume and markup, neighborhood-level market structure (monopoly indicator, number of competitors, average number of pumps, number of service islands), and local-cluster market structure (minimum distance, monopoly indicator, average number of pumps). We also compare FSA demographic characteristics across treatment and control groups. For each variable, we measure differences in the 1991 levels, and in the 1991-1995 changes. Columns (1) and (2) present unweighted averages of these differences. Columns (3) and (4) report weighted average differences obtained from matching regressions in which stations and markets are matched according to the estimated propensity scores. Each entry corresponds to a nearest neighbor matching regression coefficient, testing the null hypothesis that the difference between control and treatment groups is zero.²⁶

The results reported in the unweighted columns illustrate that the two groups exhibit significant differences in initial market structure. In Québec, local markets exhibited more stations in 1991, and stations were smaller on average (both in terms of volume and pumps). We also observe a few statistically significant differences in the trends before the policy intervention: stations in Québec grew at a slower rate, and increased their pump capacity at a faster rate. Most of these differences shrink in size and become statistically insignificant when controlling for the propensity score defined in equation 4. This means that the estimated propensity scores account for virtually all the pre-policy differences between treated and untreated markets and stations. The only outcome to exhibit a significant (but small) difference is the average markup level and change. This is likely due to the spatial correlation that exists in prices within regions, which biases downwards the standard errors.

4.4 Industry reorganization and the policy

The reorganization of the industry was evident in both treated and untreated markets. The number of stations in the selected markets decreased by about 35% across the sample. Over the eleven

²⁵The matching estimation procedure is described in detail in footnote 27.

²⁶We use the four nearest neighbors to compute this average (i.e $N = 4$). The properties of the nearest neighbor matching estimator are described by Abadie and Imbens (2006) and are incorporated into the STATA command *nnmatch*

Table 2: Differences between treated and untreated markets and stations

VARIABLES	Mean differences: QC - ROC		Weighted differences: QC - ROC	
	(1) Levels ('91)	(2) Changes ('95-'91)	(3) Levels ('91)	(4) Changes ('95-'91)
(1) Station volume	-2,246.462 ^a	-910.122 ^a	-171.179	-130.635
(2) Station markup	0.128 ^a	-0.088 ^a	0.038 ^a	-0.039 ^a
Local nbh. (200 meters):				
(3) Min. Distance	-0.067 ^b	-0.014	-0.026	-0.012
(4) Mono.	-0.045 ^b	-0.047 ^a	0.010	-0.024
(5) Avg. Pumps	-0.992 ^a	0.500 ^a	0.134	0.004
Cluster locations:				
(6) Mono.	-0.013	-0.045	0.081	0.033
(7) Nb. Comp.	0.260	0.159 ^c	-0.293	0.133
(8) Nb. Pumps	-1.634 ^a	0.374 ^c	1.663	-0.143
(9) Islands	-0.294 ^a	0.036	0.082	0.029
FSA Demographics:				
(10) Population	-3,882.280 ^a	535.186	-598.858	242.406
(11) Income	-4,550.704 ^a	-3,635.918 ^a	90.781	-857.514
(12) UR	0.803 ^a	1.571 ^a	-0.106	0.370

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Columns (1) and (2) show the differences in the mean level and mean 1995-1991 change for each variable between the markets and stations in Québec and the Rest of Canada. Columns (3) and (4) show the differences in the mean level and mean 1995-1991 change for each variable between the markets and stations in Québec and the Rest of Canada, conditional on the estimated propensity score.

Rows (1)-(5) correspond to station level variables: (1) daily sales volume (in liters), (2) price markup per liter (in cents), (3) driving distance to closest station (in minutes), (4) no competitor within a 2-minute driving distance, (5) average number of pumps of stations within a 2-minute driving distance. Rows (6)-(12) correspond to neighborhood-level variables, where each neighborhood corresponds to one cluster as explained in the Appendix as follows: (6) monopoly indicator, (7) number of competitors, (8) average number of pumps, (9) average number of islands, (10) postal code-level population, (11) postal code-level average income and (12) postal code-level unemployment rate.

The conditional means were obtained from a matching regression of each variable on the estimated propensity scores. The conditional means in rows (1)-(5) are obtained from station level regressions, whereas conditional means in rows (6)-(12) are obtained from neighborhood-level regressions. For each matching regression standard errors are robust to heteroskedasticity, and we use the bias correction method of Imbens et al. (Stata Journal 2004).

The propensity scores for the station-level regressions in rows (1)-(5) are estimated using a linear Logit specification with the following controls: postal-code area demographics (population, income and unemployment rate) in 1991, 1996 and 2001 and changes between 1991-1996 and 1996-2001; neighborhood-level market structure in 1991 (monopoly dummy, number of competitors, average number of pumps and islands); indicator variables for Ultramar and Sunoco within 2 minutes driving distances in 1991; changes for all the station-level outcomes between 1991 and 1995.

The propensity scores for the neighborhood-level regressions in rows (6)-(12) are estimated using a linear Logit specification with the following controls: postal-code area demographics (population, income and unemployment rate) in 1991, 1996 and 2001 and changes between 1991-1996 and 1996-2001; neighborhood-level market structure in 1991 (monopoly dummy, number of competitors, average number of pumps and islands); Ultramar and Sunoco presence in 1991; changes in the market structure variables between 1991 and 1995.

years of our panel, we observe a total of 229 new entrants and 670 exits out of 1,601 unique stations

Table 3: Entrant, exiting, and continuing stations

	N	(1) Nb. Pumps	(2) Nb. Islands	(3) Large CS	(4) Self-service
Mean differences in the final year:					
Entrant - Incumbent	1,125	1.234 (0.992)	0.1215 (0.138)	0.149 ^a (0.0466)	0.052 (0.0472)
Exiting - Entrant	713	-7.284 ^a (0.958)	-0.907 ^a (0.135)	-0.466 ^a (0.0461)	-0.327 ^a (0.0478)
Exiting - Incumbent	1,588	-6.050 ^a (0.333)	-0.782 ^a (0.0609)	-0.317 ^a (0.0211)	-0.275 ^a (0.0235)
Policy effects by category:					
New entrants	147	-9.161 ^a (3.035)	-1.162 ^b (0.459)	-0.462 ^b (0.183)	-0.221 (0.159)
Exiting	588	0.0974 (0.854)	0.140 (0.190)	-0.0656 (0.0730)	0.0493 (0.0837)
Reconfiguring	456	-4.496 ^b (1.796)	-0.445 (0.278)	-0.100 (0.0961)	-0.112 (0.0957)
Robust standard errors in parentheses ^a p<0.01, ^b p<0.05, ^c p<0.1					

The policy effects by category are estimated with difference-in-difference OLS regressions of each variable among the stations in each category, across regions before and after the policy. Reconfiguring stations are stations that show a change in the number of pumps larger than 4 at any point throughout the sample. In the top panel, entrants are defined as stations that enter and never exit.

in fourteen cities. The large number of exits relative to entrants is easily explained by the fact that the new “technology” corresponds to a larger capacity and requires more expensive equipment.

The upper panel in Table 3 illustrates the reorganization of all markets by comparing the characteristics of new and exiting stations with each other and with the rest of stations. The first row clearly shows that entrants and continuing firms have more or less the same size, but that entrants are slightly more likely to have a convenience store and less likely to offer full service. Exiting firms, however, are significantly smaller: on average, stations that exited before 2001 had 6 or 7 fewer pumps and less service-islands than entrants and continuing firms. Similarly, exiting firms were much more likely to offer full-service and not have a convenience store attached to it. The proportion of stations with a convenience store, and the proportion of self-service stations have each increased by more than 20% in all regions.

Although all of the markets that we study experienced similar aggregate reorganization trends, the rate at which these changes occurred differed in Québec and the other provinces. To illustrate the impact of the policy on these changes, we estimate a series of difference-in-difference regressions on three subsamples: entrants (entry date between 1992-2001), exiting (last year of activity

between 1991 and 2000), and reconfiguring stations (discrete change in number of pumps between 1992-2001). For each station in these three groups, we measure the amenities and size at the time of the change. That is, we measure its characteristics on the first year after entry occurs or after a reconfiguration, or its characteristics the last year before exit occurred. When a station leaves the market for a year before re-entering after renovations are completed, we code it as a reconfiguration, rather than an exit followed by a new entry. We then estimate the following linear regression:

$$Y_j = \beta \text{Policy}_j + \mu_t^{\text{year}} + \mu_m^{\text{market}} + u_j, \quad (5)$$

where $\text{Policy}_j = \text{QC}_j \times 1(E_j \geq 1997)$, E_j corresponds to the date of the three dynamic events (i.e. entry year, exit year, reconfiguring year), Y_j labels one of four outcomes (number of pumps, number of islands, large convenience-store indicator, and self-service indicator), and μ_t^{year} and μ_m^{market} are year and market fixed effects.

The lower panel in Table 3 shows estimates of the policy on two measures of size and on amenities. The results show that entering stations were significantly smaller in Québec in terms of number of pumps and number of service islands, and were less likely to open a store with a large convenience store. We estimate somewhat smaller but significant effects on the reconfiguration decisions of incumbent stations. However, we fail to find any significant differences in the size or amenities of exiting stations. These results suggest that the policy affected the reorganization of the industry in Québec by reducing the incentive of large scale stations to enter the market, and to a lesser extent by reducing the importance of incumbents' investments in capacity. We come back to this discussion in Section 4.6 when we study the impact of the policy on station outcomes.

4.5 Evidence of the impact of the policy on market structure

In this section, we present an econometric analysis of the change in market structure in Québec and elsewhere before and after the implementation of the policy, controlling for observed and unobserved variables that might be correlated with it. This section directly tests predictions H_1 and H_2 .

Our analysis focuses on measures of competition, spatial differentiation, station capacity and amenities. The outcome variables that we study differ based on the unit of analysis: (i) market (cities as defined by the MSA metro areas), (ii) neighborhood (or clusters), and (iii) stations. At the city level, we use two measures of concentration (HHI-index and number of stations), a measure of the average minimum distance between stations, and two measures of average station sizes (i.e. fraction of stations with more than 2 service islands). At the neighborhood (or cluster) level, we construct measures of competition (e.g. number of competitors, and local monopoly indicator), and calculate the average size and the fraction of active stations with certain amenities (e.g.

Table 4: Market-level market-structure regressions

VARIABLES	(1) Distance min.	(2) HHI	(3) Stations per cap. (log)	(4) % Large stations
A: Full-panel (ols)				
Policy	-0.037 ^b (0.02)	-130.785 ^a (34.56)	0.089 ^b (0.04)	-0.060 ^b (0.02)
Observations	154	154	154	154
B: Long-difference (ols)				
Policy	-0.049 (0.03)	-247.345 (138.52)	0.161 ^b (0.06)	-0.099 ^b (0.03)
Observations	14	14	14	14

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the market level. Specifications in panel **A** include the following controls: city population, city unemployment rate, number of companies at nearest terminal, rack price, fraction of locations with Ultramar and Sunoco in 2 minutes neighborhoods (plus interaction with “after 1996” indicator), year and market fixed effects. Panel **B** includes the following controls: change in city population, change in unemployment rate, change in region-specific rack price, change in number of upstream companies at nearest terminal. The Policy coefficients correspond to the long-run effect of the price-floor: Year \geq 1997 \times Quebec.

more than two islands, convenience store, electronic payment). Finally, at the station level we measure the distance between the nearest neighbor, the average size of stations located within a two minutes driving buffer zone, and an indicator variable for local monopolies (also within a two minutes driving distance). Notice that we define neighborhoods and choose the distance radius in a way that would approximate “homogenous” locations, and therefore measures of competition at this level also evaluate the extent of spatial differentiation.

Market-level regressions

We start by estimating the effect of the policy on the structure of markets at the market-level. We use two types of samples: (i) the full panel between 1991 and 2001, and (ii) the long-difference between 1991 and 2001. This leads to two difference-in-difference models for each outcome Y :

$$Y_{j,t} = \beta_F \text{Policy}_{j,t} + \mathbb{Z}_{j,t} \gamma_F + \mu_j^{\text{market}} + \mu_t^{\text{year}} + e_{j,t}, \quad (6)$$

$$\Delta Y_j = \beta_D \text{QC}_j + \mathbb{W}_j \gamma_D + u_j, \quad (7)$$

where j indexes a market (i.e. MSA), $Y_{j,t}$ is the variable of interest and $\Delta Y_j = Y_{j,2001} - Y_{j,1991}$ is its change between 1991 and 2001. The term $\text{Policy}_{j,t} = \text{QC}_j \times 1(t \geq 1997)$ is a policy indicator.

Table 4 presents the estimation results for the policy effects β_F and β_D for the full and long-difference specifications respectively. Panel A corresponds to the full-panel specification, and panel B corresponds to the long-difference specification. The estimates support the two theory predictions. We find that regulated cities became relatively less concentrated and more crowded

after 1997. The effects of the policy on the HHI and the number of stations per capita are sizeable. We estimate that the policy led to a 130 point decrease in the level of the HHI, and an 8.9% increase in the number of stations per capita (16.1% in panel B). Note that, on average, the market-level HHI was 301 in Québec after the introduction of the policy. The fraction of stations with more than two service islands (i.e. large stations) also decreased significantly due to the policy. The full-panel specification implies a 6% decline, while the long-difference specification implies a 10% decline.

Compared with panel A, the estimates using the long-difference sample lose some precision due to the fact that we have only 14 observations. The sign and magnitude of the implied policy impact are similar in the two samples, implying that the timing of the policy coincided with the changes estimated via the long-difference specification.

Neighborhood-level regressions

We now examine the effect of the policy at the neighborhood level in order to account for the fact that competition in retail gasoline markets is localized. Recall that neighborhoods are constructed using a clustering procedure that calculates non-overlapping areas with homogenous locations.

Before turning to the regression results, we first present the evolution of two neighborhood characteristics over the sample period. Figure 2 illustrates the evolution of the fraction of local monopolies and the average number of pumps in the Québec and control markets neighborhoods. The initial levels of both variables are different across the treatment and control groups: Québec neighborhoods initially had fewer local monopolies, and stations with fewer pumps. See the summary statistics in Table 11. To control for these initial differences, we express each variable in difference relative to their respective 1991 levels. The figures therefore show the change in the fraction of local monopolies and average station size relative to 1991.

The vertical line in each panel indicates the moment at which the policy was introduced. Both figures show that the reorganization of the industry during the 1990's was less pronounced in Québec than in the rest of Canada. Graph (a) shows that the gap between Québec and the control markets has been widening over time. Similarly, graph (b) shows that the difference in the average size of stations has been increasing after the imposition of the policy in 1997. Therefore, while the prevalence of large stations increased everywhere, the increase was substantially larger outside Québec. In both cases, the figures highlight that the implementation of the policy coincided with changes in the trends in Québec relative to control markets. In the control markets after 1997, firms entered and invested in larger-scale stations at a faster rate, consistently with the end of the underground-storage tank regulation time-line. In contrast, we observe a slowdown in investment and exit in the Québec neighborhoods, as illustrated by a slower increase in the fraction of local monopolies and average station size.

To formally quantify the effect of the policy, we estimate a set of regression models controlling for a rich set of observed covariates. As above, we consider two sub-samples: (i) the full panel

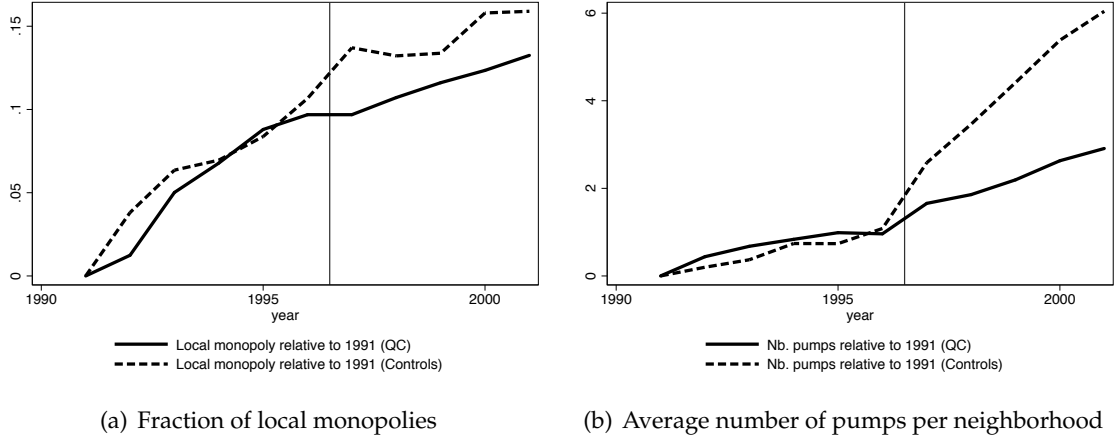


Figure 2: Evolution of neighborhood market structure characteristics control markets between 1991 and 2001

between 1991 and 2001, and (ii) the long-difference between 1991 and 2001. This leads to two difference-in-difference models for each outcome Y :

$$Y_{j,t} = \beta_F \text{Policy}_{j,t} + \mathbb{Z}_{j,t} \gamma_F + \mu_j^{\text{nbh}} + \mu_t^{\text{year}} + e_{j,t}, \quad (8)$$

$$\Delta Y_j = \beta_D \text{QC}_j + \mathbb{W}_j \gamma_D + u_j, \quad (9)$$

where $Y_{j,t}$ is the variable of interest in neighborhood j in period t , and $\Delta Y_j = Y_{j,2001} - Y_{j,1991}$ is its change between 1991 and 2001. The term $\text{Policy}_{j,t} = \text{QC}_j \times 1(t \geq 1997)$ is a policy indicator, μ_j^{nbh} and μ_t^{year} are neighborhood and year fixed-effects.

We also estimate the long-difference specification using a nearest neighbor matching estimator. We measure the distance between observations via an estimated propensity score defined in equation 4. The propensity score measures the probability that a neighborhood j was subject to the minimum price policy, conditional on the same controls variables \mathbb{W}_j .²⁷

²⁷The estimation proceeds as follows: Denote by \mathcal{Q} and \mathcal{Q}^- the set of treated and untreated markets respectively, and let $N_{\mathcal{Q}}$ and $N_{\mathcal{Q}^-}$ be the number of observations in each group. The estimated average treatment effect $\hat{\beta}_D$ is defined as the average difference between the observed long-difference in outcome Y in the treatment group, and the counter-factual long-difference estimated using observed changes in “similar” markets from the control group (rest of Canada):

$$\hat{\beta}_D = \frac{1}{N_{\mathcal{Q}}} \sum_{j \in \mathcal{Q}} \left(\Delta Y_j - \sum_{i=1}^{N_{\mathcal{Q}^-}} w_{i,j} \Delta Y_i \right), \quad (10)$$

where each treated observation $j \in \mathcal{Q}$ is matched with observations in the control group, each of which is weighted according to $w_{i,j}$. The nearest neighbor weights are equal to $1/N$ for the N observations in the control group that are “closest” to observation j in the treatment group in terms of their propensity scores, and zero otherwise. In the application, we use the four nearest neighbors to compute this average (i.e $N = 4$). For a detailed description of difference-in-differences matching estimators see (Heckman, Ichimura, and Todd 1997). The properties of the nearest neighbor matching estimator have been described by (Abadie and Imbens 2006) and are incorporated into the STATA

Table 5: Local market-structure and amenities regressions

VARIABLES	(1) Competitors	(2) Monopoly	(3) Pumps	(4) Islands	(5) No CS	(6) Card
A: Full-panel with FE (ols)						
Policy	0.111 ^c (0.06)	-0.069 ^b (0.03)	-0.940 ^b (0.42)	-0.126 ^b (0.06)	0.055 ^b (0.03)	-0.033 ^c (0.02)
Observations	5,305	5,305	5,305	5,305	5,305	5,305
B: Long difference (ols)						
Policy	0.008 (0.07)	-0.077 ^c (0.04)	-1.986 ^a (0.67)	-0.180 ^c (0.09)	0.030 (0.04)	-0.034 (0.04)
Observations	437	437	437	437	437	437
C: Long difference (match)						
Policy	0.182 (0.16)	-0.093 (0.08)	-2.628 ^a (0.82)	-0.112 (0.15)	0.038 (0.07)	0.020 (0.07)
Observations	337	337	337	337	337	337

Cross-sectional units: Non-overlapping neighborhood clusters.

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the neighborhood level (except for panel C). The sample includes “active” neighborhoods with at least one competitor. Each column represents a different regression: (1) number of competitors, (2) monopoly indicator, (3) average number of pumps, (4) average number of islands, (5) no convenience store, (6) card reader at the pump. Specifications in panel A include the following controls: postal-code area demographics (population, income and unemployment rate), city population and unemployment rate, number of companies at nearest terminal, rack price, indicator variables for Ultramar and Sunoco in 2 minutes neighborhoods (plus interaction with “after 1996” indicator), short-run policy variables ($(1996 \leq \text{Year} \leq 1997) \times \text{Quebec}$, year and neighborhood fixed effects. The Policy coefficients correspond to the long-run effect of the price-floor: $\text{Year} \geq 1997 \times \text{Quebec}$. Panels B and C include the following control variables: postal-code area demographics in 1991, 1996 and 2001 (and changes), 1991 neighborhood structure (monopoly dummy, number of competitors, average number of pumps and islands), and Ultramar and Sunoco presence in 1991. The propensity score is estimated using a linear Logit specification with additional controls for the change in all the outcomes between 1991 and 1995. For each matching regression standard errors are robust to heteroskedasticity, and we use the bias correction method of Imbens et al. (Stata Journal 2004).

Table 5 presents the estimated policy effects for six measures of competition and amenities, and for the three econometric models. Each numbered column corresponds to a market structure characteristic. Panel A contains the fixed effects regression based on the full sample with more than 5000 market-year observations. Panels B and C show the results of the long difference estimation on the cross section of neighborhoods. Panel B presents results from OLS regressions, while panel C presents the difference-in-difference matching regression results.

The results from the first two columns suggest that the policy led to more competitive local markets, and less spatial differentiation. In column (1), the full sample estimation shows a significant positive correlation of the policy with the number of competitors. The long difference estimates are also positive, but are imprecise. In column (2), the OLS estimates in panels A and B

command *nnmatch*.

show that the fraction of local monopolies decreased by around 7% after the policy; the matching estimate is qualitatively similar, but is imprecise.

The results from columns (3) and (4) show that the policy is associated with a decrease in the size of stations, both in terms of pumps and in terms of service islands. In terms of number of pumps, the estimated decrease is very robust across all specifications. The long-difference specifications suggest a decrease by more than two pumps per station on average, compared to about one pump in the first row. With the matching estimator we find a significant fall of about 2.6 pumps in the average size of stations. Recall that, on average, stations had 9.8 pumps in Québec after the introduction of the policy. These findings suggest that the changes in Québec (relative to the rest of Canada) were initiated before the policy, possibly because the policy was anticipated.

In columns (5) and (6) we estimate the effect of the policy on two amenities that we associate with the technological change: the presence of a convenience store and automated payment. In the full sample, the fraction of stations with these amenities increased at a significantly faster rate in the rest of Canada, suggesting that stations in Québec were less likely to invest in new technologies. The long-difference OLS estimates have the same signs, but are statistically insignificant.

Overall, the long-difference results, both OLS and matching, are qualitatively similar to the full-panel OLS results, but less precise, which should not be surprising given the much smaller number of observations, and the relatively larger number of controls used to estimate the propensity score functions.

Station-level regressions

We next examine the effect of the policy on the structure of local market structure, defined as a 2-minute driving distance buffer around each station. This leads to two regression specifications estimated at the station-level:

$$Y_{j,t} = \beta_F \text{Policy}_{j,t} + \mathbb{Z}_{j,t} \gamma_F + \mu_j^{\text{station}} + \mu_t^{\text{year}} + e_{j,t}, \quad (11)$$

$$\Delta Y_j = \beta_D \text{QC}_j + \mathbb{W}_j \gamma_D + u_j, \quad (12)$$

where $Y_{j,t}$ is the variable of interest for station j in period t , and $\Delta Y_j = Y_{j,2001} - Y_{j,1991}$ is its change between 1991 and 2001. The term $\text{Policy}_{j,t} = \text{QC}_j \times 1(t \geq 1997)$ is a policy indicator, μ_j^{station} and μ_t^{year} are station and year fixed-effects.

In Table 6 we show estimates of the effect of the policy on the minimum distance between stations, on the probability that an active station is a local monopoly (within two driving minutes), and on the average size of nearby competitors (within two driving minutes). Panel A uses the full year-stations panel data between 1991 and 2001 (13,228 observations), while panels B and C are first-difference regressions between 2001 and 1991 (922 stations). Column (3) is estimated on the sub-sample of stations with at least one competitor (i.e. 11,063 and 732 observations for panels A

Table 6: Local competitors' characteristics regressions

VARIABLES	(1) Minimum driving distance (min.)	(2) Monopoly 2 min. driving	(3) Avg. number of pumps 2 min. driving
A: Full-panel with FE (ols)			
Policy	-0.033 ^c (0.02)	-0.069 ^a (0.02)	-0.840 ^a (0.23)
Observations	13,228	13,228	11,063
B: Long-difference (ols)			
Policy	-0.039 ^c (0.02)	-0.045 ^c (0.02)	-1.446 ^a (0.33)
Observations	922	922	732
C: Long-difference (matching)			
Policy	-0.070 ^b (0.03)	-0.073 ^c (0.04)	-1.462 ^a (0.47)
Observations	922	922	732
Cross-sectional units: Stations			

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the station level. Each column corresponds to a different regression: (1) minimum driving distance between stations (minutes), (2) no competitor within 2 minutes, (3) average number of pumps within 2 minutes. Specifications in panel **A** include the following controls: postal-code area demographics (population, income and unemployment rate), city population and unemployment rate, number of companies at nearest terminal, rack price, indicator variables for Ultramar and Sunoco in 2 minutes neighborhoods (plus interaction with "after 1996" indicator), year and neighborhood fixed effects. The Policy coefficients correspond to the long-run effect of the price-floor: $\text{Year} \geq 1997 \times \text{Quebec}$. Panels **B** and **C** include the following control variables: postal-code area demographics in 1991, 1996 and 2001 (and changes), 1991 neighborhood structure (monopoly dummy, number of competitors, average number of pumps and islands), and Ultramar and Sunoco presence in 1991. The propensity score is estimated using a linear Logit specification with additional controls for the change in all outcomes between 1991 and 1995. For each matching regression standard errors are robust to heteroskedasticity, and we use the bias correction method of Imbens et al. (Stata Journal 2004).

and B/C respectively).

The policy effects estimated in Table 6 are statistically significant, and provide further evidence that the minimum price policy distorted the structure of local markets. The results imply that active stations in the rest of Canada became increasingly isolated compared to those in Québec. This effect is clearer when looking at the fraction of stations with no immediate neighbor within two minutes driving distance. All three specifications clearly show that this fraction decreased by about 8% because of the introduction of the policy. As a result, the average distance between stores also decreased in Québec relative to the rest of Canada.

The results regarding the size of stations are also consistent with the ones found previously using our alternative definition of competitive neighborhoods. The average number of pumps among nearby stations decreased by about one pump in Québec following the policy introduction.

The results obtained with our matching difference-in-difference estimator are significant and very similar to the ones obtained with OLS.

4.6 Evidence of the impact of the policy on station performance

As we discussed in Section 4.4, the reorganization that took place in the industry induced significant changes to the number and type of stations. In this section, we document how these changes also led to large productivity gains across all markets, and particularly in markets outside of Québec. Table 7 illustrates this point by measuring the growth in volume per station in the treated and control markets between 1991 and 2001. In the first three columns we calculate a growth rate over the 11 year period equal to 105% in the rest of Canada, compared to 58% in Québec. Much of the growth in productivity is associated with a 40% and 30% reduction in the number of stations in the rest of Canada and in Québec respectively, with most of the exiting stations being low-volume stations. Indeed, in the rest of Canada, stations that entered between 1992 and 2001 had sales' volumes that were on average 280% larger than those of stations that exited the market before 2001. The same statistic in Québec is 161%. Similarly, the growth rate of incumbent stations in Québec is equal to 32%, compared to 54% in Ontario.

In the last three columns, we decompose the growth in each market into the contribution of incumbents, entering and exiting stations respectively. To do so, we follow the strategy proposed by Foster, Haltiwanger, and Krizan (2001), and decompose the observed increase in the average volume per station into a within-station growth component, reallocation between incumbents, and relative size entering and exiting stations. In our context, the decomposition for region m is given by:

$$\begin{aligned} \Delta Q_m = & \sum_{s \in C_m} \frac{q_{s,2001} - q_{s,1991}}{N_{m,1991}} + \sum_{s \in C_m} \left(\frac{1}{N_{m,2001}} - \frac{1}{N_{m,1991}} \right) (q_{s,2001} - \bar{q}_{1991}) \\ & + \sum_{s \in C_m} \left(\frac{1}{N_{m,2001}} - \frac{1}{N_{m,1991}} \right) (q_{s,2001} - q_{s,1991}) \\ & + \sum_{s \in E_m} \frac{q_{s,2001} - \bar{q}_{1991}}{N_{m,2001}} - \sum_{s \in X_m} \frac{q_{s,1991} - \bar{q}_{1991}}{N_{m,1991}} \end{aligned} \quad (13)$$

where $q_{s,t}$ is the volume of station s in period t , \bar{q}_t is the average station-level volume in period t , C_m denotes continuing stations, E_m denotes entering stations, and X_m denotes exiting stations. In Table 7 we identify the contribution of incumbents as the sum of the first three terms (i.e. within and between firm growth terms, and the correlation term), and label the last two the contribution of firms entering with the newer technology and of "traditional" stations exiting the market. The result of this decomposition shows that the two groups of markets experienced a fairly similar transition process. In Québec 65% of the growth is due to incumbents' growth, and 35% to the net entry of more productive stations. In the rest of Canada, the contribution of incumbent is

Table 7: Decomposition of volume per station growth

	Volume per station			Growth decomposition		
	1991	2001	% Change	Incumbents	Entering	Exiting
Québec	3,775.58 (608)	5,981.27 (425)	58%	1,438.45 65%	360.16 16%	-407.08 -18%
Incumbent	4,430.36 (378)	5,850.59 (378)	32%			
Exiting/Entering	2,699.46 (230)	7,032.32 (47)	161%			
Rest of Canada	5,371.90 (551)	11,013.21 (328)	105%	3,947.83 70%	832.00 15%	-861.49 -15%
Incumbent	6,986.45 (294)	10,737.39 (294)	54%			
Exiting/Entering	3,524.90 (257)	13,398.21 (34)	280%			

The index measures the share-weighted sales volume per day per station in liters. The Industry figure corresponds to all stations in the region. Incumbent stations are stations that are present throughout the whole time span of the sample. Entering stations are stations that entered after 1991 and exiting stations are stations that exited before 2001. Numbers of stations in parentheses. The decomposition is done according to Foster, Haltiwanger, and Krizan (2001). We are interested in decomposing ΔV_t , where $V_t = \sum_{s \in I} \theta_{st} v_{st}$ is the index of industry sales volume, θ_{st} is the share of station s in industry I in period t and v_{st} is an index of station-level sales volume.

larger at 70%. The difference between the two is mostly due to exiting firms, which contributed a larger share of the growth in Québec (18% versus 15%). This is consistent with the fact that exiting stations in Québec tended to be relatively smaller on average.

In order to quantify the effect of the policy on the volume of sales and prices of individual stations, we estimate the effect of the policy using the full-panel and long-difference samples. In the first three columns of Table 8 we test whether or not the policy affected stations' sales volume, and therefore can explain the differential growth in sales volume per station observed between Québec and the rest of Canada. This corresponds to corollary C_1 : the distortions caused to market structure by the policy should indirectly lead to lower sales per station. In theory, if demand is inelastic, the policy should only affect the volume of stations through a reorganization effect; that is through changes in the characteristics of active stations, and in the degree of competition from neighboring stations.

However, in contrast to the market-structure regressions, here we are concerned about controlling for the short-run impact on the policy on sales volume and markups. This is because the year preceding the introduction of the policy exhibited severe price wars, and markups increased by more than 10% immediately following the policy introduction. We estimate the following two

econometric specifications:

$$Y_{j,t} = \beta_F^{lr} \text{Policy}_{j,t} 1(t \neq 1996, 1997) + \beta_F^{sr} \text{Policy}_{j,t} 1(t = 1996, 1997) \quad (14)$$

$$+ Z_{j,t} \gamma_F + \mu_j^{\text{station}} + \mu_t^{\text{year}} + e_{j,t},$$

$$\Delta Y_j = \beta_D \text{QC}_j + \mathbb{W}_j \gamma_D + u_j, \quad (15)$$

where $Y_{j,t}$ is the variable of interest for station j in period t , and $\Delta Y_j = Y_{j,2001} - Y_{j,1991}$ is its change between 1991 and 2001. The term $\text{Policy}_{j,t} = \text{QC}_j \times 1(t \geq 1997)$ is a policy indicator, μ_j^{station} and μ_t^{year} are station and year fixed-effects.

Panel A corresponds to linear difference-in-difference regression models estimated by OLS with alternative controls. The volume regressions exclude observations with missing entries (10,933 observations). To control for the fact that volume is not measured for about 25% of stations, all specifications in the table control for the selection probability (i.e. log and log squared). The selection probability is estimated with a Probit model, and controls for station characteristics (i.e. size, amenities, and brands), and year/market dummy variables.

In column (1) we condition only on “exogenous controls”: market fixed-effects, demographics characteristics, taxes, upstream market characteristics, and the presence of Sunoco and Ultramar in a two-minutes neighborhood. In this specification, we estimate that the policy led to a sizable decline in volume per station of nearly 1,500 liters per day. This estimate reflects both the selection and composition effect. That is, the coefficient reflects the fact that lower productivity stations remain active in Québec, and that incumbent stations exhibit different characteristics following the government intervention.

In column (2), we isolate the composition effect of the policy by conditioning on station fixed-effects. In this specification, the policy effect is identified solely from incumbent stations that remain active for a least one year in the pre and post policy period. The estimated effect decreases by about 500 liters, and implies that the composition effect alone corresponds to a 1000 liters decrease in the average volume per day of stations in Québec.

In column (3), we further control for characteristics of stations, and measures of the structure of neighborhoods and markets. The exact definitions of these “endogenous controls” are listed in the footnote of Table 8. We use the dependent variables from our market-structure regressions discussed in section 4.5 as right-hand-side variables, and use richer definitions of competitive neighborhoods by combining more than one distance radii.

The results confirm the theory prediction: after conditioning on characteristics of stations and local markets, stations in Québec did not decrease their productivity significantly following the policy introduction. Therefore, although stations in Québec became relatively less productive because of the policy, it is entirely due to a decrease in the average size of stations and an increase in the competitiveness of the markets, relative to the rest of Canada.

Table 8: Volume regressions

VARIABLES	(1) Volume	(2) Volume	(3) Volume
A: Full panel			
Policy (long run)	-1,474 ^a (359.6)	-1,012 ^a (276.7)	-220.1 (271.2)
Policy (short run)	844.1 ^a (225.0)	590.5 ^a (178.1)	238.7 (158.9)
Observations	10,933	10,933	10,933
B: Long difference			
Policy (OLS)	-875.3 ^b (374.7)		12.73 (365.0)
Policy (Matching)	-884.0 (722.9)		106.0 (610.6)
Observations	598		579
Control variables			
Demographic controls	✓	✓	✓
Station FE		✓	✓
Stations+Neighborhoods controls			✓
Market controls			✓

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the station level (for panel A). All specifications in panel A include the following demographic controls: postal-code area demographics (population, income and unemployment rate), city population and unemployment rate, number of companies at nearest terminal, rack price, indicator variables for Ultramar and Sunoco in 2 minutes neighborhoods (plus interaction with “after 1996” indicator), short-run policy variables ($(1996 \leq \text{Year} \leq 1997) \times \text{Quebec}$), year and station fixed effects. Specification (2) in Panel A adds controls for station characteristics (pumps, islands, convenience store type (3), carwash, repair shop, conventional type, self-service, major brand indicator), and neighboring station attributes using four neighborhood definitions: non-overlapping clusters (see text), driving distance radiuses (2 and 5 minutes), and common streets indicator. Competing stations’ attributes are: number of close competitors, number of pumps, number of service islands, fraction of major brands. Specification (3) in Panel A adds market-level (city) market structure variables: number of stations per capita, number of pumps and islands per capita, and the average number of pumps and islands per station (in log). Column (3) in Panel B controls for station and market characteristics by using as dependent variable the residual of two regressions of volume/markup on the demographics, stations, neighborhoods and city explanatory variables (i.e. same controls as column (3) in panel A). The propensity score in the matching regressions is estimated using a linear Logit specification with the following control variables: postal-code area demographics in 1991, 1996 and 2001 (and changes), 1991 neighborhood structure (monopoly dummy, number of competitors, average number of pumps and islands), Ultramar and Sunoco presence in 1991, and the change in volume and markup between 1991 and 1995. For each matching regression standard errors are robust to heteroscedasticity, and we use the bias correction method of Imbens et al. (Stata Journal 2004).

The short-run results in the second row of Table 8 also confirm our assumption that aggregate gasoline demand is inelastic. After controlling for stations and market characteristics, we estimate that volume per stations did not increase between 1996 and 1997 in Québec. In fact, sales volume per station increased slightly over this period, but not statistically significantly so, despite a nearly 10% increase in markups.

In Panel B, we estimate the long-run effect of the policy on sales volume using both OLS and the same nearest-neighbor matching estimator described above. In particular, for stations that were active in 1991 and 2001, we estimate the effect of the policy on the long-run change in volume. In the matching regression we compare stations in Québec with similar control stations, according to the propensity score. The long-difference estimates shown in column (1) imply that the policy had a negative effect of around 880 liters, ignoring the endogenous changes in market and station characteristics. The matching estimate is statistically insignificant but still similar to the OLS result.

In order to control for the fact that station and market characteristics changed over the same period, we obtain estimates with controls for these changes. In order to estimate the matching regression we project sales volume on our set of endogenous controls using the control sample observations. The matching difference-in-difference estimator is therefore constructed using the implied residual volumes, rather than the levels. This is an application of a regression-adjusted matching estimator as described by Heckman, Ichimura, and Todd (1997). The estimates shown in column (3) of panel B confirm the results obtained in Panel A: after conditioning on endogenous controls, we estimate that the policy did not lead to any significant change in the average volume per stations.

4.7 Evidence of the impact of the policy on prices

In light of the evidence presented so far on the effects of the policy on market structure, a naturally related question is: do these market structure effects outweigh the efficiency gains from the adoption of the new and superior technology in such a way as to have an impact on prices? In this section we attempt to answer this question by testing the effect of the policy on prices. Recall from Section 3 that our theoretical model's predictions were ambiguous as to this: if the floor is frequently binding or keeps high marginal-cost firms in the industry, then prices will increase; if the floor increases competition, then prices will fall.

Table 9 shows the estimates of the same econometric specifications used in the volume regressions. Since gasoline prices do not exhibit much dispersion within cities, we adjust the method used to calculate standard-errors. Rather than clustering observations at the station-level, we calculate standard-errors that are clustered at the city-year level.

The long-run results indicate large price declines of 3.0 cents in Québec following the introduction of the policy. This effect falls somewhat when controlling for fixed station characteristics (i.e. -2.8 cents), but it is still significantly different from zero. In column (3), we show that after conditioning on the changes in the structure of gasoline markets, the policy did not lead to a significant average price decrease. Therefore, in the full panel, controlling for the endogenous characteristics of stations and markets reduces the estimated policy effect almost entirely.

The long-difference estimates presented in Panel B of Table 9 yield more ambiguous results.

Table 9: Price regressions

VARIABLES	(1) Price	(2) Price	(3) Price
A: Full panel			
Policy (long run)	-2.996 ^a (1.116)	-2.803 ^b (1.200)	-0.783 (1.500)
Policy (short run)	3.756 ^a (1.273)	3.632 ^a (1.364)	3.057 ^b (1.351)
Observations	13,227	13,227	13,227
B: Long difference			
Policy (OLS)	-1.259 ^a (-0.206)		-1.586 ^a (-0.188)
Policy (matching)	-3.873 ^a (-0.932)		-0.948 ^c (-0.516)
Observations	922		922
Control Variables			
Demographic controls	✓	✓	✓
Station FE		✓	✓
Stations+Neighborhoods controls		✓	✓
Market controls			✓

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the city-year level (for panel A). All specifications in panel A include the following demographic controls: postal-code area demographics (population, income and unemployment rate), city population and unemployment rate, number of companies at nearest terminal, rack price, indicator variables for Ultramar and Sunoco in 2 minutes neighborhoods (plus interaction with “after 1996” indicator), short-run policy variables ($(1996 \leq \text{Year} \leq 1997) \times \text{Quebec}$), year and station fixed effects. Specification (2) in Panel A adds controls for station characteristics (pumps, islands, convenience store type (3), carwash, repair shop, conventional type, self-service, major brand indicator), and neighboring station attributes using four neighborhood definitions: non-overlapping clusters (see text), driving distance radiuses (2 and 5 minutes), and common streets indicator. Competing stations’ attributes are: number of close competitors, number of pumps, number of service islands, fraction of major brands. Specification (3) in Panel A adds market-level (city) market structure variables: number of stations per capita, number of pumps and islands per capita, and the average number of pumps and islands per station (in log). Column (3) in Panel B controls for station and market characteristics by using as dependent variable the residual of two regressions of volume/markup on the demographics, stations, neighborhoods and city explanatory variables (i.e. same controls as column (3) in panel A). The propensity score in the matching regressions is estimated using a linear Logit specification with the following control variables: postal-code area demographics in 1991, 1996 and 2001 (and changes), 1991 neighborhood structure (monopoly dummy, number of competitors, average number of pumps and islands), Ultramar and Sunoco presence in 1991, and the change in volume and markup between 1991 and 1995. For each matching regression standard errors are robust to heteroscedasticity, and we use the bias correction method of Imbens et al. (Stata Journal 2004).

The matching results are consistent with the full-panel results, and imply that once we control for endogenous characteristics the policy effect drops significantly. In contrast, the OLS results suggest that the changes in market structure and station characteristics failed to fully explain the long-run decline in prices among stations that were active between 1991 and 2001. We find that, not only do the effects not disappear after controlling for the endogenous changes in station and

market characteristics. This means that the market structure changes that we estimate in this specification are not sufficiently pronounced to explain the stark changes between 1991 and 2001. In the full-panel specifications changes were distributed over the entire period, and the slower reorganization in Québec can explain most of the observed decline in prices.

This lack of robustness across specifications is likely due to small amount of within-market price variation. Since neighboring stations often post the same price, we are left with very little cross-sectional variation to infer something about the effect of the policy on prices. Unfortunately, this is the nature of competition in gasoline markets. The lack of cross-sectional variation makes the results sensitive to the influence of market-level outlier observations. This likely explains the differences among the two long-difference specifications, since the estimation relies on long-run changes in prices among the cities (rather than 11 observations per city as in the full panel).

Our conclusions with respect to the impact of the policy on prices are also weak because of the frequency of the price data. We observe station-level prices only once a year, and we are forced to ignore the important week-to-week price variation. The existence of predictable asymmetric cycles akin to Edgeworth cycles (Edgeworth (1925) in which price increases are both simultaneous and large (relenting phase), and are followed by a sequence of small decreases (undercutting/matching phase) has been documented by Castanias and Johnson (1993) in U.S. markets, and by Eckert (2002) and Noel (2007) in Canada. If cycles were shorter in Québec because of the floor, then it would be more likely that a price at the bottom of the cycle would be sampled in Québec than in other regions. Even though we estimate effects using average prices over a period of three years, this could potentially explain the lower long-run prices we find in Québec.

To circumvent the lack of week-to-week variation in our main data set we collected weekly city-level average prices between 1991 and 2001 from MJ Ervin. Over this time period, MJ Ervin collected prices in 13 Canadian cities.²⁸ All specifications use Montreal and Québec City as treatment cities, since these are the only Québec cities in the data set. We consider three different sets of control cities. The first uses Toronto, Regina and Halifax. This choice is motivated by our selection of markets in the Kent data. In the Kent data we have information on cities in four provinces: Québec, Ontario, Nova Scotia, and Saskatchewan. We selected all available cities in the MJ data that were in these provinces, with one exception: Ottawa. We drop Ottawa, since it is linked to the Québec markets of Hull and Gatineau. The second set uses all 11 of the control cities. Finally, the third set uses all of these except for Ottawa (for the reasons explained above) and Vancouver, since it experienced an abnormally long price war in the late 1990's triggered by the entry of a big-box retailer, and since self-service stations were banned or restricted in at least one suburb of the city.²⁹

²⁸The cities are Vancouver, Edmonton, Calgary, Winnipeg, Regina, Toronto, Ottawa, Halifax, Saint John, St Johns, Charlottetown, Montreal and Québec City.

²⁹See The Conference Board (2001) and <http://www.coquitlam.ca/Libraries/Zoning.Bylaw/Part.17.-Service.Station.Zones.sflb.ashx>.

Table 10: Price difference-in-difference regressions using weekly price data

VARIABLES	(1) (t \geq 1991)	(2) (t \geq 1992)	(3) (t \geq 1993)	(4) (t \geq 1994)	(5) (t \geq 1995)
Controls: Toronto, Halifax, Regina					
Policy (LR)	-3.273 ^a (0.830)	-2.428 ^a (0.638)	-2.177 ^a (0.529)	-1.552 ^a (0.486)	-1.489 ^a (0.514)
Policy (SR)	-1.814 ^b (0.785)	-0.972 ^c (0.527)	-0.723 (0.449)	-0.0969 (0.410)	-0.0333 (0.451)
Observations	2,765	2,500	2,250	1,995	1,780
Controls: All cities					
Policy (LR)	-1.929 ^b (0.753)	-1.093 ^c (0.620)	-0.786 (0.630)	-0.206 (0.596)	-0.114 (0.666)
Policy (SR)	-1.778 ^a (0.630)	-0.943 ^b (0.450)	-0.637 (0.472)	-0.0574 (0.446)	0.0343 (0.557)
Observations	6,636	6,000	5,400	4,788	4,272
Controls: All cities except Ottawa and Vancouver					
Policy (LR)	-2.862 ^a (0.749)	-1.979 ^a (0.609)	-1.656 ^a (0.613)	-1.021 ^c (0.565)	-0.928 (0.622)
Policy (SR)	-2.230 ^a (0.637)	-1.351 ^a (0.451)	-1.030 ^b (0.467)	-0.395 (0.414)	-0.303 (0.499)
Observations	5,530	5,000	4,500	3,990	3,560

Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Standard-errors are clustered at the city-year level. Results correspond to OLS estimates of the effect of the policy on weekly prices. Exogenous controls in all regressions include the rack price, and city and year fixed effects. In columns (2) through (5) we test the robustness of our results to the variation of the start date of the “before” period in our analysis.

We also test the robustness of these results to the variation of the start date of the ‘before’ period in our analysis. More precisely, using the weekly MJ Ervin data we re-estimate the price regressions successively dropping the earliest years of data (first dropping 1991 and estimating using 1992-2001, then dropping 1992 and estimating with 1993-2001, and so on). These results are also presented in Table 10

To summarize, using our main data set, we find that two of three specifications suggest that the majority of the price drop in Québec was due to endogenous market structure changes. The weekly results provide mixed evidence that prices fell unconditionally of these changes. Recall that the effect of the floor on price operates through three channels: (i) an increase in competition (negative), (ii) a drop in efficiency and productivity (positive), and (iii) a positive probability of a binding floor (positive). While our findings may not provide robust results, we believe that they credibly establish that the two positive effects (efficiency and binding constraint) *do not* dominate the negative force (competition). Every specification that we have considered yields a zero or negative effect of the regulation on prices, and the point estimate is *never* positive. Given that price-floor regulations are typically viewed as exerting upward pressure on prices, and that it is reasonable to expect the efficiency effect to be important this result is, in itself, interesting. Our

price results are also consistent with those of Skidmore, Peltier, and Alm (2005), who find that prices tend to fall after the adoption of sales-below-cost laws in US gasoline markets.

5 Conclusion

We have shown that the price floor regulation established in the Québec retail gasoline market has had a substantial effect on market structure. We also showed that the policy caused distortions to the sales of gasoline at individual stations. We identified these effects comparing the long-run behavior (before and after 1997) in the local gasoline markets in Québec and other provinces where the policy was never implemented.

We find that the policy significantly affected the reorganization of the markets. In Québec, there were more stations after the policy was introduced compared to the rest of Canada, after controlling for unobserved market and time-specific effects. Moreover, stations outside Québec became bigger and offered a wider variety of products. After the policy was introduced, Québec stations became relatively more homogeneous in terms of the type of services that they offered, mostly because new stations entering in the rest of Canada were very different from the stations that stayed in the market. Moreover, there is evidence that the policy caused a long-run decrease in station-level sales, but that despite the observed increase in efficiency, prices did not rise.

These results are consistent with our interpretation of the effect of the price floor on market structure. As we have shown in our theoretical model, even when it does not bind, the policy can block the entry of more efficient firms that must incur larger operating costs and increase competition. Sales and prices can therefore both be lower.

Moreover, our interpretation is in line with reports from impartial sources critical of Québec's price floor regulation that also suggest that the implementation of the floor slowed down the rationalization relative to other North-American markets. In a 2002 report aimed at opposing a proposed increase in the price floor, the Canadian Automobile Association (CAA) argued that between 1997 and 2002 the price floor policy generated a productivity gap between stations in Québec and those in Ontario and the US. This is evidence that industry analysts consider the floor as the leading cause of Ontario's faster rationalization.³⁰

We have mentioned that there are other factors not considered in our static model that may influence the reorganization of markets. In particular, there are other possible means by which the price floor could negatively affect the profits that stations expect to earn upon entry in a repeated game framework. For instance, it may be that the presence of a price floor makes it difficult for firms engaged in tacit collusion (as in Porter (1983)) to revert to a "punishing" stage. By limiting

³⁰The CAA is a consumer protection group present throughout all of Canada. In 1997, this organization was in favor of the price floor regulation as a way to reduce the frequency of price wars. Their reports are available on the Board's website: <http://www.regie-energie.qc.ca/audiences/3499-02/main.DocDepotAudience3499.htm> and <http://www.regie-energie.qc.ca/audiences/3499-02/main.Prew3499.htm>

the extent to which firms can punish defectors, Québec's price floor may restrict the severity of price wars. In doing so it may make pricing strategies less stable and make it increasingly difficult to sustain this type of equilibrium. The expected reduction in profit may deter the entry of new firms, as we observe in the data.³¹

However, the floor may actually serve as a facilitating device. It clearly provides a focal price to coordinate price changes, and can facilitate communication because it permits firms to sue their competitors if they charge low prices. Indeed, collusion has taken place in Québec since the arrival of the floor. Stations in four cities in Québec were charged with price fixing in 2006.³² However, collusion in Sherbrooke does not seem to have started until after our sample period as the market experienced severe price wars in 1998 and 1999. The price wars are consistent with the market being in excess capacity, and evidently help to explain some of our markup results. It may also facilitate collusion when firms are asymmetric in terms of costs. The floor may allow high-cost firms to punish low-cost firms when this would not be possible otherwise. The existence of the floor means that punishment is at the floor rather than at some even lower price. So if the floor is sufficiently low that reverting to it represents a punishment, but is high enough that it is above the marginal cost of the high-cost firm, then it may in fact facilitate collusion.

There might also be additional effects of Ultramar's low-price guarantee, beyond the effects that we control for. In our empirical analysis we control for Ultramar's low-price guarantee within local markets, but it is possible that it had global effects that were not picked up by the controls. Specifically, the concern is that the low-price guarantee, if credible, might in fact be responsible for the entry distortion that we observe. New firms may be reluctant to enter the market with the "large" technology – low variable, but high fixed costs – since this technology demands that firms set lower-prices than their competitors in order to grab market share. Ultramar's low-price guarantee might prevent entrants from selling enough volume to cover their fixed costs.

However, since this policy is essentially a price-matching guarantee, it may not actually deter entry. Arbatskaya (2001) shows that with price-matching guarantees an incumbent cannot deter entry into the market.³³ Furthermore, Ultramar's commitment to the low-price guarantee may not be credible. Therefore, if new firms actually enter with the "large" technology, Ultramar will retreat from its guarantee. That is, should a sufficient number of "large"-type entrants actually enter the market, it may not actually be optimal for Ultramar to stick with its low-price guarantee.

What does our analysis say about welfare? The evidence presented here suggests that the policy decreased prices at the pump for consumers. On the other hand, after the policy, stations in

³¹See Gagné, Van Norden, and Versaevél (2003) for a study of punishment under the Québec gasoline price regulation.

³²See Clark and Houde (2013a) and Clark and Houde (2013b) for detailed analyses of these cartels

³³Entry occurs in any subgame perfect equilibrium of the sequential move entry game and the incumbent is accommodating. However, the price-matching guarantee is shown to be valuable for the incumbent as an incentives management device. In any subgame perfect equilibrium the firms share the market equally and the price-matching guarantee serves to facilitate collusion.

Québec became increasingly different from stations in the other cities in the sample. Therefore, after the policy, consumers in the control markets were paying for different services than consumers in Québec, in the sense that new stations in the controls were generally bigger and offered a wider variety of products.

On the firms' side, we find that after the policy stations in Québec were charging lower prices and had lower sales. This increase in net revenue for stations outside Québec may reflect an increase in rents due to decreased competition. On the other hand, the higher prices may just reflect higher fixed or entry costs which reflect the expanded services they provide. From an environmental perspective, if the policy kept older stations in the market, an indirect consequence is increased risk of hazardous tank leakage. Evaluating all these welfare effects requires the structural estimation of a precise market model, which is something we leave for future research.

Notice finally that our results are specific to Québec gasoline markets during a particular time period. Even though extrapolation of our results to other markets is not possible, our findings provide evidence that such effects might be present in *any* market. As discussed in the Introduction, similar price regulations are not uncommon. For example, agricultural price controls that provide insurance to local producers against excessively low prices are a form of price floor. In Europe, regulation aimed at protecting small retailers from the aggressive pricing of big retailers is common. And throughout the world, anti-dumping regulation aimed at protecting local producers by forbidding foreign firms from setting price below average variable costs are a subtle form of price floor. All these policies have longer run effects on market structure and performance that are not always fully recognized.

References

- Abadie, A. and G. W. Imbens (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica* 74(1), 235–267.
- Anderson, R. and R. Johnson (1999). Antitrust and sales-below-cost laws: The case of retail gasoline. *Review of Industrial Organization* 14(3), 189–204.
- Arbatskaya, M. (2001). Can low-price guarantees deter entry? *International Journal of Industrial Organization* 19(9), 1387–1406.
- Armstrong, M., J. Vickers, and J. Zhou (2009). Consumer protection and the incentive to become informed. *Journal of the European Economic Association* 7(2-3), 399–410.
- Asker, J. and H. Bar-Isaak (2010). Exclusionary minimum resale price maintenance.
- Biscourp, P., X. Boutin, and T. Vergé (2008, May). The effects of retail regulations on prices: Evidence from the loi galland. CREST working paper.

- Brown, J., J. Hastings, E. T. Mansur, and S. B. Villas-Boas (2008). Reformulating competition? gasoline content regulation and wholesale gasoline prices. *Journal of Environmental economics and management*.
- Busse, M. and N. O. Keohane (2009). Market effects of environmental regulation: Coal, railroads, and the 1990 clean air act. forthcoming, *Rand Journal of Economics*.
- Campbell, J. R. and T. N. Hubbard (2010). The economics of of 'radiator springs:' industry dynamics, sunk costs, and spatial demand shifts.
- Castanias, R. and H. Johnson (1993). Gas wars: Retail gasoline price fluctuations. *Review of Economics and Statistics*.
- Clark, R. (2007). Advertising restrictions and competition in the children's breakfast cereal market. *Journal of Law and Economics* 50(4), 757–780.
- Clark, R. and J.-F. Houde (2013a). Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case. forthcoming *American Economic Journal: Microeconomics*.
- Clark, R. and J. F. Houde (2013b). The effect of explicit communication on pricing: Evidence from the collapse of a gasoline cartel. Forthcoming *Journal of Industrial Economics*.
- Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty*. Princeton University Press.
- Eckert, A. (2002). Retail price cycles and response asymmetry. *Canadian Journal of Economics*.
- Eckert, A. and D. S. West (2005). Rationalization of the retail gasoline station networks in canada. *Review of Industrial Organization* 26, 1–25.
- Eckert, H. and A. Eckert (2008). Environmental regulation and rationalization in the retail gasoline industry. Working paper, University of Alberta.
- Edgeworth, F. Y. (1925). The pure theory of monopoly. *Papers Relating to Political Economy*.
- Fenili, R. and W. Lane (1985). Thou shalt not cut prices. *Regulation* 9(5), 31–35.
- Foster, L., J. Haltiwanger, and C. Krizan (2001). Aggregate productivity growth: Lessons from microeconomic evidence. In *New Developments in Productivity Analysis*. University of Chicago Press.
- Foster, L., J. Haltiwanger, and C. J. Krizan (2006, November). Market selection, reallocation, and restructuring in the u.s. retail trade sector in the 1990s. *The Review of Economic and Statistics*.
- Gagné, R., S. Van Norden, and B. Versaevel (2003, September). Testing optimal punishment mechanisms under price regulation: The case of the retail market for gasoline. working paper, HEC Montréal.
- Heckman, J. J., H. Ichimura, and P. E. Todd (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The Review of Economic Studies* 64(4), 605–654.

- Holmes, T. (2010). The diffusion of wal-mart and economies of density. forthcoming, *Econometrica*.
- Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review* 102(5), 2147–82.
- Jia, P. (2008, November). What happens when wal-mart comes to town: An empirical analysis of the discount industry. *Econometrica*.
- Johnson, R. (1999). The impact of sales-below-cost laws on the u.s. retail gasoline market. Report Prepared for Industry Canada, Competition Bureau.
- Johnson, R. N. and C. J. Romeo (2000). The impact of self-service bans in the retail gasoline market. *The Review of Economic and Statistics* 82(4), 625–633.
- Milyo, J. and J. Waldfogel (1999). The effect of price advertising on prices: Evidence in the wake of 44 liquormart. *American Economic Review* 89(5), 1081–1096.
- Noel, M. (2007). Edgeworth price cycles, cost-based pricing and sticky pricing in retail. *Review of Economics and Statistics*.
- Porter, R. H. (1983, Autumn). A study of cartel stability: The joint executive committee, 1880-1886. *The Bell Journal of Economics* 14(2), 301–314.
- Ryan, S. P. (2006). The costs of environmental regulation in a concentrated industry. Working paper, MIT.
- Skidmore, M., J. Peltier, and J. Alm (2005). Do state motor fuel sales-below-cost laws lower prices? *Journal of Urban Economics* 57(1), 189–211.
- Tan (2006). The effects of taxes and advertising restrictions on the market structure of the u.s. cigarette market. *Review of Industrial Organization* 28(3), 231 – 251.
- The Conference Board (2001). *The final fifteen feet of hose: The Canadian gasoline industry in the year 2000*. Ottawa: Conference Board of Canada.
- Yin, H., H. Kunreuther, and M. White (2007). Do environmental regulations cause firms to exit the market? evidence from underground storage tank (ust) regulations. Risk Management and Decision Processes Center working paper # 2007-10-17.

A Additional tables

Table 11: Summary statistics of the key variables for markets in Québec and the rest of Canada before and after the policy change

VARIABLES	Before 1997				After 1997				Diff-in-Diff %
	Quebec		Rest		Quebec		Rest		
	N	Avg. (sd)	N	Avg. (sd)	N	Avg. (sd)	N	Avg. (sd)	
Price	3873	29.48 (3.082)	3926	27.62 (3.964)	2739	34.12 (6.815)	2689	35.72 (7.934)	-0.118 ^a
Markup	3873	0.27 (0.114)	3926	0.19 (0.120)	2739	0.17 (0.103)	2689	0.17 (0.073)	-0.290 ^a
Sales vol (x1000 lt/day)	3495	4.29 (2.971)	3130	6.70 (4.535)	2447	5.73 (4.043)	1979	10.12 (6.792)	-0.483 ^a
Nb. of pumps	3873	8.14 (5.723)	3926	9.80 (5.914)	2739	9.85 (7.307)	2690	13.73 (9.739)	-0.270 ^a
Nb. Islands	3873	2.11 (1.268)	3926	2.42 (1.275)	2739	2.26 (1.341)	2690	2.77 (1.453)	-0.085 ^a
Islands > 4	3873	0.17 (0.380)	3926	0.20 (0.396)	2739	0.22 (0.412)	2690	0.30 (0.457)	-0.334 ^a
No Conv. store	3873	0.58 (0.494)	3926	0.59 (0.492)	2739	0.40 (0.490)	2690	0.36 (0.480)	0.079 ^a
Carwash	3873	0.19 (0.392)	3926	0.18 (0.383)	2739	0.19 (0.390)	2690	0.22 (0.413)	-0.221 ^a
Pay at the pump	3873	0.00 (0.000)	3926	0.02 (0.124)	2739	0.00 (0.019)	2690	0.02 (0.140)	0
Repair shop	3873	0.19 (0.393)	3926	0.08 (0.265)	2739	0.16 (0.362)	2690	0.07 (0.247)	-0.142 ^b
Self service	3873	0.37 (0.482)	3926	0.31 (0.463)	2739	0.50 (0.500)	2690	0.40 (0.491)	0.131 ^a
Local comp.	3873	4.30 (0.459)	3926	4.18 (1.291)	2739	4.03 (0.893)	2690	3.43 (1.332)	0.113 ^a
Street comp.	3873	10.20 (9.423)	3926	13.36 (15.677)	2739	7.43 (6.245)	2690	11.58 (13.575)	-0.090 ^a

The policy effects by category are estimated with difference-in-difference OLS regressions of each variable among the stations in each category, across regions before and after the policy. Significance level: ^a 0.01, ^b 0.05, ^c 0.1. Reconfiguring stations are stations that show a change in the number of pumps larger than 4 at any point throughout the sample. In the top panel, entrants are defined as stations that enter and never exit.

B Numerical example of model

In this appendix we provide a numerical example in support of the claims made in the model of entry and price competition presented above. That is, we show for a given demand characterization and a given set of parameter values, that the outcomes described in the text can arise.

To show that our results are not driven by the asymmetric strategy sets of the two firms we assume here that both firms have access to both the s and the l technologies. Therefore, sets of available actions for both types of firms are summarized by $\{A_1, A_2\} = \{(0, s, l), (0, s, l)\}$.³⁴

To characterize demand at each store, we use a logit model, and so when both firms are active, demand for store $j \in \{1, 2\}$ is given by:

$$D_j(p_1, p_2|\omega) = M \frac{e^{(\delta_j - p_j)}}{e^\nu + \sum_k e^{(\delta_k - p_k)}}, \quad (16)$$

where k is the set of firms that chose action s or l . Demand for firms choosing action 0 is assumed to be zero.

Without a price floor constraint, a symmetric price equilibrium is described by two prices $p = \{p_1(\omega), p_2(\omega)\}$ solving the following FOCs:

$$D_j(p_1, p_2|\omega) - M(p_j(\omega) - c_j) s_j (1 - s_j) = 0, \quad j \in \{s, l\} \quad (17)$$

where $s_j = \frac{e^{(\delta_j - p_j)}}{e^\nu + \sum_k e^{(\delta_k - p_k)}}$. When firms are constrained by a price floor $p_f > 0$ the equilibrium is characterized by four Kuhn-Tucker conditions:

$$D_j(p) - M(p_j(\omega) - c_j) s_j (1 - s_j) + \lambda_j = 0 \quad (18)$$

$$\lambda_j (p_j(\omega) - p_f) = 0 \quad (19)$$

for each $j \in \{1, 2\}$ and $\lambda_j \geq 0$.

In what follows we set $M = 1000$, normalize the value of the outside good (ν) to be zero, and set $\delta_j = 1$ for $j = \{1, 2\}$. We also assume that $c_s = 0.2$ and that $c_l = 0$.

With these parameter values, the unconstrained game is as in Table 12.

Table 12: Unconstrained game

Player		2		
		0	s	1
1	0	0.0, 0.0	0.0, 497.7 - F_s	0.0000, 567.1 - F_l
	s	497.7 - F_s , 0.0	362.9 - F_s , 362.9 - F_s	348.4 - F_s , 417.5 - F_l
	l	567.1 - F_l , 0.0	417.5 - F_l , 348.4 - F_s	401.1 - F_l , 401.1 - F_l

³⁴If we think of one of the firms as being the incumbent, then the large technology would represent its reconfiguration.

B.1 Excessive crowding

In the excessive crowding case we assume that the price floor is high and binds even in the (s, s) market structure. So we assume that $p_f > 1.563$ which is the Nash equilibrium price in the (s, s) outcome of the unconstrained game. We set $F_s = 360$ and we suppose that $F_l = 413$. Then the constrained and unconstrained games can be seen in Table 13.

Table 13: Excessive Crowding

Unconstrained game				
Player		2		
		0	s	l
1	0	0.0, 0.0	0.0, 137.7	0.0, 153.8
	s	137.7, 0.0	2.9, 2.9	-11.6, 3.0
	l	153.8, 0.0	3.0, -11.6	-12.2, -12.2

Constrained game				
Player		1		
		0	s	l
1	0	0.0, 0.0	0.0, 137.7	0.0, 153.8
	s	137.7, 0.0	3.1*, 3.1*	3.1*, 3.0*
	l	153.8, 0.0	3.0*, 3.1*	3.0*, 3.0*

$p_f = 1.565$, $F_s = 360$ and $F_l = 413$.

The equilibria of the unconstrained game (shown in bold) are $(0, l)$ and $(l, 0)$, while the equilibrium of the constrained game is (s, s) . The price floor is binding for both the s - and the l -type in all of the oligopoly outcomes (as indicated by the $*$'s). Price is 1.567 (the l -type monopoly price) in the unconstrained game, and 1.565 (the value of the floor) in the constrained game.

B.2 Blockaded entry (High floor)

Again we assume that $p_f > 1.563$ which is the Nash equilibrium price of the unconstrained game. We set $F_s = 335$ and we suppose that $F_l = 388.5$. The constrained and unconstrained games can be seen in Table 14.

The equilibria of the unconstrained game are (s, l) and (l, s) , while the equilibrium of the constrained game is (s, s) . Prices are 1.548 and 1.417 (the Nash prices for the s - and l -types respectively in the (s, l) and (l, s) markets) in the unconstrained game, and 1.565 (the value of the floor) in the constrained game.

Table 14: Blockaded entry

Unconstrained game				
Player	1			
	0	s	l	
1	0	0.0, 0.0	0.0, 162.7	0.0, 178.6
	s	162.7, 0.0	27.9, 27.9	13.4, 29.0
	l	178.6, 0.0	29.0, 13.4	12.6, 12.6

Constrained game				
Player	1			
	0	s	l	
1	0	0.0, 0.0	0.0, 162.7	0.0, 178.6
	s	162.7, 0.0	28.1*, 28.1*	28.1*, 27.8*
	l	178.6, 0.0	27.8*, 28.1*	27.8*, 27.8*

$p_f = 1.565$, $F_s = 335$ and $F_l = 388.5$.

B.3 Blockaded entry (Low floor)

We assume that the floor is set to a lower level so that it is only binding for the l -type in the (s, l) (or (l, s)) equilibrium. Specifically, we set the floor to be $1.417 < p_f < 1.548$, where the bounds are given by the unconstrained equilibrium prices in the (s, l) outcome for the s - and l -types respectively. We set $F_s = 350$ and $F_l = 404.5$. The unconstrained and constrained games can be seen in Table 15.

Table 15: Blockaded entry

Unconstrained game				
Player	1			
	0	s	l	
1	0	0.0, 0.0	0.0, 147.7	0.0, 162.6
	s	147.7, 0.0	12.9, 0.0129	-1.6, 13.0
	l	162.6, 0.0	13.0, -0.0016	-3.4, -3.4

Constrained game				
Player	1			
	0	s	l	
1	0	0.0, 0.0	0.0, 147.7	0.0, 162.6
	s	147.7, 0.0	12.9, 12.9	10.5, 12.3*
	l	162.6, 0.0	12.3*, 10.5	-44.1*, -44.1*

$p_f = 1.565$, $F_s = 350$ and $F_l = 404.5$.

The equilibria of the unconstrained game are $(l, 0)$ and $(0, l)$, while the equilibrium of the constrained game is (s, s) . Price is 1.567 (the l -type monopoly price) in the unconstrained game, and 1.563 (the Nash price in the (s, s) market) in the constrained game. Note that in this case the price floor is only binding for the l -type, and so not in equilibrium.

C Description of the clustering algorithm

In this appendix we describe the construction of neighborhoods. Consider an isolated metropolitan area composed of L potential store locations. In the data we define n as the set of geographic coordinates and street intersection pairs that were ever occupied by a gasoline station between 1991 and 2001. Because of entry and exit, n is thus larger than the total number of active station at any point in time.

The clustering algorithm proceeds iteratively by grouping stations with similar spatial characteristics until the allocation of stores in groups is stable. We define the degree of similarity between two locations using the euclidian distance (d_{ij} expressed in Km), and an indicator variable equal to one if they share at least one street. Each location can be characterized by up to two streets.

The key parameter of the algorithm is δ . It determines a threshold distance such that two locations are considered in the same local market even if they do not have street in common. This parameter is important since two stores can be very close in euclidian distance, but the survey company does not locate them along the same street. Intuitively this parameter is a penalty added to the euclidian distance between two stores that are not connected by a common street. We fix the value of δ to 1/4 Km, which is a very small distance. Note that the number of stable clusters is rapidly decreasing in δ .

We initiate the algorithm by defining initial neighborhoods, as the set of possible street intersections in the city. Let \mathcal{M}^t be the allocation at iteration t . \mathcal{M}^t is a mapping from locations to neighborhoods:

$$\mathcal{M}^t = \{m_1^t, m_2^t, \dots, m_L^t\}, \quad (20)$$

where m_i^t is the neighborhood id associated with location i .

1. At iteration t , update the assignment of store $i \in L$:

(a) Calculate the distance between location i and the center of neighborhood m_i denoted by l_{m_i} ³⁵:

$$D(i, m_i) = \begin{cases} \delta & \text{if } |m_i| = 1, \\ d(l_i, l_{m_i}) & \text{otherwise.} \end{cases} \quad (21)$$

³⁵We definite l_m as the average latitude/longitude coordinate of locations belonging to m .

- (b) Calculate the distance from l_i to all other neighborhoods. For each neighborhood $m \neq m_i$:

$$D(i, m) = \begin{cases} d(l_i, l_m) + \delta & \text{if } s_i \notin S_m, \\ d(k_i, l_m) & \text{otherwise,} \end{cases} \quad (22)$$

where s_i is the vector of street indices of location i and $S_m = \bigcup_{j \in m} s_j$ is the union of streets for all locations in market m . Let $m_i^* \neq m_i^t$ the closest neighborhood for location i .

- (c) If $D_i^* < D(i, l_{m_i^t})$ set $m_i^{t+1} = m_i^*$. Otherwise leave location i assignment unchanged.
2. Repeat the previous steps for all $i \in L$.
 3. If $\mathcal{M}^{t+1} \neq \mathcal{M}^t$ repeat step (1) and (2). Otherwise stop.